

Independent Submission
Request for Comments: 6139
Category: Informational
ISSN: 2070-1721

S. Russert, Ed.
Unaffiliated
E. Fleischman, Ed.
F. Templin, Ed.
Boeing Research & Technology
February 2011

Routing and Addressing in Networks with
Global Enterprise Recursion (RANGER) Scenarios

Abstract

"Routing and Addressing in Networks with Global Enterprise Recursion (RANGER)" (RFC 5720) provides an architectural framework for scalable routing and addressing. It provides an incrementally deployable approach for scalability, provider independence, mobility, multihoming, traffic engineering, and security. This document describes a series of use cases in order to showcase the architectural capabilities. It further shows how the RANGER architecture restores the network-within-network principles originally intended for the sustained growth of the Internet.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This is a contribution to the RFC Series, independently of any other RFC stream. The RFC Editor has chosen to publish this document at its discretion and makes no statement about its value for implementation or deployment. Documents approved for publication by the RFC Editor are not a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc6139>.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

Table of Contents

1. Introduction	3
2. Terminology	4
3. Approach	7
4. Scenarios	11
4.1. Global Concerns	11
4.1.1. Scaling the Global Inter-Domain Routing Core	11
4.1.2. Supporting Large Corporate Enterprise Networks	13
4.2. Autonomous System Concerns	16
4.3. Small Enterprise Concerns	16
4.4. IPv4/IPv6 Transition and Coexistence	18
4.5. Mobility and MANET	21
4.5.1. Global Mobility Management	21
4.5.2. First-Responder Mobile Ad Hoc Networks (MANETs)	23
4.5.3. Tactical Military MANETs	24
4.6. Provider Concerns	27
4.6.1. ISP Networks	27
4.6.2. Cellular Operator Networks	28
4.6.3. Aeronautical Telecommunications Network (ATN)	28
4.6.4. Unmanaged Networks	31
5. Mapping and Encapsulation Concerns	32
6. Problem Statement and Call for Solutions	32
7. Summary	33
8. Security Considerations	33
9. Acknowledgements	34
10. References	34
10.1. Normative References	34
10.2. Informative References	34

1. Introduction

The Internet is continually required to support more users, more internetwork connections, and increasing complexity due to diverse policy requirements. This growth and change strains the infrastructure and demands new solutions. Some of the complementary approaches to transform Internet technology are being pursued concurrently within the IETF: translation (including Network Address Translation (NAT)), tunneling (map and encapsulate), and native IPv6 [RFC2460] deployment. Routing and Addressing in Networks with Global Enterprise Recursion (RANGER) [RFC5720] describes the architectural elements of a "map and encapsulate" approach that also facilitates the other two approaches. This document discusses RANGER operational scenarios.

RANGER provides an architectural framework for scalable routing and addressing. It provides for scalability, provider independence, mobility, multihoming, and security for the next-generation Internet. The RANGER architectural principles are not new. They can be traced to the deliberations of the ROAD group [RFC1380], and also to still earlier works including NIMROD [RFC1753] and the Catenet model for internetworking [CATENET] [IEN48] [RFC2775]. [RFC1955] captures the high-level architectural aspects of the ROAD group deliberations in a "New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG".

The Internet has grown tremendously since these architectural principles were first developed, and that evolution increases the need for these capabilities. The Internet has become a critical resource for business, for government, and for individual users throughout the developed world. RANGER carries forward these historic architectural principles, creating a ubiquitous enterprise network structure that can represent collections of network elements ranging from the granularity of a singleton router all the way up to an entire Internet. This enterprise network structure uses border routers that configure tunnel endpoints to connect potentially recursively nested networks. Each enterprise network may use completely independent internal Routing Locator (RLOC) address spaces, supporting a virtual overlay network connecting edge networks and devices that are addressed with globally unique Endpoint Interface iDentifiers (EIDs). The RANGER virtual overlay can transcend traditional administrative and organizational boundaries. In its purest form, this overlay network could therefore span the entire Internet and restore the end-to-end transparency envisioned in [RFC2775].

The RANGER architecture drew early observations from the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [RFC5214] [RFC5579] but now uses Virtual Enterprise Traversal (VET) [RFC5558], the Subnetwork

Encapsulation and Adaptation Layer (SEAL) [RFC5320], and other mechanisms including IPsec [RFC4301] as its functional building blocks. This document describes use cases and shows how the RANGER mechanisms apply. Complementary mechanisms (e.g., DNS, DHCP, NAT, etc.) are included to show how the various pieces can work together. It expands on the concepts introduced in "IPv6 Enterprise Network Scenarios" [RFC4057] and "IPv6 Enterprise Network Analysis - IP Layer 3 Focus" [RFC4852], and shows how the enterprise network model generalizes to a broad range of scenarios. These use cases are included to provide examples, invite criticism and comment, and explore the potential for creating the next-generation Internet using the RANGER architecture. Familiarity with RANGER, VET, SEAL, and ISATAP are assumed.

2. Terminology

Internet Topology Hierarchy

The Internet Protocol (IP) natively supports a topology hierarchy comprised of increasing aggregations of networked elements. Network interfaces of devices are grouped into subnetworks, and subnetworks are grouped into larger aggregations. Subnetworks can be optionally grouped into areas and the areas grouped into an autonomous system (AS). Alternatively, subnetworks can be directly grouped into an AS. The foundation of the IP Topology Hierarchy is the AS, which determines the administrative boundaries of a network deployment including its routing, addressing, quality of service, security, and management. Intra-domain routing occurs within an autonomous system, and inter-domain routing links autonomous systems into a "network of networks" (Internet).

Routing Locator (RLOC)

an address assigned to an interface in an enterprise-interior routing region. Note that RLOC space is local to each enterprise network.

The IPv4 public address space currently in use today can be considered as the RLOC space for the global Internet as a giant "enterprise network".

Endpoint Interface iDentifier (EID)

an address assigned to an edge network interface of an end system. Note that EID space is global in scope, and must be separate and distinct from any RLOC space.

commons

an enterprise-interior routing region that provides a subnetwork for cooperative peering between the border routers of diverse organizations that may have competing interests. An example of a commons is the Default-Free Zone (DFZ) of the global Internet. The enterprise-interior routing region within the commons uses an addressing plan taken from RLOC space.

enterprise network

the same as defined in [RFC4852], where the enterprise network deploys a unified RLOC space addressing plan within the commons, but may also contain partitions with disjoint RLOC spaces and/or organizational groupings that can be considered as enterprises unto themselves. An enterprise network therefore need not be "one big happy family", but instead provides a commons for the cooperative interconnection of diverse organizations that may have competing interests (e.g., such as the case within the global Internet Default-Free Zone).

Historically, enterprise networks are associated with large corporations or academic campuses. However, in RANGER an enterprise network may exist at any IP Topology Hierarchy level. The RANGER architectural principles apply to any networked entity that has some degree of cooperative active management. This definition therefore extends to home networks, small office networks, a wide variety of Mobile Ad hoc Networks (MANETs), and even to the global Internet itself.

site

a logical and/or physical grouping of interfaces within an enterprise network commons, where the topology of the site is a proper subset of the topology of the enterprise network. A site may contain many interior sites, which may themselves contain many interior sites in a recursive fashion.

Throughout the remainder of this document, the term "enterprise" refers to either enterprise or site; i.e., the RANGER principles apply equally to enterprises and sites of any size or shape. At the lowest level of recursive decomposition, a singleton Enterprise Border Router can be considered as an enterprise unto itself.

Enterprise Border Router (EBR)

a node at the edge of an enterprise network that is also configured as a tunnel endpoint in an overlay network. EBRs connect their directly attached networks to the overlay network, and connect to other networks via IP-in-IP tunneling across the commons to other EBRs. This definition is intended as an

architectural equivalent of the functional term "EBR" defined in [RFC5558], and is synonymous with the term "xTR" used in other contexts (e.g., [LISP]).

Enterprise Border Gateway (EBG)

an EBR that also connects the enterprise network to provider networks and/or to the global Internet. EBGs are typically configured as default routers in the overlay, and provide forwarding services for accessing IP networks not reachable via an EBR within the commons. This definition is intended as an architectural equivalent of the functional term "EBG" defined in [RFC5558], and is synonymous with the term "default mapper" used in other contexts (e.g., [APT]).

overlay network

a virtual network manifested by routing and addressing over virtual links formed through automatic tunneling. An overlay network may span many underlying enterprise networks.

6over4

"Transmission of IPv6 over IPv4 Domains without Explicit Tunnels" [RFC2529]; functional specifications and operational practices for automatic tunneling of unicast/multicast IPv6 packets over multicast-capable IPv4 enterprise networks.

