Internet Engineering Task Force
INTERNET DRAFT

T. Przygienda/P. Droz
Bell Labs/IBM
5 March 1998

OSPF over ATM and Proxy PAR
<draft-ietf-ospf-atm-01.txt>

Status of This Memo

This document is an Internet Draft, and can be found as draft-ietf-ospf-atm-01.txt in any standard internet drafts repository. 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. Internet Drafts may be updated, replaced, or obsoleted by other documents at any time. It is not appropriate to use Internet Drafts as reference material, or to cite them other than as a ``working draft'' or ``work in progress.''
Please check the I-D abstract listing contained in each Internet Draft directory to learn the current status of this or any other Internet Draft.

Abstract

This draft specifes for OSPF implementors and users mechanisms describing how the protocol operates in ATM networks over PVC and SVC meshes with the presence of Proxy PAR. These recommendations do not require any protocol changes and allow for simpler, more efficient and cost-effective network designs. It is recommended that OSPF implementations should be able to support logical interfaces, each consisting of one or more virtual circuits and used as numbered logical point-to-point links (one VC) or logical NBMA networks (more than one VC) where a solution simulating broadcast interfaces is not appropriate. PAR can help to distribute configuration changes of such interfaces when OSPF capable routers are reconfigured on the ATM cloud.

1. Introduction

1.1. Introduction to PAR

PNNI Augmented Routing (PAR) [For98] is an extension to PNNI [AF96b] routing to allow information about non-ATM services to be distributed in an ATM network as part of the PNNI topology. The content and format of the information is specified by PAR but is transparent to PNNI routing. A PAR-capable device, one that implements PNNI and the PAR extension, is able to create PAR PTSEs that describe the non-ATM services located on or behind that device. Because this information is flooded by PNNI routing, PAR-capable devices are also able to examine the PAR PTSEs in the topology database that were originated by other nodes to obtain information on desired services reachable through the ATM network. An important example of how PAR can be used is provided by overlay routing on ATM backbones. If the routers are PAR-capable, they can create PTSEs to advertise the routing protocol supported on the given interface (e.g., OSPF, RIP, or BGP), along with their IP address and subnet, and other protocol-specific details. The PAR-capable routers can also automatically learn about "compatible" routers (e.g., supporting the same routing protocol, in the same IP subnet) active in the same ATM network. In this manner the overlay routing network can be established automatically on an ATM backbone. The mechanism is dynamic, and does not require configuration. One potential drawback of PAR is that a device must implement PNNI in order to participate. Therefore an additional set of optional protocols called Proxy PAR has been defined to allow a client that is not PAR-capable to interact with a server that is PAR-capable and thus obtain the PAR capabilities. The server acts as a proxy for the client in the operation of PAR. The client is able to register its own services, and query the server to obtain information on compatible services available in the ATM network. A key feature of PAR and Proxy PAR is the ability to provide VPN support in a simple yet very effective manner. All PAR information is tagged with a VPN ID and can therefore be filtered on that basis. This can be used for example, in a service provider network. Each customer can be provided with a unique VPN ID that is part of all Proxy PAR registrations and queries. Usage of the correct VPN ID can easily be enforced at the Proxy PAR server. In this way the services of a given customer will be available only to clients in that customer's network.

1.1.1. Overview of PNNI Augmented Routing (PAR)

PNNI Augmented Routing (PAR) is an extension to PNNI to allow the flooding of information about non-ATM devices. PAR uses a new PTSE

type to carry this non-ATM-related information. The current version of PAR specifies IGs for the flooding of IPv4-related protocol information such as OSPF or BGP. In addition PAR also allows the use of the System Capabilities IG, which can be used to carry proprietary or experimental information.

PAR supports extensive filtering possibilities, which allow the implementation of virtual private networks (VPN). As PAR is a PNNI extension, it can reuse existing PNNI routing level scopes. In addition, PAR provides filtering in terms of a VPN ID, IP address, including a subnet mask, as well as protocol flags. The correct filtering according to these parameters is part of a PAR implementation.

