NGTRANS Working Group
INTERNET-DRAFT

F. Templin
Nokia
T. Gleeson
Cisco Systems K.K.
M. Lehman
Microsoft

Expires 1 May 2003
1 November 2002

ISATAP Transition Scenario for Enterprise/Managed Networks

draft-ietf-ngtrans-isatap-scenario-01.txt

Status of this Memo

Abstract

This document discusses application of the Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) as a transition tool for enterprise/managed networks. The enterprise/manged network problem space is described, and the ISATAP transition scenario for enterprise/managed network environments is presented.

1. Introduction

The Intra-Site Automatic Tunnel Addressing Protocol [ISATAP] is an NGTRANS mechanism intended for use within arbitrarily large enterprise/managed networks. Examples include corporate intranets and academic campus networks. This document describes the enterprise/managed network problem space and the role of ISATAP within that space.

2. Enterprise/Managed Network Problem Space

Enterprise/managed networks include corporate and academic campus networks (sometimes known as "intranets") that fall under the control of a single administrative authority. The administrative authority may be centralized or distributed, but is ordinarily governed by a commonality of interests and/or policies. These networks typically attach to the global Internet as "stub" networks such that all communications originate and/or terminate internally, i.e., "transit" for foreign traffic is normally blocked by policy restrictions.

Enterprise/managed networks may be arbitrarily large (in both the topological and geographical sense) and may peer with the global Internet at multiple border gateways. They may deploy security mechanisms such as firewalls, proxies, packet filters, etc. to protect intellectual property and other private assets, but the same protocols and services available in the global Internet are typically supported. Enterprise/managed networks are usually constrained by a conservative change/upgrade policy, and in contrast to many service provider networks they are characterized by having a large number of leaf nodes which are often difficult to manage. This makes automation of transition mechanisms critical.

Many nodes (hosts and routers) in existing enterprise/managed networks still communicate using the legacy IPv4 Internet protocol, with IPv4 addresses allocated from either globally assigned prefixes or prefixes from private address spaces. Such networks require transition scenario analysis and transition mechanisms to enable seamless migration to IPv6. In the following sections, we present the anticipated transition scenario for ISATAP in Enterprise/Managed networks and demonstrate ISATAP's applicability for such environments.

3. Enterprise/Managed Network Transition Scenario for ISATAP

[ISATAP] specifies a "simple, scalable approach that enables incremental deployment of IPv6 within IPv4-based sites". The document does not define "site", nor does it place any limits on the topological or geographical scope that a site might cover. But, [ISATAP,2]

("Applicability Statement") and [ISATAP,6] ("ISATAP Deployment Considerations") describe functional and operational aspects of ISATAP that appear to provide a good fit for the enterprise/managed network problem space.

The transition scenario for ISATAP in an enterprise/managed network begins with an administrative decision to enable the service. First, the administrative authority identifies (or deploys) one or more router(s) to carry the ISATAP service. Each such router must configure one or more native IPv4 link(s) and zero or more native IPv6 link(s). An interface for automatic IPv6-in-IPv4 tunneling is configured over one or more IPv4 links that will support ISATAP routing. (Some of these links may also configure IPv6 natively, but this is not required.) Thus, the links configured for IPv6 may include any combination of native IPv6, IPv6-over-IPv4 tunnels, or (in some instances) no IPv6 links at all."

After ISATAP routers have been deployed as described above, the administrative authority for the enterprise/managed network enters the IPv4 address(es) of each ISATAP routing interface into the DNS as described in [ISATAP, 5.2.1]. Following this action, hosts that enable ISATAP will begin to automatically discover ISATAP routers and thus gain access to the global IPv6 network. (Hosts may actually enable ISATAP prior to the administrative deployment of the service, but their ISATAP interfaces will have IPv6 link-local operation only until the first router becomes available.) No other administrative actions are necessary.