ISATAP

Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) [RFC5214] [RFC5579]; functional specifications and operational practices for automatic tunneling over unicast-only enterprise networks.

VET

Virtual Enterprise Traversal (VET) [RFC5558]; functional specifications and operational practices that provide a functional superset of 6over4 and ISATAP. In addition to both unicast and multicast tunneling, VET also supports address/prefix autoconfiguration as well as additional encapsulations such as IPsec, SEAL, UDP, etc.

SEAL

Subnetwork Encapsulation and Adaptation Layer (SEAL) [RFC5320]; a functional specification for robust packet identification and link MTU adaptation over tunnels. SEAL supports effective ingress filtering and adapts to subnetworks configured over links with diverse characteristics.

Within the RANGER architectural context, the SEAL "subnetwork" and RANGER "enterprise" should be considered as identical abstractions.

Provider-Independent (PI) prefix

an EID prefix (e.g., 2001:DB8::/48, 192.0.2/24, etc.) that is routable within a limited scope and may also appear in enterprise network mapping tables. PI prefixes that can appear in mapping tables are typically delegated to a BR by a registry, but are not aggregated by a provider network.

Provider-Aggregated (PA) prefix

an EID prefix that is either derived from a PI prefix or delegated directly to a provider network by a registry. Although not widely discussed, it bears specific mention that a prefix taken from a delegating router's PI space becomes a PA prefix from the perspective of the requesting router.

Customer Premises Equipment (CPE) Router

a residential or small office router that provides IPv4 and/or IPv6 support. The user or the service provider may manage the router.

Carrier-Grade NAT (CGN)

a special (usually high capacity) IPv4-to-IPv4 NAT deployed within the service provider network that serves multiple subnets.

3. Approach

The RANGER [RFC5720] architecture seeks to fulfill the objectives set forth in [RFC1955]:

- o No Changes to Hosts
- o No Changes to Most Routers
- o No New Routing Protocols
- o No New Internet Protocols
- o No Translation of Addresses in Packets
- o Reduce the Routing Table Size in All Routers
- o Use the Current Internet Address Structure

The RANGER enterprise network is a cooperative networked collective sharing a common (business, social, political, etc.) goal. An enterprise network can be simple or complex in composition and can operate at any IP Topology Hierarchy level. Although RANGER focuses on encapsulation, it is also compatible with both native and translated routing and addressing.

RANGER enables a protocol and/or addressing system to be connected in a virtual overlay across an untrusted transit network, or "commons". While it does not show all possible uses, Figure 1 illustrates that RANGER supports the creation of a distributed network across an intervening commons, which could implement a dissimilar IP version, routing protocol, or addressing system.

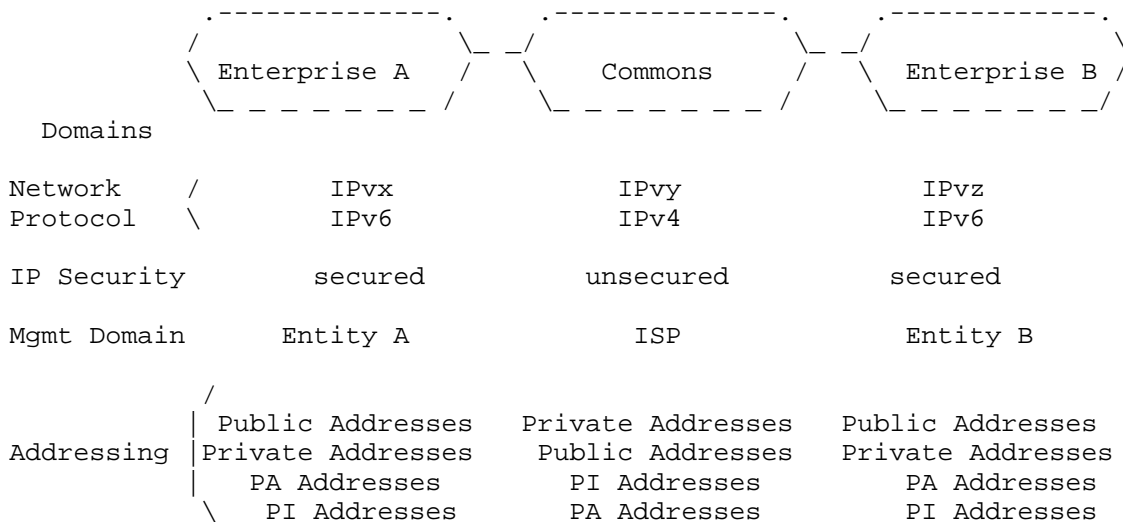


Figure 1. RANGER Links Distributed Enterprise Networks

The RANGER concepts can be applied recursively. They can be implemented at any level within the IP Topology Hierarchy to create an enterprise-within-enterprise organizational structure extending traditional AS, area, or subnetwork boundaries. This structure uses border routers that configure tunnel endpoints to enable communications between potentially recursively nested enterprise networks in a virtual overlay network that transcends traditional administrative and organizational boundaries. In its purest form, this overlay network could therefore span the entire Internet and restore end-to-end transparency [RFC2775].

The RANGER architecture applies the best current practice insights from previous encapsulation systems as they are currently articulated within the Virtual Enterprise Traversal [RFC5558], and Subnetwork Encapsulation and Adaptation Layer [RFC5320] functional specifications. The result is an architecture and protocol system that can be used to create arbitrarily complex, scalable IP deployments that support both unicast and multicast routing and addressing systems.

RANGER supports scalable routing through a recursively nested enterprise-within-enterprise network capability. The fundamental building block is the Enterprise Border Router (EBR) (see Figure 2). The EBR is the limiting factor for RANGER recursion, and in certain contexts a singleton EBR can be viewed as an enterprise network unto itself. Traditional network infrastructures can be extended to support complex structures solely with the addition of EBRs with no other modification to any networked entity.

An EBR can be a commercial off-the-shelf router, a tactical military radio, an aircraft mobile router, etc., but it can also be an end system (e.g., a laptop computer, a soldiers' handheld device, etc.) with an embedded gateway function [RFC1122].

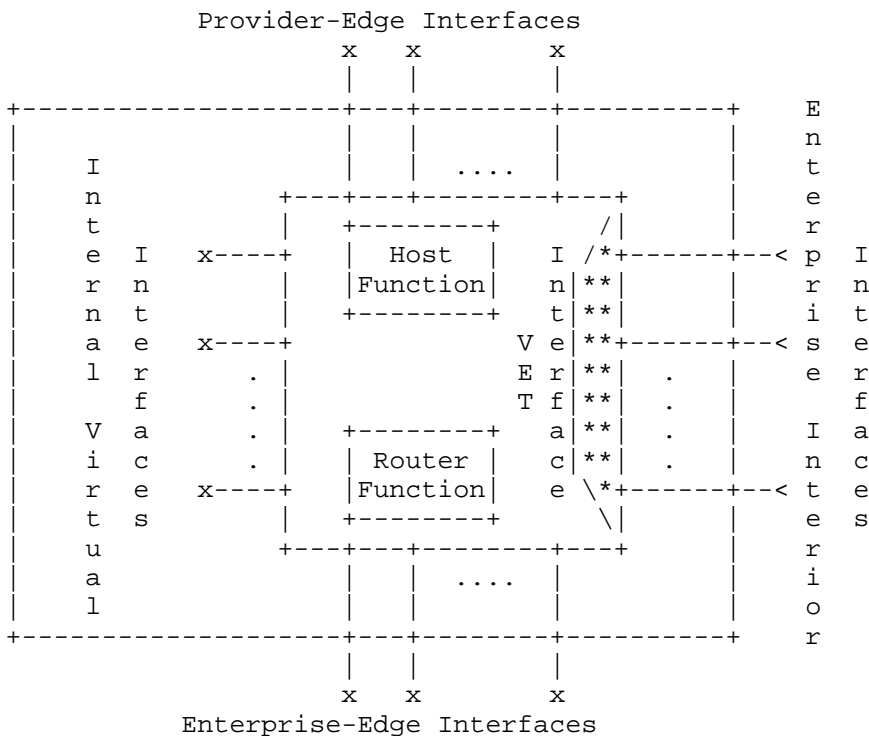


Figure 2. Enterprise Border Router (EBR)

EBRs connect networks and end systems to one or more enterprise networks via a repertoire of interface types. Enterprise-interior interfaces attach to a commons. Provider-edge interfaces support

traditional routing relationships up the IP Topology Hierarchy, and enterprise-edge interfaces support traditional relationships down the IP Topology Hierarchy. Internal virtual interfaces are typically loopback interfaces or VMware-like host-in-host interfaces.