1.1.2. Overview of Proxy PAR

Proxy PAR is a protocol that allows for different ATM attached devices to interact with PAR-capable switches and obtain information about non-ATM services without executing PAR themselves. The client side is much simpler in terms of implementation complexity and memory requirements than a complete PAR instance and should allow easy implementation in, for example, existing IP routers. Clients can use Proxy PAR to register different non-ATM services and protocols they support. This protocol has deliberately not been included as part of ILMI [AF96a] owing to the complexity of PAR information passed in the protocol and the fact that it is intended for integration of non-ATM protocols and services only. A device executing Proxy PAR does not necessarily need to execute ILMI or UNI signaling, although this will normally be the case.
The protocol does not specify how the distributed service registration and data delivered to the client are supposed to drive other protocols. For example, OSPF routers finding themselves through Proxy PAR could use this information to form a full mesh of P2P VCs and communicate using RFC1483 [Hei93] encapsulation. In terms of the discovery of other devices such as IP routers, Proxy PAR is an alternative to LANE [AF95] or MARS [Arm96]. It is expected that the guidelines defining how a certain protocol can make use of Proxy PAR and PAR should come from the group or standardization body that is responsible for the particular protocol.

PAR and Proxy PAR have the ability to provide ATM address resolution for IP attached devices, but such resolution can also be achieved by other protocols under specification in IETF e.g. [CH97a, CH97b]. However, the main purpose of the protocol is to allow the automatic

detection of devices over an ATM cloud in a distributed fashion, not relying on a broadcast facility. Finally, it should be mentioned that the protocol complements and coexists with server detection via ILMI extensions.

1.2. Introduction to OSPF

OSPF (Open Shortest Path First) is an Interior Gateway Protocol (IGP) and described in [Moy94, Moy97] from which most of the following paragraphs has been taken almost literally. OSPF distributes routing information between routers belonging to a single Autonomous System. The OSPF protocol is based on link-state or SPF technology. It was developed by the OSPF working group of the Internet Engineering Task Force. It has been designed expressly for the TCP/IP internet environment, including explicit support for IP subnetting, and the tagging of externally-derived routing information. OSPF also utilizes IP multicast when sending/receiving the updates. In addition, much work has been done to produce a protocol that responds quickly to topology changes, yet involves small amounts of routing protocol traffic.

To cope with the needs of NBMA and demand circuits capable networks such as Frame Relay or X.25, [Moy95] has been made available that standardizes extensions to the protocol allowing for efficient operation over on-demand circuits.
OSPF supports three types of networks today:

- Point-to-point networks: A network that joins a single pair of routers. Point- to-point networks can either be numbered or unnumbered in the latter case the interfaces do not have IP addresses nor masks. Even when numbered, both sides of the link do not have to agree on the IP subnet.

- Broadcast networks: Networks supporting many (more than two) attached routers, together with the capability to address a single physical message to all of the attached routers (broadcast). Neighboring routers are discovered dynamically on these nets using OSPF's Hello Protocol. The Hello Protocol itself takes advantage of the broadcast capability. The protocol makes further use of multicast capabilities, if they exist. An Ethernet is an example of a broadcast network.

- Non-broadcast networks: Networks supporting many (more than two) attached routers, but having no broadcast capability.

Neighboring routers are maintained on these nets using OSPF's Hello Protocol. However, due to the lack of broadcast capability, some configuration information is necessary for the correct operation of the Hello Protocol. On these networks, OSPF protocol packets that are normally multicast need to be sent to each neighboring router, in turn. An X.25 Public Data Network (PDN) is an example of a non-broadcast network.

OSPF runs in one of two modes over non-broadcast networks. The first mode, called non-broadcast multi-access (NBMA), simulates the operation of OSPF on a broadcast network. The second mode, called Point-to-MultiPoint, treats the non-broadcast network as a collection of point-to-point links. Non-broadcast networks are referred to as NBMA networks or Point-to-MultiPoint networks, depending on OSPF's mode of operation over the network.

2. OSPF over ATM

2.1. Model

Contrary to broadcast-simulation based solutions such as LANE [AF95] or RFC1577 [Lau94], this RFC elaborates on how to handle virtual OSPF interfaces over ATM such as NBMA, point-to-multipoint or point-to-point and allow for their auto-configuration in presence of Proxy PAR. One advantage is the circumvention of server solutions that often present single points of failure or hold large amounts of configuration information. The other main benefit is the possibility to execute OSPF on top of partially meshed VC topologies.

