

SCALABLE AND SECURE IPv6 SOLUTIONS FOR CONNECTING MOBILE NETWORKS TO THE GIG

Işıl Sebüktekin* and Anthony McAuley
Telcordia Technologies Inc.,
Piscataway, NJ 08854

ABSTRACT¹

Future military programs are mandated to use IPv6; however, little emphasis has been placed on exploiting the potential in IPv6 to more efficiently support mobile networks. Current approaches mimic IPv4 solutions, which may prevent the full benefits of IPv6 from being realized in dynamic networks. These IPv4-copypat solutions may, for example, degrade routing performance and scalability. In this paper we analyze the alternatives available within IPv6 to improve the interconnection of mobile user networks with the GIG, while addressing the stringent application and security requirements of future military networks. The benefits apply to both the mobile network and its more stable transit backbones. We show that much better scalability, performance and autonomy can be achieved in supporting mobile user networks. No new protocols are required; only exploiting advanced IPv6 features such as autoconfiguration, large address space, address summarization for routing, and advanced mobility support. We also show that our proposed solutions can be made compatible with military security needs, such as the use of HAIPE.

1. INTRODUCTION

Over 10 years of IPv6 protocol development, testing, and standardization have lead to mature IPv6 standards. IPv6 offers a massive increase in the number of addresses, not only allowing all nodes to have their own globally routable addresses, but to enable much simpler address administration. IPv6 offers improvements in many important networking functions, notably in mobility, autoconfiguration, quality of service, multicasting, and security. IPv6 is now mandated for use in future military

networks and is increasingly used in commercial networks.

As commercial and military networks move towards IPv6, many of the approaches are naturally mimicking IPv4 solutions. Although this has advantages in terms of gradual IPv6 knowledge build up, and ease of transition, it does not take full advantage of the IPv6 enhancements over IPv4. Other than the large address space and packet header conversions, little emphasis has been placed on exploiting the potential in IPv6 to provide more robust, more manageable and more efficient solutions.

The driving need to move to more novel IPv6-type solutions does not exist for most commercial applications. Even cellular and wireless access networks can perform well with existing IPv4-copypat solutions. In the commercial space, mobility is in general confined to single-node events and is supported by high-bandwidth low-latency wireline backbone networks. For example, networks can still use Dynamic Host Configuration Protocol (DHCP) for the dynamic allocation of a single address to each host and node mobility can be handled by Mobile IP.

The premise of this paper, however, is that future military networks would be significantly handicapped by using only IPv4-copypat solutions. In particular, the many mobile user networks (large and small), envisioned for the future battlefield networks could have much greater performance, efficiency and flexibility in their use of the backbone by leveraging the advance features of IPv6. The benefits are especially notable for large multi-homed networks (e.g., WIN-T), with multiple border gateways and inter-domain connectivity to the GIG, and for the GIG serving them as their backbone network. Such multi-homed IP Autonomous Systems (ASes) would enjoy improved range extension (e.g., healing fragmented networks) and reach-back (e.g., communicating between the battlefield and CONUS) services through the GIG, while similarly, GIG routing scalability and performance would be greatly enhanced.

In Section 2, we describe the types of scenarios where mobile user networks connect to the GIG and the challenges they face. In Section 3, we propose alternatives

¹ Prepared through collaborative participation in the Communications and Networks Consortium sponsored by the U.S. Army Research Laboratory under the Collaborative Technology Alliance (CTA) Program, Cooperative Agreement DAAD19-2-01-0011. The U.S. Government is authorized to reproduce and distribute reprints for Government purposes notwithstanding any copyright notation thereon.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 01 NOV 2006		2. REPORT TYPE N/A		3. DATES COVERED -	
4. TITLE AND SUBTITLE Scalable And Secure Ipv6 Solutions For Connecting Mobile Networks To The Gig				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Telcordia Technologies Inc., Piscataway, NJ 08854				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited					
13. SUPPLEMENTARY NOTES See also ADM002075., The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 6	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

available within IPv6 to support the interconnection of the mobile networks with the GIG. In Section 4, we discuss how the use of enhanced IPv6 features can help improve the performance when nodes move, networks fragment and gateways are lost. We also investigate the effects on mobility management, routing and security in this paper, but do not compare the alternate location management, routing, or security approaches.