VET interfaces support RANGER recursion and IP-in-IP encapsulation. VET interfaces are configured over provider-edge, enterprise-interior, or enterprise-edge interfaces to allow recursion horizontally or vertically within the IP Topology Hierarchy. A VET interface may be configured over several underlying interfaces that all connect to the same enterprise network. This creates a link-layer multiplexing capability that can provide several advantages (see [RFC1122], Section 3.3.4). One important advantage is continuous operation across failovers between multiple links attached to the same enterprise network, without any need for readdressing.

Figure 3 shows two enterprise networks (each with their own internal addressing and routing systems) that communicate over a virtual overlay network across a commons. The virtual overlay is manifested by tunneling, which links enterprise networks separated by geographical remoteness, protocol incompatibility, or both. An ingress EBR (iEBR) within the left enterprise network seeks to forward encapsulated packets across the commons to the egress EBR (eEBR) within the right enterprise network.

The figure shows that the eEBR assigns a Routing Locator (RLOC) address on its interface to the commons' interior IP routing and address space, while the destination host assigns an Endpoint Interface Identifier (EID) on its enterprise-edge interface. The iEBR uses a mapping system to discover the RLOC of an eEBR on the path to the destination EID address. A distinct mapping system is maintained within each recursively nested enterprise network instance operating at a specific level of the IP Topology Hierarchy. RANGER uses the mapping system to join peer enterprise networks via a virtual overlay across a commons.




Figure 3. The RANGER Model

EBRs must configure both RLOC and EID addresses and/or prefixes. Autoconfiguration is coordinated with Enterprise Border Gateways (EBGs) that connect to the next-higher layer in the recursive hierarchy, as specified in VET. Standard mechanisms including DHCP [RFC2131] [RFC3315] and Stateless Address Autoconfiguration (SLAAC) [RFC4862] are used for this purpose.

Similarly, EBRs require a means to discover other EBRs and EBGs that can be used as enterprise network exit points. VET specifies mechanisms for border router discovery using the global DNS and/or enterprise-local name services such as Link-Local Multicast Name Resolution (LLMNR) [RFC4795].

The mapping system is a distributed database that is synchronized among a limited set of mapping agents. Database synchronization can be achieved by many different protocol alternatives. The most commonly used alternatives are either the Border Gateway Protocol (BGP) [RFC4271] or the Domain Name System (DNS) [RFC1035]. Mapping-system databases can be populated by many different mechanisms including administrative configuration and automated prefix registrations.

EBRs forward initial packets for which they have no mapping to an EBG. The EBG in turn forwards the packet toward the final destination and returns a redirect to inform the EBR of a better next hop if necessary. The EBR then receives a mapping reply that it can use to populate its Forwarding Information Base (FIB). It then encapsulates each forwarded packet in an outer IP header for transmission across the commons to the remote RLOC address of an eEBR. The eEBR in turn decapsulates the packets and forwards them to the destination EID address. The Routing Information Base (RIB) within the commons only needs to maintain state regarding RLOCs and not EIDs. The synchronized EID-to-RLOC mapping state is not subject to oscillations due to link state changes within the commons. RANGER supports scalable addressing by selecting a suitably large EID addressing range that is distinct from any enterprise-interior RLOC addressing ranges.

4. Scenarios

4.1. Global Concerns

4.1.1. Scaling the Global Inter-Domain Routing Core

Growth in the Internet has created challenges in routing and addressing that have been recognized for many years [RADIR-PROB-STATE]. IPv4 [RFC0791] address space is limited, and Regional Internet Registry (RIR) allocation is passing the "very

painful" Host Density (HD) ratio threshold of 86% (that is, 192M allocated addresses) [RFC3194]. As a result, exhaustion of the IPv4 address pool is predicted within the next two years [IPv4POOL], [HUSTON-END]. IPv6 promises to resolve the address shortage with a much larger address space, but transition is costly and could exacerbate BGP problems described below. Richer interconnection, increased multihoming (especially with provider-independent (PI) addresses), and a desire to support traffic engineering via finer control of routing has led to super-linear growth of BGP routing tables in the Default-Free Zone, or "DFZ", of the Internet. This growth is placing increasing pressures on router capacities and technology costs that are unsustainable for the longer term within the current Internet routing framework.

RANGER allows the coordinated reuse of addresses from enterprise to enterprise by making RLOC address spaces independent of one another. Figure 4 shows how the RANGER architecture allows the use of separate address spaces for RLOC and EID addressing in the Internet. This yields more endpoint address space, especially with the use of IPv6, and also reduces the load on BGP in the Internet routing core. Note that Figure 4 could represent variants of RFC 4057 scenarios 1 and 2.




Figure 4. Enterprise Networks and the Internet

RLOC address spaces are entirely independent of one another, as they are used only within an enterprise network (recall that an enterprise network can exist at any level of the IP Topology Hierarchy). Such an arrangement allows each RLOC space to maintain an independent routing system and thereby avoid the inherent scaling issues if a single monolithic routing system were used for all.

EID address space can be provider-aggregated (PA) or PI, and taken from either IPv4 or IPv6. EID addresses (barring the use of Network Address Translation (NAT)) are globally unique, even when routable only within a more limited scope (e.g., in their own edge networks).

The IRTF routing research group is investigating a Preliminary Recommendation for a routing architecture [RFC6115] that provides a taxonomy for routing scaling solutions for the global Internet inter-domain routing core. RANGER presents a core/edge separation architecture within this taxonomy that uniquely shows applicability from the core all the way out to edge networks via its recursive enterprise-within-enterprise framework. RANGER is further compatible with a number of schemes intending to address routing scaling issues, including "APT: A Practical Transit Mapping Service" [APT], "FIB Suppression with Virtual Aggregation" [GROW-VA], "Locator/ID Separation Protocol (LISP)" [LISP], and others.

4.1.2. Supporting Large Corporate Enterprise Networks

Each enterprise network operator must be able to manage its internal networks and use the Internet infrastructure to achieve its performance and reliability goals. Enterprise networks that are multihomed or have mobile components frequently require provider-independent addressing and the ability to coordinate with multiple providers without renumbering "flag days" [RFC4192] [RFC5887]. RANGER provides a way to coordinate addressing plans and inter-enterprise routing, with full support for scalability, provider independence, mobility, multihoming, and security.




Figure 5. Enterprise Networks within the Internet Commons

Figure 5 depicts enterprise networks E1 through Em connected to the global IPv4 Internet via Enterprise Border Routers (EBRs) X1 through X9. These same border nodes also act as Enterprise Border Gateways (EBGs) that provide default routing services for nodes within their respective enterprise networks. The global Internet forms a commons across which the various enterprise networks connect as cooperating yet potentially competing entities. Within each enterprise network there may be arbitrarily many hosts, routers, and networks (not shown in the diagram) that use addresses taken from that enterprise network's RLOC space and over which both encapsulated IP packets with (global-scoped) EID addresses and unencapsulated IP packets with (enterprise-local) RLOC addresses can be forwarded.

Each enterprise network may encompass lower-tier networks; for instance, the singleton EBR "W" in network E2 resides in a lower-tier network (say E2.1), and (along with any of its attached devices) may be considered as an enterprise unto itself. W sees Y3 and Y4 as EBGs, which in turn see X5 and X6 as EBGs that connect to a common provider network (in this case, the Internet). Each enterprise network has one or more Endpoint Interface iDentifier (EID) address prefixes used for addressing nodes on edge networks. RANGER's map-and-encaps approach separates the mapping of EIDs to Routing Locators (RLOCs) from the Routing Information Base (RIB) in the Internet commons that are assigned to EBR router interfaces. Not only does

BGP in the Internet commons only need to maintain state regarding RLOCs in the Internet commons, it has fewer unique routes to maintain because only routes to EBRs are needed; traffic engineering can therefore be accommodated via the mapping database.

In Figure 5, enterprise network E2 represents a corporation that has multiple locations and connections to multiple ISPs. The corporation has recently merged with another corporation so that its internal network has two disjoint RLOC address spaces, but neither of the formerly separate entities can bear the burden of address renumbering. Enterprise network E2 can use a suitably large IPv4 and/or IPv6 EID addressing range (that is distinct from any enterprise-interior RLOC addressing range) to support end systems on enterprise-edge networks with no disruption to preexisting address numbering.