4. Applicability

[ISATAP,2] provides an applicability statement that shows direct relevance for enterprise/managed networks. We discuss each aspect of the applicability statement in the following subsections:

4.1. Treats site's IPv4 infrastructure as an NBMA link layer using automatic IPv6-in-IPv4 tunneling (i.e., no configured tunnel state)

No configuration of tunnel endpoints is required - ISATAP is an "automatic tunneling" mechanism whereby the layer 2 (i.e. IPv4) address of other nodes within the ISATAP network is encoded in the layer 3 (i.e. IPv6) address.

IPv6 destinations outside the enterprise/managed network are reached via a router within the enterprise/managed network, the latter being reached by the same ISATAP mechanisms.

Since ISATAP effectively forms an NBMA overlay on the enterprise/managed network, router discovery cannot proceed via standard broadcast discovery mechanisms. Instead, the recommended method is to use DNS resource records to store and distribute the list of routers. (Other mechanisms are also allowed, but not currently specified.)

4.2. **Enables incremental deployment of IPv6 hosts within IPv4 sites** with no aggregation scaling issues at border gateways

Additional hosts can be added with no need for manual configuration (though this is possible, if desired). Such hosts will (when using the recommended mechanism) discover the set of ISATAP routers via a lookup of DNS resource record. These routers are polled (using ISATAP encapsulation) and auto-configuration can be performed, resulting in aggregation efficiency when many hosts configure addresses from prefixes advertised by the routers.

4.3. **Requires no special IPv4 services within the site (e.g., multi-**cast)

IPv4 unicast connectivity within the enterprise/managed network is, of course, required. ISATAP recommends the use of the DNS for establishing essential state, (the list of site ISATAP routers) but apart from this, no other special IPv4 services are required.

4.4. **Supports both stateless address autoconfiguration and manual con-**figuration

Stateless address configuration has many benefits, and ISATAP enables this by the establishment of a list of potential routers in every node within the enterprise/managed network participating in the service.

4.5. **Supports networks that use non-globally unique IPv4 addresses** (e.g., when private address allocations [PRIVATE] are used), but does not allow the virtual ISATAP link to span a Network Address Translator [NAT]

ISATAP uses IPv4 as a layer 2 transport mechanism, but only within the enterprise/managed network itself. Thus the only requirement that ISATAP imposes on these addresses is that they be unique within the local scope - non-global addresses are perfectly usable. Off-site connectivity is achieved via IPv6 routing.

4.6. Compatible with other NGTRANS mechanisms (e.g., [6TO4])

ISATAP encodes the layer 2 (i.e. IPv4) addresses within the interface identifier portion of an IPv6 address, so ISATAP is unconcerned with the higher-order part of an address. Thus ISATAP can be perfectly well used with global unicast addresses in general, and 6to4 addresses in particular.

Two different enterprise/managed networks, both using the same non-globally unique IPv4 addresses internally, and each provided with a single globally-unique IPv4 address for external connectivity through a NAT, can employ 6to4 for external connectivity and ISATAP for internal connectivity.

6to4 encodes the globally-unique IPv4 address (representing the external point of connectivity) within the 6to4 prefix. ISATAP encodes the unique-within-the-site IPv4 address of a node within the interface identifier.

5. IANA Considerations

See [ISATAP, 7].

6. Security Considerations

See [ISATAP, 8].

Acknowledgments

The authors acknowledge Alain Durand, Bob Hinden, and Margaret Wasserman for their helpful comments and/or guidance.

Normative References

[ISATAP] Templin, F., Gleeson, T., Talwar, M., and D. Thaler, "Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)", draft-ietf-ngtrans-isatap-06.txt, (work in progress).

Author's Address:

Fred L. Templin
Nokia
313 Fairchild Drive
Mountain View, CA 94043 USA
Phone: (650)-625-2331
Email: ftemplin@iprg.nokia.com

Tim Gleeson
Cisco Systems K.K.
Shinjuku Mitsu Building
2-1-1 Nishishinjuku, Shinjuku-ku
Tokyo 163-0409, JAPAN
Email: tgleeson@cisco.com

Matthew Lehman
Microsoft
One Microsoft Way
Redmond, WA 98052 USA
Phone: (206)-826-5160
Email: mlehman@microsoft.com