2. LARGE MOBILE NETWORKS MULTIHOMED TO THE GIG

This section describes the challenges in connecting large dynamic future military mobile user networks, such as WIN-T, to the Global Information Grid (GIG).

The GIG will interconnect many military IPv6 user network domains, much like the public Internet interconnects many ISPs today. The GIG backbone will include, for example, the high-speed GIG – Bandwidth Expansion (GIG-BE), Transformational Satellite Communication System (TSAT), Army LandWarNet, the Navy/Marine Corps FORCEnet, Air Force Constellation Net, and possibly the networks of Allied Forces.




Figure 1: GIG Provider – User Network Connectivity

Figure 1 shows the envisioned complex hierarchy of the future military networks (GIG), in which the larger gray clouds represent Tier1 military backbones and the dark teal clouds represent Tier2 military provider networks. Smaller clouds in light turquoise indicate the mobile user networks (some are multi-homed to their providers, many others are single-homed) and the red boxed figures represent individual users in MILS enclaves behind HAIPE. BGP speaking routers manage the inter-domain (inter-AS) connectivity, where they implement the routing policies (e.g., advertise/filter route prefixes) in accordance with the Service Level Agreements (SLAs)

between the provider and the user networks. Double lines in Figure 1 represent multiple independent connections, hence would be supported by multiple BGP routers (i.e., group of routers as indicated in Figure 1 legend). In particular, two BGP routers per connection are needed, one at each end of the inter-domain link. Future military networks, organized as the Global Information Grid, will not be any simpler than today's Internet. In fact, they will be much more complex due to the highly mobile nature of the agile Joint Tactical Forces, governing mobile provider networks that will support several black network enclaves and mobile nodes and red end hosts behind HAIPEs. As we try to show in Figure 1 by our vision of Army's future LandWarNet, mobility is not solely a single-node specific event; the GIG will support large numbers of highly dynamic mobile networks including provider networks, user networks, and lone users that may move away from their home networks.

We believe that the future backbone is going to be relatively stable in both the Internet and the GIG, but achieving routing stability in the mobile parts of the GIG will be a harder goal to attain. Mobile network support requires careful network architecture planning and additional capabilities to enhance the stability and reliability of connections. Both the commercial Internet and the tactical GIG will similarly use:

- IPv6-only. In the GIG this is mandated by the federal government to prevent the predicted IPv4 address exhaustion. Connections to IPv4-only user networks will be through tunnels or NAT-PTs.
- Service level agreements to govern their relations with the attaching user networks.
- Special Autonomous System (AS) Border Routers (ASBR) or Points of Presence (PoPs) to interconnect with the user networks.
- Policy-based BGP routing.

The dynamics of the Joint Tactical Forces, however, will impose tougher requirements and challenges than the networks of the public Internet. In contrast to the fixed commercial infrastructure, the future military networks will span Tier 1 and Tier 2 peer domains that perform transit services as well access services for predominantly ad hoc mobile users and user networks. These user networks will include not only mobile nodes, but mobile networks (terrestrial, airborne, and under-water) that will want to seamlessly maintain their connections to other networks and to the GIG as they move. The mobile user networks of the Joint Forces add a lot of requirements for the military networks:

- Reach-back service across many alternative Points of Presence (PoPs) to the GIG. Each node (and

ideally each application) may want to have multiple connections to the backbone and choose the best possible connectivity alternatives (e.g., multi-homing).

- User Network Fragmentation. Arbitrary network splits/merges in intra-domain communication. Multi-homed user networks may also require VPN-based range extension through the GIG with potentially significant routing table impacts.
- Mobile Gateways: Points of Presence (PoPs) with the GIG will vary over time, due to gateway and/or user network mobility.
- Heterogeneous radio links between the GIG and user networks. Use of terrestrial, airborne and satellite links may create significant jitter and packet reordering when routing switches over different heterogeneous radio links.
- Intermittent Links between the GIG and user networks. Terrestrial, airborne and satellite wireless links will have highly variable impairments that make the inter-network connectivity to and from the GIG-BE much more intermittent and lossy.
- Enhanced Policy-based BGP routing. DoD technical working groups are investigating enhanced BGP features for security and better resiliency to intermittent wireless interconnections.
- Application Persistence Requirements. C4ISR applications need to seamlessly connect from anywhere to anywhere irrespective of mobility, impairments and attacks. While some tolerate intermittent connectivity (e.g., situational awareness), others don't (e.g., Netfires).
- Stringent Security Requirements. Forcing a separation between applications and users in MILS red enclaves and the intermittent wireless black networks.