As EBRs are deployed to connect enterprise networks together, ordinary routers within the enterprise network continue to function as normal and deliver both ordinary and encapsulated packets across the existing Internet infrastructure and the network's own RLOC commons. Legacy IPv4 services that bind to RLOC addresses continue to be supported even as EID-based services are rolled out. Where a legacy IP client and server are within the same RLOC address space, they simply communicate by using RLOC-based routing across the enterprise network commons. If the client and server are not within the same RLOC address space, they communicate through some form of network address and/or protocol translation (see [RFC5720], Section 3.3.4 for details). EBRs from the various enterprise networks publish their EID prefixes to an enterprise-specific mapping system, so that other EBRs from the various enterprise networks can consult the mapping system to receive the RLOC address of one or more EBRs that serve the EID prefix.

As an example, when an end system connected to W in E2.1 has a packet to send to node Z in enterprise network E3, W sends the packet to EBR Y4, which encapsulates the packet in an outer IP packet with its own source address and the RLOC address of the next-hop EBR as the destination -- in this case, X6. X6 decapsulates the packet and looks up the destination EID prefix, obtaining the RLOC of X7 as next-hop. X6 then encapsulates the IPv6 packet in a packet with RLOC address X6 as the source and X7 as the destination. X7 decapsulates the packet on receipt and forwards it via its enterprise-edge interface to node Z.

This example uses one thread out of many that are possible using RANGER; see [RFC5720] and [RFC5558] for other options and details. Many enterprise networks that use proxies and firewalls at their border routers today will wish to maintain that control over their enterprise borders, and the use of RANGER does not preclude such configurations (for example, see Section 4.3).

4.2. Autonomous System Concerns

An enterprise network such as E2 in Figure 5 above can represent an AS within the IP Topology Hierarchy. A possible configuration for enterprise network E2 is for each of its enterprise components to also be recursive ASs linked together using the RANGER constructs. Such a configuration is increasingly commonplace today for the networks of very large corporations (e.g., Boeing's corporate enterprise network). These networks support an internal instance of the BGP linking many corporate-internal ASs and independent from the BGP instance that maintains the RIB within the global Internet Default-Free Zone (DFZ). Such configurations are often motivated by scaling or administrative requirements.

Such a corporate entity is internally an Internet unto itself, albeit with separate default routes leading to the true global Internet. The enterprise network E2 therefore appears to the rest of the Internet as if it were a traditional IP Topology Hierarchy AS. Since RANGER supports recursion, each AS within such a network may itself use BGP internally in place of an IGP, and can therefore also internally be composed of a locally internal Internet in a recursive fashion. This enterprise-within-enterprise framework can recursively be extended as broadly and as deeply as required in order to achieve the specific requirements of the deployment (e.g., scaling, unique administration, and/or functional compartmentalization).

4.3. Small Enterprise Concerns

Global enterprise networks operating at the autonomous system level of the IP Topology Hierarchy include multiple geographical regions, multiple ISPs, and complex internal structures that naturally benefit from the application of RANGER techniques. However, all other enterprise network instances (both large and small) can also be served by RANGER. For example, Small and Home Office (SOHO) networks may comprise only a few computers on a single network segment or may extend to larger configurations with security islands, internal routers and switches, etc.

An important concern of the small enterprise network is the ability to grow the network, change ISPs, or expand to more locations without readdressing the existing network. Consider a small company that has

a single location in California. The ISP connection is via a router that acts as a Network Address Translator and firewall for the company. Addresses of the few computers ("Wksta") are taken from the [RFC1918] private address space.




Figure 6. Simple SOHO Network

This configuration has been adequate for the few employees performing software development work, since there is no need to expose services within the site to the outside world. But now a web presence is required as product introduction approaches. The network manager deploys an EBR either as a co-resident function on the existing NAT/firewall platform (as depicted in Figure 7) or on a separate platform.

The EBR has a provider-edge interface connected to the ISP; the preexisting workstations; the preexisting enterprise-edge interfaces connecting the workstations; and enterprise-edge interfaces connecting several network segments connected by routers that host web servers, workstations, and other enterprise network services. A VET interface is configured over the new service network to allow the servers to be addressed from the public Internet.

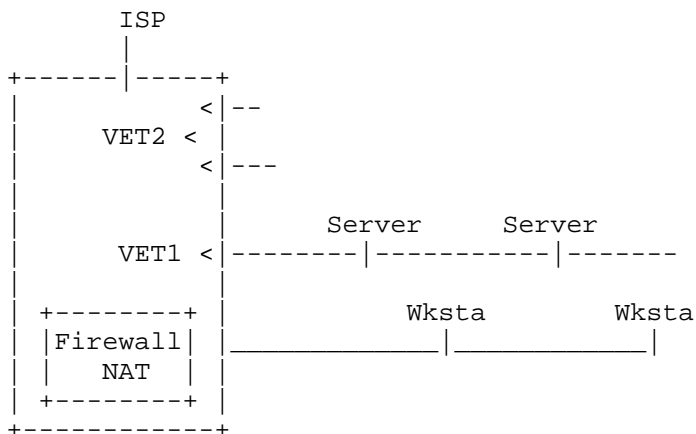


Figure 7. RANGER Serving the Small Company

In this new configuration, the EBR maintains the services within a "demilitarized zone (DMZ)" that is accessible from the public Internet without exposing other corporate assets that are still protected by the preexisting firewall/NAT functions.

Shortly afterward, an infusion of venture capital allows acceleration of the product development and marketing work by adding programmers in Tokyo and sales offices in New York and London. These new branches connect via Virtual Private Network (VPN) links across the Internet, and a new VET interface (VET2) is configured over these links to form a new sub-enterprise:

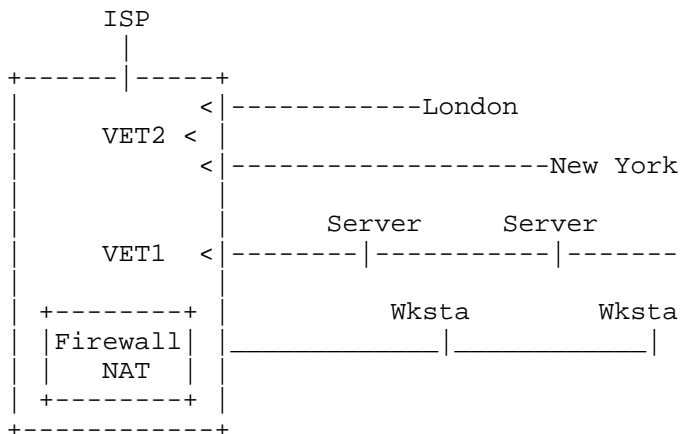


Figure 8. RANGER for Multiple Locations

4.4. IPv4/IPv6 Transition and Coexistence

End systems and networks need to accommodate long-term support for both IPv4 and IPv6. Requirements for transition include support for IPv4 applications running over IPv4 protocol stacks, IPv4 applications over IPv6 stacks, IPv4 applications over dual stacks, and IPv6 or IPv4/IPv6-capable applications over both IPv6 and dual stacks. Both encapsulation and translation will likely be needed to allow applications, enterprises, and providers to incorporate IPv6, including all intermediate states, without global coordination or a "flag day".

The RANGER architecture facilitates the addition of IPv6 addressing to existing IPv4 end systems and routers (i.e., via dual stack) as well as the addition of IPv6 networks to the existing set of IPv4 networks. RANGER (with VET and SEAL) makes it possible to carry packets originated in one protocol across a network infrastructure supporting another protocol or routing system. Figure 1 shows how

RANGER further provides a compatible framework for incorporating supporting mechanisms including protocol translation, application-layer aspects of IPv4/IPv6 transition discussed in [RFC4038], and DNS issues for IPv6 from [RFC4472]. For instances where IPv4 applications remain in use, RANGER expects that IPv4<->IPv6 translation will be supported via network-based [BEHAVE-v6v4] and/or end system stack-based (e.g., [RFC2767]) protocol translation systems. Figure 10 shows the NAT - Protocol Translation (NAT-PT)-equivalent translation in the VET router, and Figure 11 shows the "Bump-In-the-Stack" (BIS)-equivalent translation in end systems ([RFC2767]). These examples address scenarios not mentioned in RFC 4852.