Parallel to [dR94] that describes the recommended operation of OSPF over Frame Relay networks, a similar model is assumed where the underlying ATM network can be used to model single VCs as point-to-point interfaces or collections of VCs can be accessed as an non-broadcast interface in NBMA or point-to-multipoint mode. Such a VC or collection of VCs is called a logical interface and specified through its type (either point-to-point, NBMA or point-to-point), IP instance (presenting an incarnation of IP with its own address spaces), address and mask. Layer 2 specific configuration such as address resolution method, class and quality of service of used circuits and other must be also included. As logical consequence thereof, a single, physical interface could encompass multiple IP subnets or even multiple, independent IP instances. In contrary to layer 2 and IP addressing information, when running Proxy PAR, most of the OPSF information needed to operate such a logical interface

does not have to be configured into routers statically but can be provided through Proxy PAR queries. This allows for much more dynamic configuration of VC meshes in OSPF environments than e.g. in Frame Relay solutions.

2.2. OSPF Configuration Interaction with Proxy PAR

To achieve the goal of simplification of VC mesh reconfiguration, Proxy PAR allows the router to learn automatically most of the configuration that has to be provided to OSPF. Non-broadcast and point-to-point interface information can be learned across an ATM cloud as described in the ongoing sections. It is up to the implementation to possibly allow for a mixture of Proxy PAR autoconfiguration and manual configuration of neighbor information. Moreover, manual configuration could e.g. override or complement information derived from a proxy PAR client. Additionally, OSPF extensions to handle on-demand circuits [Moy95] can be used to allow for graceful tearing down of VCs not carrying any OSPF traffic over prolonged periods of time. The different interactions are described in sections 2.2.1, 2.2.2 and 2.2.3.

Even after autoconfiguration of interfaces has been provided, the problem of VC setups in an ATM network is unsolved since none of the normally used mechanisms such as RFC1577 [Lau94] or LANE [AF95] are assumed to be present. Section 2.5 describes the behavior of OSPF routers to allow for router connectivity necessary.

2.2.1. Autoconfiguration of Non-Broadcast Interfaces

Proxy PAR allows to autoconfigure the list of all routers residing on the same IP network in the same IP instance by simply querying the Proxy PAR server. Each router can easily obtain the list of all OSPF routers on the same subnet with their router priorities and ATM address bindings. This is the precondition for OSPF to work properly across such logical NBMA interfaces. Note that the memberlist, when learned through Proxy PAR queries, can dynamically change with PNNI (in)stability and general ATM network behavior. It maybe preferable for an implementation to withdraw list membership e.g. much slower than detect new members. Relying on OSPF mechanism to discover lack of reachability in the overlaying logical IP network could alleviate the risk of thrashing DR elections and excessive information flooding. Once the DR registration is completed and the router has not been elected DR or BDR, an implementation of [Moy95] can ignore the fact that all routers on the specific NBMA subnet are

available in its configuration since it only needs to maintain VCs to the DR and BDR.

Traditionally, router configuration for a NBMA network provides the list of all neighboring routers to allow for proper protocol operation. For stability purposes, the user may choose to provide a list of neighbors through such static means but additionally enable the operation of Proxy PAR protocol to complete the list. It is left to specific router implementations whether the manual configuration is used in addition to the information provided by Proxy PAR, used as filter of the dynamic information or whether a concurrent mode of operation is prohibited. In any case it should be obvious that allowing for more flexibility may facilitate operation but provides more possibilities for misconfiguration as well.

2.2.2. Autoconfiguration of Point-to-Multipoint Interfaces

Point-to-Multipoint interfaces in ATM networks only make sense if no VCs can be dynamically set up since an SVC-capable ATM network normally presents a NBMA cloud to OSPF. This is e.g. the case if the intended use of the network is only to execute OSPF in presence of a partial PVC or SPVC mesh or pre-determined SVC meshes. Such a collection could be modeled using the point-to-multipoint OSPF interface and the neighbor detection could be provided by Proxy PAR or other means. In Proxy PAR case the router queries for all OSPF routers on the same network in the same IP instance but it installs in the interface configuration only routers that are already reachable through preset PVCs. The underlying assumption is that a router understands the remote NSAP of a PVC and can compare it with appropriate Proxy PAR registrations. If the remote NSAP of the PVC is unknown, alternative autodiscovery mechanisms have to be used e.g. inverse ARP [BB92, LH96].