3. PROPOSED APPROACH

This section describes how we propose to leverage the power of IPv6 to enhance the mobility experience of tactical user networks, while also ensuring routing performance and scalability GIG-wide. The following subsections describe the primary IPv6 features we will utilize.

3.1 IPv6 Address Summarization

We propose that large mobile user networks be allocated a single address prefix to allow easy summarization. It should also be a large enough address space to allow for division into smaller address prefixes that allow the flexible hierarchy within the AS.

The 128-bit IPv6 (Deering 1998) unicast addresses (Deering 2003) are aggregatable with prefixes of arbitrary bit-length (CIDR). For example, the interface with IP address:

<12AB:0:0:CD30:123:4567:89AB:CDEF>

is part of the subnet *<12AB:0:0:CD30::/60>*. Informational RFC 3177 (IAB, 2001) provides IETF recommendations to the addressing registries (e.g., ARIN) for assigning IPv6 address blocks to end sites. In general, allocation must balance address conservation with renumbering costs: They propose assignment of /48 to end systems in the general case, but it is also possible to assign large networks with a shorter prefix or multiple /48's.

For our application, we propose that large mobile user networks be allocated a larger block of addresses (e.g., a /40) to allow the organization of the AS into a hierarchy.

The larger IPv6 address space also allows division of address prefixes based on the ASBR (section 4 describes why this is helpful to deal with node mobility or network splits).

3.2 IPv6 Multihoming

IPv4 has always allowed multihoming by assigning different IP address prefixes to different interfaces. Unlike IPv4, however, a single IPv6 interface (e.g., on an ad hoc node) can have multiple addresses simultaneously (of the same or different type) to support functions such as soft handover and multi-homing. IPv6 thus allows a node with a single interface to send or receive packets through different ASBRs simply by switching among its assigned addresses.

A key idea is to leverage the large IPv6 address space and the capability of an IPv6 interface to own multiple IPv6 addresses in order to cope with large-scale, frequent, and potentially-chaotic mobility within the user network without impacting the routing scalability in the GIG

3.3 IPv6 Prefix Delegation

We propose that any ASBR can use IPv6 prefix delegation to dynamically request a pool of addresses for handing out to mobile nodes. This address pool can either come from the GIG or from the backbone of the user network.

Unlike IPv4, IPv6 allows the transfer of whole pools of addresses, as represented by address prefixes. This address allocation can work across AS boundaries; in fact, it is primarily designed for a service provider to assign a prefix to a Customer Premise Equipment (CPE) device

acting as a router between the subscriber's internal network and the service provider's core network. Specifically, there is an IETF standards-track DHCPv6-based protocol (Troan, 2003) for automated delegation of IPv6 prefixes from a “delegating router” to a “requesting router.” across an administrative boundary.

The prefix delegation process begins when the requesting router requests configuration information through a DHCPv6 Solicit message option. When the delegating router receives the DHCPv6 messages from the requesting router, it selects an available prefix or prefixes for delegation to the requesting router. The delegating router then returns the prefix or prefixes to the requesting router in the options field of a DHCP Advertise message.

3.4 IPv6 Node Autoconfiguration

We propose to exploit IPv6 autoconfiguration facilities to allow dynamic addressing and readdressing of individual interfaces (possibly using the pools handed out by IPv6 prefix delegation). We propose to flexibly choose among the different options in IPv6: Stateless Autoconfiguration (Thomson 1998), Stateless DHCPv6 and Stateful DHCPv6 (Droms 2003); only the latter is available in IPv4.

Every IPv6 interface must generate a link-local address. It is used to reach neighboring nodes attached to the same link. The interface forms a link-local address by simply appending a unique interface's identifier (e.g., based on a 48-bit MAC address (EUI-48)) to the well-known link-local prefix (FE80::10). This “tentative” address is then checked using the Duplicate Address Detection (DAD) protocol that is part of the IPv6 Neighbor Discovery Protocol.

The default autoconfiguration of hosts is to use the Stateless Autoconfiguration. Routers send a Router Advertisement message (ICMPv6 type 134), unicast to an inquiring host's link local address or periodically multicast. The Router Advertisements contain **IPv6 Prefixes** to which each host simply attaches the least significant 64 bits of its link-local address to this prefix to generate its global IPv6 address.