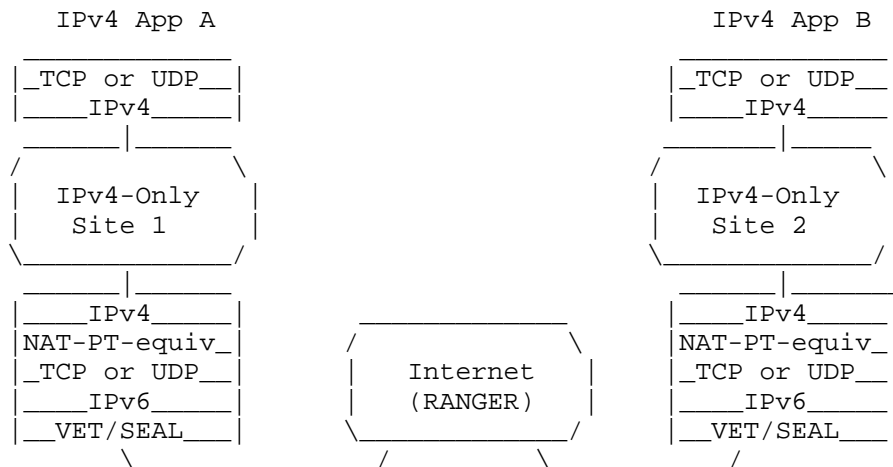


Figure 10. Translation in Routers

In Figure 10, an IPv4 application on end system A operates normally, and the end system sends IPv4 packets on the IPv4-only site network. The IPv4 packets are received by an Enterprise Border Router (EBR) that translates them into IPv6 packets by a NAT-PT-equivalent process. The EBR then encapsulates the packets into IPv4 and sends them across the RANGER-enabled Internet to Site 2 where they are received and decapsulated by an EBR for Site 2. The EBR uses NAT-PT-equivalent translation to translate the resulting IPv6 packet back to an IPv4 packet that is delivered across the Site 2 IPv4-only network to an IPv4 application on end system B.




Figure 11. BIS-Style Translation in Dual-Stack End Systems

Figure 11 shows the simplified approach using a BIS translation process within dual-stack end systems ([RFC2767]). In this case, the IPv4 application on dual-stack end system A forms an IPv4 payload, which is then transformed into an IPv6 packet within the end system protocol stack itself. The IPv6 packet can then be encapsulated and sent across the Internet to be decapsulated and sent to the dual-stack end system hosting IPv4 application B. The BIS-equivalent process on end system B reverses the translation, yielding an IPv4 packet for consumption by the IPv4-only application.

Other issues besides IP protocol translation may arise during IPv4-IPv6 transition; [RFC4038] points out issues including IPv4/IPv6-capable applications running on IPv4-only protocol stacks, DNS responses that include addresses of both IP versions, and the difficulty of supporting multiple application versions. It also advises that applications be converted to dual support as a preferred solution. These issues are outside the scope of this document.

4.5. Mobility and MANET

4.5.1. Global Mobility Management

Ubiquitous wireless access enables connection to network infrastructure nearly anywhere. Vehicles and even persons can host networks that move around with them. For example, commercial aircraft networks include requirements for nomadic networks, local mobility, and global mobility where the connection point between airplane and ground station can move from one continent to another. Mobile networks need to be able to use provider-independent (PI) as well as provider-aggregated (PA) address prefixes. Some applications such as voice require rapid or seamless connection handoffs -- also known as session survivability. Internet routing should not be unduly disrupted by mobility, so movement of mobile nodes or edge networks should not cause large ripples of routing protocol traffic, especially in the DFZ.

enterprise network (as indicated by the two different links from the aircraft to ACSP1 in Figure 12), the aircraft announces its PI prefixes at its new point of attachment and withdraws them from the old. The worldwide Internet sees no change, and mapping-system churn is confined to ACSP1, since the prefixes need not be announced or withdrawn within the parent aero enterprise network; i.e., the churn is isolated to lower tiers of the recursive hierarchy. This can be contrasted with the deprecated mobility solution previously fielded by Connexion, which propagated disruptive BGP changes into the Internet routing system to support mobile onboard networks.

4.5.2. First-Responder Mobile Ad Hoc Networks (MANETs)

Many emerging network scenarios require autoconfiguration of Mobile Ad hoc Networks (MANETs). Where first responders need networking for communications and coordination between teams, RANGER allows each team or agency to quickly stand up a network and then use the autoconfiguration described in [RFC5558] to coordinate address/prefix autoconfiguration and discover border routers needed for teams and agencies to interconnect.

For example, Figure 13 shows how police units arriving on a scene with no network infrastructure can create a wireless network using vehicle-mounted 802.11 hotspots with one or more cellular, 802.16, or satellite links in order to reach the Internet. In this example, the California Highway Patrol sets up an incident management center with a satellite link to the Internet and vet1 serving network L1. The Los Angeles County Sheriff team sets up network L1.1 at their field headquarters, and the Altadena police force creates the L1.2 network with their mobile units. R2 is the router that serves as an EBR for border routers X3 and X4, which connect networks L1.2 and L1.1, respectively. X3 serves vet3, and X4 serves vet2.

In like manner, the Angeles National Forest creates enterprise network F1, with the San Gabriel Ranger District setting up enterprise network F1.1 and the Fire Response Team Enterprise Network F1.2. R1 and R2 discover one another and become peer EBRs across the Internet by means of manual configuration. In network L1, individual PI address prefixes are announced from L1.2 and L1.1 to L1, and R2 advertises them to the satellite ISP. R1 receives a PA prefix from its WiMAX provider and delegates parts of the prefix to X1 and X2. R2 also runs an IGP with R1, advertising the PI prefixes to R1 and learning the PA prefixes there.

Figure 3 shows that RANGER supports creation of a VET interface between the enterprise-interior (network) interface of two Enterprise Border Routers (EBR) located within separate enterprise networks, A and B. When this concept is applied to enterprise networks operating above the subnetwork level of the IP Topology Hierarchy, then this VET interface uses IP-in-IP encapsulation. This corresponds with a popular COMSEC approach (IPsec -- [RFC4301]). When this same RANGER concept is applied to enterprise networks operating at the subnetwork level of the IP Topology Hierarchy, then this corresponds to an older form of COMSEC (Link Layer Encryption). When the same RANGER concept is applied to enterprise networks being singleton EBR nodes (i.e., the interface level of the IP Topology Hierarchy), then this corresponds to a third military COMSEC alternative (Link Encryption).

The previous paragraph shows the flexibility of the RANGER architecture to describe COMSEC approaches in terms of IP Topology Hierarchy structured relationships. The power of the RANGER architecture becomes apparent when one recognizes that each of the entities in Figure 3 may themselves be simple or complex network structures operating at any specific level of the IP Topology Hierarchy. (Complex structures refer to architectures that have been extended by RANGER recursion.) For example, the commons in the figure may itself be an interface, a subnetwork, an autonomous system, or an Internet. Enterprise networks A and B can be a single end system, a subnetwork, an autonomous system, or an Internet.

Tactical military MANETs differ from traditional networks in many ways, the most obvious being the high mobility of tactical deployments and self-forming-network attributes of MANETs themselves. Because each networked tactical entity supports a radio/router, the numbers of routers within military MANETs can be orders of magnitude more numerous (denser) than traditional civilian networks. This means that even small deployments have comparatively large router populations when compared to non-MANET deployments. Larger router populations directly create greater sensitivity to protocol scalability issues. Router scalability issues are further exacerbated because IP protocols react unfavorably to signal intermittence, which effectively dampens and constrains router scaling even when mitigation techniques are employed. Signal intermittence itself is a characteristic of mobility and the radio signal propagation attributes of local deployment environments (e.g., such issues as terrain, foliage, buildings, weather, distance, etc.). War fighting also encourages war fighters to locate into more defensible terrain features, many of which naturally reduce radio signal propagation, further increasing the probability of signal intermittence.

RANGER recursion enables MANETs that naturally encourage route aggregation and scaling through simple "plug and play" hierarchical arrangements that parallel organizational structures and do not entail complex manual configurations. For example, a MANET autonomous system may benefit from RANGER recursion by being physically comprised of enterprise networks that are autonomous systems themselves. This relationship can be recursively extended vertically as deep as required in order to create route aggregation between entities having common mission assignments at differing levels of abstraction. Since MANET routing is an active research topic, it is helpful to realize that these structures may or may not use routing protocols similar to their civilian IP Topology Hierarchy peers. For example, because of the behavior of BGP within highly mobile environments, the Exterior Gateway Protocol (EGP) used to link ASs may or may not be BGP and, if it is BGP, it may have unusual timer settings. However, whatever IGP and EGP is used, RANGER constructs can increase route aggregation between entities sharing common mission assignments to enable route scaling.