2.2.3. Autoconfiguration of Numbered Point-to-Point Interfaces

OSPF point-to-point links do not necessarily have an IP address assigned and even when having one, the mask is undefined. As a precondition to successfully register a service with Proxy PAR, IP address and mask is required. Therefore, if a router desires to use Proxy PAR to advertise the local end of a point-to- point link to the router it intends to form an adjacency with, an IP address has to be provided and a netmask set or a default of 255.255.255.254 (this gives as the default case a subnet with 2 routers on it) assumed. To allow the discovery of the remote end of the interface, IP address

of the remote side has to be provided and a netmask set or a default of 255.255.255.254 assumed. Obviously the discovery can only be successfull when both sides of the interface are configured with the same network mask and are within the same IP network. The situation where more than two possible neighbors are discovered through queries and the interface type is set to point-to-point presents a configuration error.

2.2.4. Autoconfiguration of Unnumbered Point-to-Point Interfaces

For reasons given already in [dR94] using unnumbered point-to-point interfaces with Proxy PAR is not a very attractive alternative since the lack of an IP address prevents efficient registration and retrieval of configuration information. Relying on the numbering method based on MIB entries generates conflicts with the dynamic nature of creation of such entries and is beyond the scope of this work.

2.3. Proxy PAR Interaction with OSPF Configuration

To allow other routers to discover an OSPF interface automatically, the IP address, mask, Area ID, interface type and router priority information given must be registered with the Proxy PAR server at an appropriate scope. A change in any of these parameters has to force a reregistration with Proxy PAR.

It should be emphasized here that since the registration information can be used by other routers to resolve IP addresses against NSAPs as explained in section 2.4 already, whole IP address of the router must be registered. It is not enough to just indicate the subnet up to the mask length but all address bits must be provided.

2.3.1. Registration of Non-Broadcast Interfaces

For an NBMA interface the appropriate parameters are available and can be registered through Proxy PAR without further complications.

2.3.2. Registration of Point-to-Multipoint Interfaces

In case of a point-to-multipoint interface the router registers its information in the same fashion as in the NBMA case except that the interface type is modified accordingly.

2.3.3. Registration of Point-to-Point Interfaces

In case of point-to-point numbered interfaces the address mask is not specified in the OSPF configuration. If the router has to use Proxy PAR to advertise its capability, a mask must be defined or a default value of 255.255.255.254 used.

2.3.4. Registration of Unnumbered Point-to-Point Interfaces

Due to the lack of a configured IP address and difficulties generated by this fact as described earlier, registration of unnumbered point-to-point interfaces is not covered in this document.

2.4. IP address to NSAP Resolution Using Proxy PAR

As a byproduct of Proxy PAR presence, an OSPF implementation could use the information in registrations for the resolution of IP addresses to ATM NSAPs on a subnet without having to use static data or mechanisms such as ATMARP [LH96]. This again should allow for drastic simplification of number of mechanisms involved in operation of OSPF over ATM to provide an IP overlay.

2.5. Connection Setup Mechanisms

This sections describes OSPF behavior in an ATM network under different assumptions in terms of signaling capabilities and preset connectivity.

2.5.1. OSPF in PVC Environments

In environments where only partial PVCs (or SPVCs) meshes are available and modeled as point-to-multipoint interfaces, the routers see reachable routers through autodiscovery provided by Proxy PAR. This leads to expected OSPF behavior. In cases where a full mesh of PVCs is present, such an interface should preferably be modeled as broadcast and Proxy PAR discovery should be superfluous.

2.5.2. OSPF in SVC Environments

In SVC-capable environments the routers can initiate VCs after having discovered the appropriate neighbors, preferably driven by the need to send data such as Hello-packets. Since this can lead to race conditions where both sides can open a VC and it is desirable to minimize this valuable resource, if the router with lower Router ID

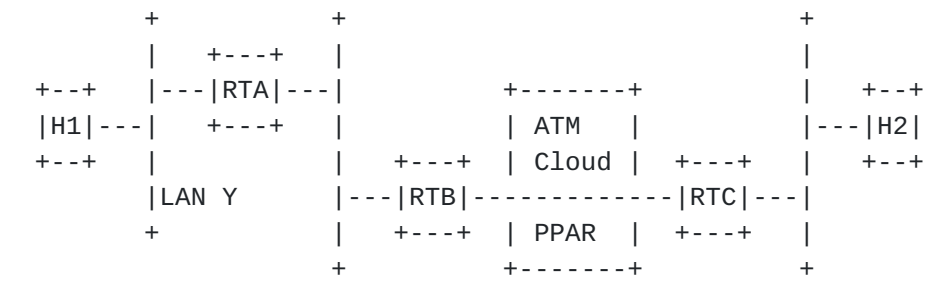

Figure 1: Simple Topology with Router B and Router C operating across NBMA ATM interfaces with Proxy PAR

detects that the VC initiated by the other side is bidirectional, it is free to close its own VC and use the detected one. Observe that this behavior operates correctly in case OSPF over Demand Circuits extensions are used [Moy95] over SVC capable interfaces.