IPv6 clients that are configured to use DHCPv6 autoconfiguration instead of SLAC, perform a link-local multicast to solicit DHCPv6 servers who can meet their requirements. After receiving an Advertise message from at least one DHCP server, the client can then exchange unicast messages with a server of its choice, to request addresses and other configuration information, and later to update the configuration as needed. Servers can (and do by default) maintain state that keeps track of each client (hence DHCPv6 is called the stateful

autoconfiguration in contrast to the SLAC stateless autoconfiguration). However, Servers can also not maintain state (e.g., when just configuring a DNS server location).

3.5 Mobile IPv6

Mobile IP allows an IP node to arbitrarily change its IP address and still maintain existing sessions. Mobile IPv4 forces all packets to pass through a Home Agent. Thus, a node sending to the mobile node uses the home address of the mobile node to send packets. These packets are intercepted by the home agent, which tunnels the packets to the mobile node's care-of address. In contrast, Mobile IPv6 includes additional mechanisms to facilitate more efficient communication after changing IPv6 address. In particular it allows the mobile node to send a binding update packet to the corresponding node to eliminate the need to go to the home agent for all communication.

4. ANALYSIS OF PROPOSED IPV6 ARCHITECTURE

This section describes how the proposed IPv6 enhancements may effectively solve problems expected in interconnecting future mobile user networks with the GIG.

Our goal is to enhancing the mobility experience of the mobile nodes and networks within multi-homed networks, as they move between different intra-AS domains of their network. Also we want to efficiently support multiple gateways and connections to the GIG in order to provide more robust connectivity. The presumably high mobility of nodes and sub-networks between the gateway nodes may break IPv6 address summarization and threaten routing scalability in the GIG. Our proposed solutions help preserving the address summarization and routing scalability in the GIG. The three issues we address are:

1. **Mobility between different routing domains within the user network.** In multi-homed user networks, mobility between domains may break the address summarization in the GIG as a result of the consequent routing updates and advertisements of more specific routes (prefixes) to the provider network(s) and that requires routing to converge for the mobile node to continue its sessions.
2. **Potential splitting (and merging) of intra-AS domains.** Network fragmentation may potentially causing large BGP update traffic to the provider network. The BGP update volume may be especially significant if the provider

network is offering BGP/MPLS based VPN services to the user network to enhance its intra-domain connectivity.

3. **Loss of gateways that connect the mobile network to the GIG.** Loss of gateway nodes may take place due to many reasons, such as mobility, intermittent links, jamming, and physical destruction, which may all result in significant routing updates to the GIG, degrading the connectivity experience of the mobile nodes.

Our proposed solution to problems 1 and 2 above is primarily based on IPv6 re-addressing. Our proposal is to readdress the mobile node that move into another domain. This ensures that address summarization is maintained in the route advertisements to the provider network, which has tremendously positive impact on the GIG routing scalability. Otherwise, it is not hard to imagine that routing tables may explode and GIG routing may be uncontrollably fluctuating due to frequent and significant amounts of mobility events in chaotic, unorganized patterns. This may lead to significant amount of /128 routes (e.g. per node routes) appearing in the GIG routing tables as address summarization may not be practically achievable.

The mobile nodes will also benefit from this solution through reduced latency handoffs and better performance (reduced packet loss during handoffs) in this solution. We believe performing a mobility registration of the new address binding and informing the current corresponding hosts is going to be much faster than waiting for the routing update to converge across the GIG. Mobility registration requires implementation of a Mobility Solution by the Provider Network, and any standard Mobility Support Protocol, e.g. Mobile IPv6, can be used to achieve this functionality.

We will discuss solution to problem 3 further below, as it is a different approach independent from and not relying on IPv6 readdressing solutions

As an example to illustrate the possibilities made possible by IPv6, Figure 2 shows how a multi-homed network (e.g., WIN-T) with multiple gateways to the provider network (e.g., GIG-BE) may preserve address summarization and routing scalability in the GIG routing tables, despite potentially intense route fluctuations between the intra-AS (IGP) domains of the mobile user network due to mobility of nodes between domains.