Tactical military MANETs often have requirements to communicate with stationary infrastructures. By localizing mobility into an enterprise network, the specific mobility-friendly protocols can then be localized and their aggregation results presented to the stationary network using a protocol supported by the stable network. This also reduces the impact of mobility upon routing and addressing systems as reported to the stationary infrastructure. Mobility-induced route fluctuations (e.g., routing flaps) can still occur, but their impact can be dampened if RANGER constructs are used to localize them in lower tiers of the IP Topology Hierarchy. For example, enterprise network A in Figure 3 can be a military MANET, and enterprise network B may be a stationary military entity. Recall that enterprise networks A and B interface at a specific IP Topology Hierarchy level, but they may be physically extended by RANGER mechanisms. For example, enterprise network A can be a MANET enterprise that is physically a network-of-networks Internet that interfaces to enterprise network B as if it were an autonomous system. This gives enterprise network B a more stable and aggregated view of the enterprise network A Internet than would be the case if it were directly aware of A's various sub-enterprise components.

Another key distinctive feature of tactical military networks is that, because radio networks operate at a different classification level than the data they convey, tactical military networks have several orders of magnitude more COMSEC devices than do equivalently sized stationary military deployments (i.e., the number of COMSEC devices is a function of the number of mobile war-fighting entities). This can create significant scalability issues within the overlay COMSEC network relationships themselves. COMSEC scaling problems are

manifested in several dimensions. It is important to recognize, however, that just as RANGER recursion was used vertically to create IP Topology enterprise-within-enterprise structures in order to improve routing aggregation and scaling, so RANGER recursion allows for authorization of route-optimized transactions between peer enterprises (within the same IP Topology Hierarchy level) to improve COMSEC aggregation and scaling of the network overlay system. The RANGER use of VET also combines with the Subnetwork Encapsulation and Adaptation Layer (SEAL) to provide robust packet identification and maximum transmission unit (MTU) link adaptation services over tunnels. These capabilities protect against both source address spoofing and black holes caused by MTU limitations.

4.6. Provider Concerns

Network providers must have a way to support the protocol transitions and network types mentioned above and still remain reliable and financially sound. The RANGER architecture provides ways to support general Internet Service Providers (ISPs), cellular operator networks, and specialized networks such as the Aeronautical Telecommunications Network (ATN).

4.6.1. ISP Networks

Internet service provider networks provide a commons for the connection of Customer Premises Equipment (CPE) routers [CPE-RTRS] that connect arbitrarily complex customer networks. This is true whether the ISP permits direct customer-to-customer communications, or whether all communications are forwarded through ISP provider-edge equipment.

The ISP commons must potentially support hundreds of thousands of CPE routers (or more); hence the ISP may be obliged to assign private IPv4 address allocations (i.e., instead of public) as RLOCs for CPE routers. This gives rise to a "nested NATs" scenario, which can increase the overall brittleness brought on by NAT traversal.

To address this brittleness, the ISP can deploy "Carrier-Grade NATs" (CGNs) [INCR-CGN] that provide a second level of RLOC address translation on the path from the CPE to the Internet. When the CGNs are also configured as EBGs, CPE routers can discover them as default routers for reaching EID-based services using the EBG discovery mechanisms specified in VET.

"Scenarios and Analysis for Introducing IPv6 into ISP Networks" [RFC4029] discusses both ISP backbone network and customer connection transition considerations; however, this document considers router-to-router tunneling use cases. Therefore the ISATAP mechanism (which

only supports host-to-router or host-to-host tunneling) is not mentioned as a candidate technology. Early point solutions (e.g., the Tunnel Setup Protocol (TSP) [RFC5572], the Simple IPv6-in-IPv4 Tunnel Establishment Procedure (STEP) [STEP], etc.) were recommended. This document suggests that RANGER, VET, and SEAL would also be suitable solutions in these networks.

4.6.2. Cellular Operator Networks

[RFC4215] provides a (dated) "Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks". It envisions an extended period of support for both IPv4 and IPv6 protocols in the operator network. User Equipment (UE) uses the Packet Data Protocol (PDP) context to establish tunnels through the operator network to a Gateway General Packet Radio Service (GPRS) Support Node (GGSN). RANGER could be used in 3GPP transition; when the UE uses IPv6, and the PDP context is established across an IPv4 provider network, the UE can configure itself as an EBR and contact the GGSN (as a RANGER EBG) through VET tunneling.

Other [RFC4215] scenarios examine IPv4-only UEs, IPv6-only UEs, and various combinations of IPv4 and IPv6 within the operator network. Also to be considered are scenarios in which the UE is configured as a router or bridge that connects an end system such as a laptop computer. In that case, the UE could be the first-hop router/bridge into the cellular provider network, and the laptop computer could be configured as an EBR in the RANGER model. Again, the GGSN or a device reachable through the GGSN could serve as a RANGER EBG.

4.6.3. Aeronautical Telecommunications Network (ATN)

The Aeronautical Telecommunications Network (ATN) is currently based on the OSI and IPv4 protocols and is deployed only in limited areas. The future ATN under consideration within the civil aviation industry will be IPv6-based. The IP variant of ATN is expected to take the form of a worldwide enterprise network that internally comprises an aeronautical-only Internet that has additional external interfaces to the global Internet. Within the ATN, there may be many Air Communications Service Provider (ACSP) and Air Navigation Service Provider (ANSP) networks that are internally organized either as autonomous systems or internets within the ATN, i.e., as depicted in Figure 5. Each of these entities may themselves be further internally subdivided into lower-tier enterprise networks organized as regional, organizational, or functional compartments. It is important to note that while ACSPs and ANSPs within the ATN will share a common objective of safety-of-flight for civil aviation services, enterprise networks may have competing business, social, or political interests that require that components be distinct ASs.

The RANGER principles therefore support collaborative objectives while allowing very diverse local policy distinctions. In this manner, entities that do not trust each other can create collaborative infrastructures to achieve common goals.

Operational associations like this will characterize many future deployments, including the US Department of Defense's Global Information Grid (GIG). In particular, although the routing and addressing arrangements of all enterprise networks require a mutual level of cooperative active management at a certain level, scaling issues, security policy differences, free market forces, organizational differences, political distinctions, or other factors may create internal competition among entities that otherwise share common goals. This will require different enterprise networks within that association to be separated into distinct ASs that are linked within their own functional Internet relationship.

The ATN illustrates transition from OSI protocols to IPv6. It must support mobility (see Section 4.5.1), and it serves many government and private entities that cooperate to provide safe and efficient air travel while often competing with one another. One possible way to meet these needs with RANGER is to create an overlay using IP-in-IP tunneling across the Internet, as illustrated in Figure 14. The aero overlay forms an enterprise network, so that inner packets from ACSP and ANSP edge networks that travel between VET interfaces on EBRs see their passage across the Internet as only one hop.

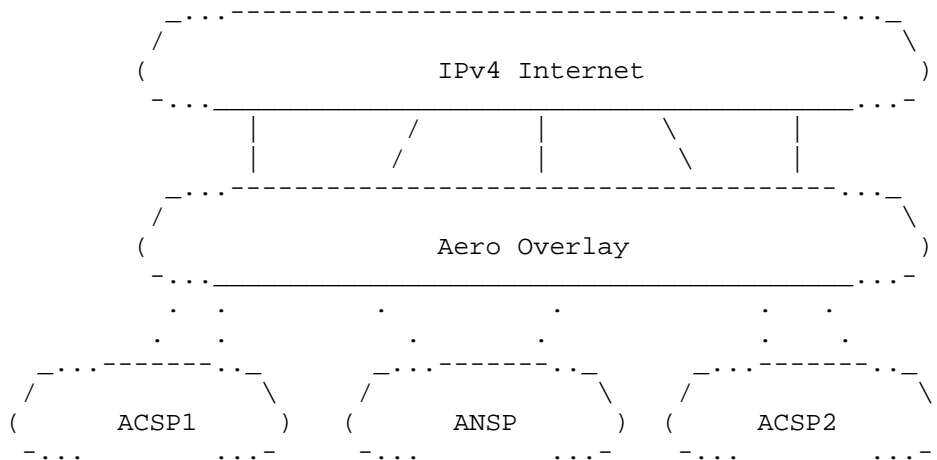


Figure 14. Aeronautical Overlay

Each Aeronautical Communications Service Provider (ACSP), and Aeronautical Navigation Service Provider (ANSP) constitute an enterprise network recursively nested below the aero overlay. Relationships between the various enterprise networks can vary from slight to tight integration. In the example, the ACSP and ANSP might choose to exchange full routing information for their edge networks using a coordinated global-scope RLOC address space across which ACSP and ANSP EBRs can route traffic without further mapping lookups or re-encapsulation at intermediate EBRs. Other enterprise networks that have the aero network as a common parent may not have any knowledge of each other's interior routing but will merely forward packets on a default route up to the aero overlay.

