

Intra-area IP unicast among routers over legacy ATM
[<draft-ietf-ion-intra-area-unicast-01.txt>](#)

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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

To learn the current status of any Internet-Draft, please check the ``1id-abstracts.txt'' listing contained in the Internet-Drafts Shadow Directories on ftp.is.co.za (Africa), nic.nordu.net (Europe), munnari.oz.au (Pacific Rim), ds.internic.net (US East Coast), or ftp.isi.edu (US West Coast).

Abstract

This document describes how IP unicast can be efficiently implemented among routers belonging to the same area of a routing domain, where the connectivity is provided by a legacy ATM network as defined by the ATM Forum or ITU. This proposal is designed to be complementary to IP multicast solutions such as the one described in [1].

1. Introduction

This document describes how a set of routers (such as the access/edge routers of an ISP or enterprise) connected to a legacy ATM network can in a dynamic and plug-and-play fashion optimize ATM connections for efficient forwarding of unicast IP packets. The method can be used in situations where the number of routers is so large that a full mesh of point-to-point ATM VCs is not practical from technical or economic reasons. In addition, it can be applied to smaller router networks to automate the setup of a full mesh of ATM connectivity between the routers.

The set of routers must belong to the same area of a link state routing protocol, such as OSPF or IS-IS, that floods topology and reachability information to every router in the area.

This proposal only deals with IP unicast, but it complements and can be used in conjunction with IP multicast solutions such as the one described in [1].

2. Router configuration and behavior

Initialization

As introduced above, this document defines a method of dynamically managing ATM connectivity among a set of routers that belong to the same area of a routing domain, where a link state protocol, such as OSPF or IS-IS, is used to exchange topology and reachability information.

Before the dynamic management of ATM VCs can begin, the routers of the area must be manually configured to exchange routing information among themselves. There must thus be enough initial connectivity so that at least one IP path exists from each router to each other router in the area. This initial connectivity is also used to forward IP packets when dynamic ATM VCs don't exist.

Note that the initial connectivity doesn't necessarily need to be implemented over ATM and that not all routers of the area need to be ATM attached. Furthermore, even if a router is ATM attached, it doesn't need to participate in the dynamic management of ATM VCs. The ATM routers of an area can thus be upgraded one at a time to support the method described in this document.

Connection setup and teardown

After the initial connectivity has been established, ATM attached routers that participate in this method start to dynamically create and delete dynamic shortcut ATM VCs among themselves based on traffic volumes. This can be accomplished, for example, as follows.

An ATM attached router R measures, how many bytes it has received during the past M seconds, whose final hop router S within the area is also ATM attached. If the number of bytes is less than N, R forwards the packets according to its routing table. When the number of bytes equals or exceeds N, and R doesn't yet have a dynamic ATM VC to S, R creates such a VC and starts to forward S bound packets directly.

Once the dynamic ATM VC from R to S has been created, R starts to

measure the traffic along it. When R detects that during the past K seconds the number of bytes along the dynamic ATM VC to S has fallen below L, it deletes the dynamic ATM VC and returns to the initial mode of operation that was described above.

The values of the constants K, L, M, and N control the rate of dynamic ATM VC creation and teardown. They are assigned by the network administrator and may differ from one ATM attached router to another.

Setting the VC creation and deletion limits N and L to zero, turns off the measurement process and causes the router to create a dynamic VC to every other participating router. That can be the default in small router networks that want to use this method to automate the setup of a full mesh of ATM VCs.

If a router doesn't want to set up any dynamic ATM VCs, the VC creation limit N is set to a positive value and the measurement interval M is set to zero. Finally, if a router doesn't want to be a destination of dynamic ATM VCs, it doesn't make its ATM address available to the other routers for the purpose of this application.

Note that if a router is not capable in measuring traffic, it can still participate as a destination of dynamic ATM VCs and can itself set up dynamic VCs non-selectively to every other router.

Characteristics of dynamic ATM VCs

In order to keep the number of routing peers small and in order to avoid frequent changes in topology information, the router that establishes a dynamic ATM VC does not use it for the exchange of routing information nor does it advertise the dynamic VC to its routing peers.