Figure 2 illustrates two inter-domain connections only, mainly for simplicity but it is fair to assume that a user network such as WIN-T in reality will support thousands of IP routers (e.g., 5000 nodes) and several tens of gateways (e.g., 100 BGP border routers) to connect to

the GIG backbone. It will also have hundreds of IPv6 subnets, operating over heterogeneous waveforms, and possibly hierarchically organized according to military structures and mission tasks.




Figure 2: Multi-homed Battlefield Network

In this example in Figure 2, a provider may be delegating multiple IPv6 address prefixes to a user network (e.g., one /48 per gateway). Each mobile node in the user network is assigned an IPv6 address from each /48 address prefix delegated to the user network border routers (e.g. via DHCPv6 prefix delegation).

The networks in this case also preserve address summarization and routing scalability for the GIG, in spite of high rate of topology changes, that force (or make it desirable) for nodes or sub-networks to use different gateway nodes. This approach leverages the large IPv6 address space and an IPv6 interface's ability to own multiple IPv6 addresses. A mobile node or a group of nodes may arbitrarily move between the two gateways without breaking route summarization in routing tables of provider routers. This alleviates any need to keeping track of /128 host addresses. Upon movement to a different gateway, the address bindings on the interface will need to be updated and registered with the nearest location management server so as to inform the remote end users of active sessions of the address binding change of the mobile node (e.g. as in Mobile IPv6).

It is possible to have multiple Security Associations (SAs) in advance between the communicating nodes for each address, minimizing the time to perform this locally-confined Layer 3 handoff, without sacrificing from security, but achieving almost seamless mobility that does not suffer large route convergence delays. An alternative

and more efficient solution may be to make the Security Associations independent of the interface addresses that change upon mobility, and to use permanent identifiers (e.g., Home Network IPv6 Address) in establishing and refreshing SAs. A standards based implementation of this is discussed in IETF RFC 4423 (Moscowitz, 2006) and in the IETF HIP Working Group Internet drafts.

Figure 3 illustrates how we may also overcome the 3rd problem of gateway loss, when border routers and direct BGP peering connections between the provider and user network border routers disappear. Gateway losses are unfortunate, but potentially expected events, and may be overcome exploiting multiple redundant connections to the GIG and assigning them as secondary BGP routes (for use in the event that primary route fails). This can be achieved through the configuration of tunnels as shown in Figure 3. Figure 3 simply illustrates an old networking solution proposed by IETF RFC 3178 (Bates, 1998) for IPv4 networks. The same solution is directly applicable in the case of IPv6 networks.




Figure 3: Exploiting Redundant Connectivity

CONCLUSIONS

Our main motivation in this paper is to show that leveraging the power of IPv6 will enhance the mobility experience of tactical user networks, while also ensuring routing performance and scalability GIG-wide.

We see effective use of IPv6 as a critical need in the face of enormous growth predictions for mobile users and mobile user networks, and in anticipation of the various, and potentially chaotic, mobility patterns within and amongst these networks, especially in battlefield situations. We believe IPv6 facilities such as autoconfiguration, scalable routing and summarizable addressing, efficient mobility, and security will mutually serve the future user and provider networks, commercial and military alike.

Future work will investigate dividing the nodes within a large AS into routing domains based on their proximity to particular ASBRs. We will also investigate the use of intra-domain routing protocols such as OSPF to facilitate the discovery of the best ASBR

REFERENCES

- Bates, T., Rekhter, Y., 1998: Scalable Support for Multi-homed Multi-provider Connectivity in RFC 2260.
- Deering, S and Hinden, R., 1998: Internet Protocol, Version 6 (IPv6) Specification in RFC 2460.
- Deering, S and Hinden, R., 2003 Internet Protocol Version 6 (IPv6) Addressing Architecture in RFC 3513.
- Droms, R., Bound, J, Volz, B., Lemon, T., Perkins, C., Carney, M., 2003: Dynamic Host Configuration Protocol for IPv6 (DHCPv6), RFC 3315.
- Droms, R., 2004: Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6 in RFC 3736
- IAB and IESG, 2001: IAB/IESG Recommendations on IPv6 Address Allocations to Sites in RFC 3177.
- Thomson, S., Narten, T., 1998: IPv6 Stateless Address Autoconfiguration in RFC 2462.
- Troan, O., Droms, R., 2003: IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6 in RFC 3633.²
- Moscowitz, R., Nikander, P., 2006: Host Identity Protocol (HIP) Architecture in RFC 4423.³

² The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied of the Army Research Laboratory or the U.S. Government.