The ATN is currently an OSI network but is projected to transition to IPv6 over time. RANGER can bridge OSI networks together across the IPv4 (or IPv6) Internet, or bridge IPv4 or IPv6 networks across an OSI network. A pair of EBRs that have IP interfaces on a common enterprise network (whether it is the Internet, the aero network, or another parent or child enterprise network) can support communications between their attached OSI edge networks by looking up ISO network service access point (NSAP) addresses [IS8348] instead of IP addresses for RLOC mappings. OSI ConnectionLess Network Protocol (CLNP) [IS8473] packets can therefore be encapsulated within IPv4 (or IPv6) headers for transmission across an Internet Protocol enterprise network. Some OSI networks may transition to IPv6 addressing [RFC4548] while applications are adapted by using RFC 2126 [RFC2126] to carry OSI upper layers over TCP/IP, with the resulting IP packets carried across and between RANGER enterprises in the normal way. Another approach is to use subnetwork convergence to tunnel OSI network protocol data units over Internet Protocol networks [RFC1070].

Figure 15 depicts an ACSP and ANSP connected via an IPv4 aero overlay. Host H represents a system onboard an aircraft that has a wireless link to the ACSP, connected via an enterprise-edge network interface on EBR F within the ACSP enterprise network. H resides on an IPv6 edge network, and its EID is taken from the ACSP IPv6 prefix. H needs to send a query to server S in the ANSP enterprise network. H starts by sending a DNS query to the server at G, and in return it receives the EID of server S. H then creates an IPv6 packet with source EID(H) and destination EID(S) and forwards it to its default router, F. F consults G for a mapping from EID(S) to the appropriate RLOC. In this case, EBR F encapsulates the IPv6 packet in an IPv6 outer packet and forwards the packet to its default EBG, A. A decapsulates the packet and looks up the destination EID(S) by querying the DNS server at EBR B. B returns a mapping with the RLOC of EBR E. A encapsulates the IPv6 inner packet in an IPv4 outer packet with source RLOC(A) and destination RLOC(E). The packet is

forwarded via EBRs C and D in the aero overlay until it reaches E, where it is decapsulated. E consults its cache of EID/RLOC mappings and finds that the EBR for S is N. E encapsulates the packet in an IPv6 packet with source RLOC(E) and destination RLOC(N). When the packet reaches N, it is decapsulated, and the inner IPv6 packet is forwarded on the edge network to the server, S.

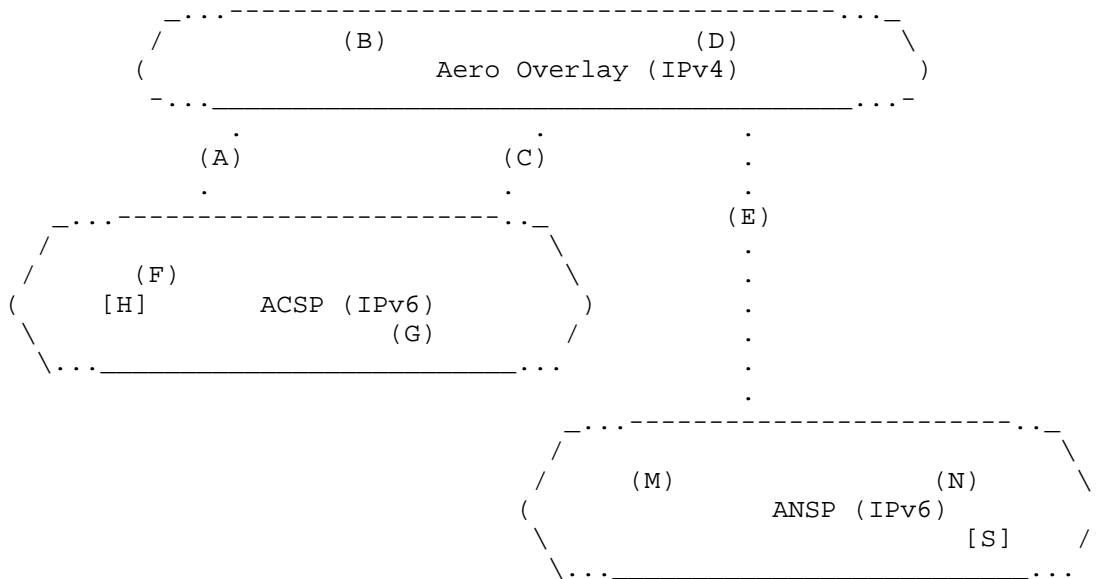


Figure 15. Packet Forwarding for Aeronautical Networks

4.6.4. Unmanaged Networks

"Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks" [RFC3904] considers four cases for support of IPv6-enabled routers and end systems connected to an ISP network via a gateway:

- a. a gateway that does not provide IPv6 at all;
- b. a dual-stack gateway connected to a dual-stack ISP;
- c. a dual-stack gateway connected to an IPv4-only ISP; and
- d. a gateway connected to an IPv6-only ISP.

Case a is typified by the widespread practice of customer networks using IPv4-only NAT boxes to connect to their service providers. RANGER does not address this scenario directly; however, the Teredo mechanism [RFC4380] can provide a sufficient solution in many cases.

Case d is a scenario that has not yet seen widespread adoption. In this scenario, the customer network could be configured as IPv6 only, and the deployment could be considered as an IPv6-only extension to a RANGER enterprise-edge network. End systems in this scenario would still require support for legacy IPv4-only applications, and if the customer network contained IPv4-only routers and end systems the RANGER encapsulation mechanisms would still apply.

Cases b and c correspond to the scenario of the customer gateway to the ISP becoming an IPv6 router. In that case, the gateway could become a RANGER EBR, and the scenario becomes the same as the SOHO network use cases discussed in Section 4.3. In particular, when traditional home network IPv4 NAT boxes are updated to also support IPv6 routing, the NAT box becomes a RANGER EBR.

5. Mapping and Encapsulation Concerns

Mapping and encapsulation concerns related to RANGER have been discussed in Section 3.7 of [RFC5720]. These include effects of mapping systems to application traffic, the need to secure the mapping system, MTU effects, and the ability of legacy Internet networks to connect to those employing RANGER.

6. Problem Statement and Call for Solutions

The scenarios discussed in this document have not closely examined future growth of the native IPv6 and IPv4 Internets independently of any growth in RANGER overlay networking. For example, it is likely that current-day major Internet services that support millions of customers simultaneously (e.g., Google, Yahoo, eBay, Amazon, etc.) will continue to be served best by native Internet routing and addressing rather than by overlay network arrangements that require dynamic mapping state coordination. At the same time, however, more and more small end user networks will wish to use provider-independent addressing for multihoming via multiple ISPs as well as support traffic engineering and mobility management.

These requirements call for an overlay network solution that is compatible with both RANGER and the IPv6 and IPv4 native Internet routing system without adversely affecting Internet routing scaling. The solution must avoid the mapping and encapsulation concerns discussed in Section 3.7 of [RFC5720]; for example, it must provide generally shortest path routing without imparting unacceptable delays for initial packets. The solution must further provide mobility management capabilities for mobile end user networks that can take

advantage of route optimization while requiring no modifications to end systems. Finally, the solution must be based on a business model that allows end user networks to obtain Internet access services from multiple ISPs simultaneously with support for traffic engineering and fault tolerance.

7. Summary

The Internet today can be considered as a giant enterprise network, with nodes in the Internet addressed from the public IPv4 address space as RLOCs. Due to the 32-bit addressing limitations of IPv4, however, continued expansion has occurred through the widespread deployment of IPv4 Network Address Translators (NATs) while IPv6 has yet to see wide adoption.

In many senses, however, this has resulted in a degenerate manifestation of the network-of-networks model originally envisaged, e.g., in the Catenet model. Indeed, these NATed domains have the external appearance of being a simple host within the global Internet RLOC space even though they may be proxying for arbitrarily large networks of end systems. The end result is a loss of transparency in the end-to-end model; it is no longer true that any node in the Internet can directly address any other node.

RANGER enables a true network-within-network (or enterprise-within-enterprise) framework. This is true even across a wide array of deployment scenarios as documented here, and even for networks-within-networks that may be recursively nested to an arbitrary depth. RANGER therefore brings a unifying architecture applied consistently across all layers of recursion, rather than a mixed bag of point solutions that may or may not be mutually compatible. When coupled with an overlay network solution that supports coexistence with the IPv6 and IPv4 native Internet routing systems, a unified future Internet architecture is possible.