Dynamic ATM VCs are unidirectional, because the source router that established a dynamic ATM VC does not have information about the traffic volumes to the reverse direction. Unidirectionality also simplifies the method, since it allows the source router to manage the dynamic ATM VC autonomously without coordination with the destination router. Yet another advantage of unidirectionality is that unidirectional VCs can be merged if more than one source router sets up a connection to the same destination router.

The traffic category, traffic parameters, and protocol encapsulation of dynamic ATM VCs are a local matter of the routers that establish them. The default traffic category is UBR with peak cell rate set to the link rate and minimal acceptable cell rate (if applicable) set to zero. The default protocol encapsulation method is LLC Encapsulation

as defined in [2]. Signalling is as specified in [3] for UNI 3.1 and in [4] for UNI 4.0.

3. Address resolution

Since all the routers belong to the same area of a link state routing domain, they learn each others' router IDs and the IP address prefixes that are reachable via each router. In order to dynamically create an ATM VC from one router to another, the source router also needs to learn the ATM address of the destination router.

A router that wants to participate as a destination in the dynamic management of ATM VCs, makes its ATM address known to the other routers of the area by including in its link state advertisements a field that contains an ATM address of the advertising router.

In OSPF, the router advertises its ATM address(es) using the Address Resolution Advertisement (ARA) option [5]. The Opaque Type of the ARA is Intra-area Router ARA (Opaque Type-1). One or more ATM addresses or LIJ identifications (one per Vertex Association) can be advertised in a single ARA. The Resolution Type of a Vertex Association is ATM Address, if the Link Service Type is ATM Point-To-Point, or ATM LIJ Call Identification, if the Link Service Type is ATM MultiPoint-To-Point. See [5] for further details regarding the ARA option.

A field similar to Opaque LSA could be easily defined for IS-IS. Further, it could be possible to use a well-known discretionary non-transitive attribute of BGP to carry the address resolution information, but the use of inter-domain routing protocols is outside the scope of this document.

4. Discussion

The method proposed in this document allows efficient interconnection of a set of routers over a legacy ATM network. After small amount of manual configuration, the routers will automatically optimize direct connectivity among themselves based on dynamic traffic load. Network administrators can control the number of ATM VCs created by the method taking into account scalability and cost.

As shown above, the method can readily exploit a multipoint-to-point ATM signalling capability, which will reduce the number of ATM VCs terminating at the destination routers. The method also benefits from the capability to dynamically renegotiate the traffic parameters of active ATM VCs. Both of these new capabilities are currently under study in the ATM Forum.

5. Security Considerations

Since the method described in this document allows data paths to be established that bypass the normal hop-by-hop control path, the location of any access filters should be decided carefully. To ensure proper enforcement of filter policies, filters should be moved to the edges of an area so that they may be applied on entry or exit from the short-cut data path.

References

- [1] Farinacci, D., Meyer, D., and Rekhter, Y., "Intra-LIS IP multicast among routers over ATM using Sparse Mode PIM". [draft-ietf-ion-intralis-multicast-01.txt](#), August 1997.
- [2] Heinanen, J., "Multiprotocol Encapsulation over ATM Adaptation Layer 5". [RFC 1483](#), July 1993.
- [3] Perez, M, et al., "ATM Signalling Support for IP over ATM". [RFC 1755](#), February 1995.
- [4] Maher, M, "ATM Signalling Support for IP over ATM - UNI Signalling 4.0 Update". [draft-ietf-ion-sig-uni4.0-04.txt](#), May 1997.
- [5] Coltun, R. and Heinanen, J., "The OSPF Address Resolution Advertisement Option". Internet Draft, November 1997.

Acknowledgements

I would like to thank Rob Coltun and Lou Berger of Fore Systems for their comments on earlier versions of this document.

Author Information

Juha Heinanen
Telia Finland, Inc.
Myrmaentie 2
01600 VANTAA
Finland

Phone +358 303 994 808
Email jh@telia.fi