8. Security Considerations

Security considerations are addressed in [RFC5720], [RFC5558], and [RFC5320]. While the RANGER architecture does not in itself address security considerations, it proposes an architectural framework for functional specifications that do. Security concerns with tunneling, along with recommendations that are compatible with the RANGER architecture, are found in [TUNNEL-SEC]. Security considerations for specific use cases are discussed there.

9. Acknowledgements

This work was inspired by the original architectural principles of the Internet supplemented with "lessons learned" by many peers from actual Internet deployments and experience developing encapsulation protocols. The editors acknowledge that they are furthering work initiated by many.

10. References

10.1. Normative References

- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC5720] Templin, F., "Routing and Addressing in Networks with Global Enterprise Recursion (RANGER)", RFC 5720, February 2010.

10.2. Informative References

- [APT] Jen, D., Meisel, M., Massey, D., Wang, L., Zhang, B., and L. Zhang, "APT: A Practical Transit Mapping Service", Work in Progress, November 2007.
- [BEHAVE-v6v4] Baker, F., Li, X., Bao, C., and K. Yin, "Framework for IPv4/IPv6 Translation", Work in Progress, August 2010.
- [BELL-LaPADULA] Bell, D. and L. LaPadula, "Secure Computer Systems: Mathematical Foundations and Model", October 1974.
- [CATENET] Pouzin, L., "A Proposal for Interconnecting Packet Switching Networks", May 1974.
- [CPE-RTRS] Singh, H., Beebee, W., Donley, C., Stark, B., and O. Troan, Ed., "Basic Requirements for IPv6 Customer Edge Routers", Work in Progress, December 2010.
- [GROW-VA] Francis, P., Xu, X., Ballani, H., Jen, D., Raszuk, R., and L. Zhang, "FIB Suppression with Virtual Aggregation", Work in Progress, August 2010.

- [HUSTON-END] Huston, G., "The End of the (IPv4) World is Nigher!", July 2007.
- [IEN48] Cerf, V., "The Catenet Model for Internetworking", July 1978.
- [INCR-CGN] Jiang, S., Guo, D., and B. Carpenter, "An Incremental Carrier-Grade NAT (CGN) for IPv6 Transition", Work in Progress, March 2009.
- [IPv4POOL] Hain, T., "The IPv4 Address Pool Projection", April 2009.
- [IS8348] International Organization for Standardization, International Electrotechnical Commission, "Open Systems Interconnection -- Network service definition", 2002.
- [IS8473] International Organization for Standardization, International Electrotechnical Commission, "Protocol for providing the connectionless-mode network service: Protocol specification", 1998.
- [LISP] Farinacci, D., Fuller, V., Meyer, D., and D. Lewis, "Locator/ID Separation Protocol (LISP)", Work in Progress, March 2009.
- [RADIR-PROB-STATE] Narten, T., "On the Scalability of Internet Routing", Work in Progress, February 2010.
- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [RFC1070] Hagens, R., Hall, N., and M. Rose, "Use of the Internet as a subnetwork for experimentation with the OSI network layer", RFC 1070, February 1989.
- [RFC1122] Braden, R., Ed., "Requirements for Internet Hosts - Communication Layers", STD 3, RFC 1122, October 1989.
- [RFC1380] Gross, P. and P. Almquist, "IESG Deliberations on Routing and Addressing", RFC 1380, November 1992.
- [RFC1753] Chiappa, N., "IPng Technical Requirements Of the Nimrod Routing and Addressing Architecture", RFC 1753, December 1994.

- [RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, February 1996.
- [RFC1955] Hinden, R., "New Scheme for Internet Routing and Addressing (ENCAPS) for IPNG", RFC 1955, June 1996.
- [RFC2126] Pouffary, Y. and A. Young, "ISO Transport Service on top of TCP (ITOT)", RFC 2126, March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.
- [RFC2529] Carpenter, B. and C. Jung, "Transmission of IPv6 over IPv4 Domains without Explicit Tunnels", RFC 2529, March 1999.
- [RFC2767] Tsuchiya, K., Higuchi, H., and Y. Atarashi, "Dual Stack Hosts using the "Bump-In-the-Stack" Technique (BIS)", RFC 2767, February 2000.
- [RFC2775] Carpenter, B., "Internet Transparency", RFC 2775, February 2000.
- [RFC3194] Durand, A. and C. Huitema, "The H-Density Ratio for Address Assignment Efficiency An Update on the H ratio", RFC 3194, November 2001.
- [RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.
- [RFC3775] Johnson, D., Perkins, C., and J. Arkko, "Mobility Support in IPv6", RFC 3775, June 2004.
- [RFC3904] Huitema, C., Austein, R., Satapati, S., and R. van der Pol, "Evaluation of IPv6 Transition Mechanisms for Unmanaged Networks", RFC 3904, September 2004.
- [RFC4029] Lind, M., Ksinant, V., Park, S., Baudot, A., and P. Savola, "Scenarios and Analysis for Introducing IPv6 into ISP Networks", RFC 4029, March 2005.
- [RFC4038] Shin, M-K., Ed., Hong, Y-G., Hagino, J., Savola, P., and E. Castro, "Application Aspects of IPv6 Transition", RFC 4038, March 2005.

- [RFC4057] Bound, J., Ed., "IPv6 Enterprise Network Scenarios", RFC 4057, June 2005.
- [RFC4192] Baker, F., Lear, E., and R. Droms, "Procedures for Renumbering an IPv6 Network without a Flag Day", RFC 4192, September 2005.
- [RFC4215] Wiljakka, J., Ed., "Analysis on IPv6 Transition in Third Generation Partnership Project (3GPP) Networks", RFC 4215, October 2005.
- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4380] Huitema, C., "Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)", RFC 4380, February 2006.
- [RFC4472] Durand, A., Ihren, J., and P. Savola, "Operational Considerations and Issues with IPv6 DNS", RFC 4472, April 2006.
- [RFC4548] Gray, E., Rutenmiller, J., and G. Swallow, "Internet Code Point (ICP) Assignments for NSAP Addresses", RFC 4548, May 2006.
- [RFC4795] Aboba, B., Thaler, D., and L. Esibov, "Link-local Multicast Name Resolution (LLMNR)", RFC 4795, January 2007.
- [RFC4852] Bound, J., Pouffary, Y., Klynsmas, S., Chown, T., and D. Green, "IPv6 Enterprise Network Analysis - IP Layer 3 Focus", RFC 4852, April 2007.
- [RFC4862] Thomson, S., Narten, T., and T. Jinmei, "IPv6 Stateless Address Autoconfiguration", RFC 4862, September 2007.
- [RFC5214] Templin, F., Gleeson, T., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", RFC 5214, March 2008.
- [RFC5320] Templin, F., Ed., "The Subnetwork Encapsulation and Adaptation Layer (SEAL)", RFC 5320, February 2010.

- [RFC5558] Templin, F., Ed., "Virtual Enterprise Traversal (VET)", RFC 5558, February 2010.
- [RFC5572] Blanchet, M. and F. Parent, "IPv6 Tunnel Broker with the Tunnel Setup Protocol (TSP)", RFC 5572, February 2010.
- [RFC5579] Templin, F., Ed., "Transmission of IPv4 Packets over Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) Interfaces", RFC 5579, February 2010.
- [RFC5887] Carpenter, B., Atkinson, R., and H. Flinck, "Renumbering Still Needs Work", RFC 5887, May 2010.
- [RFC5944] Perkins, C., Ed., "IP Mobility Support for IPv4, Revised", RFC 5944, November 2010.
- [RFC6115] Li, T., Ed., "Recommendation for a Routing Architecture", RFC 6115, February 2011.
- [STEP] Savola, P., "Simple IPv6-in-IPv4 Tunnel Establishment Procedure (STEP)", Work in Progress, January 2004.
- [TUNNEL-SEC] Krishnan, S., Thaler, D., and J. Hoagland, "Security Concerns With IP Tunneling", Work in Progress, October 2010.

Authors' Addresses

Steven W. Russert (editor)
1078 Ridge Crest Dr.
Wenatchee, WA 98801
USA

EEmail: russerts@hotmail.com

Eric W. Fleischman (editor)
Boeing Research & Technology
P.O. Box 3707 MC 7L-49
Seattle, WA 98124
USA

EEmail: eric.fleischman@boeing.com

Fred L. Templin (editor)
Boeing Research & Technology
P.O. Box 3707 MC 7L-49
Seattle, WA 98124
USA

EEmail: fltemplin@acm.org