The existence of VCs used for OSPF exchanges is orthogonal to the number and type of VCs the router chooses to use within the logical interface to forward data to other routers. OSPF implementations are free to use any of these VCs (1) to send packets if their endpoints are adequate and must accept hello packets arriving on any of the VCs belonging to the logical interface even if OSPF operating on such an interface is not aware of their existence. An OSPF implementation may not accept or close connections being initiated by another router that has either not been discovered by Proxy PAR or whose Proxy PAR registration is indicating that it is not adjacent.
As an example consider the topology in figure 2.5.2 where router RTB and RTC are connected to a common ATM cloud offering Proxy PAR services. Assuming that RTB's OSPF implementation is aware of SVCs initiated on the interface and RTC only makes minimal use of Proxy PAR information the following sequence could develop illustrating some of the cases described above:

1. RTC and RTB register with ATM cloud as Proxy PAR capable and discover each other as adjacent OSPF routers.

---

1. in case they are aware of their existence

2. RTB sends a hello which forces it to establish a SVC connection to RTC.

3. RTC sends a hello to RTB but disregards the already existing VC and establishes a new VC to RTB to deliver the packet.

4. RTB sees a new bi-directional VC and assuming here that RTC's OSPF Id is higher, closes the VC originated in step 2.

5. Host H1 sends data to H2 and RTB establishes a new data SVC between itself and RTC.

6. RTB sends a Hello to RTC and decides to do it using the newly establish data SVC. RTC must accept the hello despite the minimal implementation.

3. Acknowledgments

Comments and contributions from several sources, especially Rob Coltun, Doug Dykeman and John Moy are included in this work.

4. Security Consideration

Several aspects are to be considered when talking about security of operating OSPF over ATM and/or Proxy PAR. The security of registered information handed to the ATM cloud must be guaranteed by the underlying PNNI protocol. Extensions to PNNI are available and given their implementation spoofing of registrations and/or denial-of-service issues can be addressed [PB97]. The registration itself through proxy PAR is not secured and appropriate mechanisms are for further study. However, even if the security at the ATM layer is not guaranteed, OSPF security mechanisms can be used to verify that detected neighbors are authorized to interact with the entity discovering them.

References

[AF95] ATM-Forum. LAN Emulation over ATM 1.0. ATM Forum af-lane-0021.000, January 1995.

[AF96a] ATM-Forum. Interim Local Management Interface (ILMI) Specification 4.0. ATM Forum 95-0417R8, June 1996.

[AF96b] ATM-Forum. Private Network-Network Interface Specification Version 1.0. ATM Forum af-pnni-0055.000, March 1996.

[Arm96] G. Armitage. Support for Multicast over UNI 3.0/3.1 based ATM Networks, RFC 2022. Internet Engineering Task Force, November 1996.

[BB92] T. Bradley and C. Brown. Inverse Address Resolution Protocol, RFC 1293. Internet Engineering Task Force, January 1992.

[CH97a] R. Coltun and J. Heinanen. Opaque LSA in OSPF. Internet Draft, 1997.

[CH97b] R. Coltun and J. Heinanen. The OSPF Address Resolution Advertisement Option. Internet Draft, 1997.

[dR94] O. deSouza and M. Rodrigues. Guidelines for Running OSPF Over Frame Relay Networks, RFC 1586. Internet Engineering Task Force, March 1994.

[For98] ATM Forum. PNNI Augmented Routing (PAR) Version 1.0. ATM Forum PNNI-RA-PAR-01.04, 1998.

[Hei93] J. Heinanen. Multiprotocol Encapsulation over ATM Adaptation Layer 5, RFC 1483. Internet Engineering Task Force, July 1993.

[Lau94] M. Laubach. Classical IP and ARP over ATM, RFC 1577. Internet Engineering Task Force, January 1994.

[LH96] M. Laubach and J. Halpern. Classical IP and ARP over ATM. Internet Draft, 1996.

[Moy94] J. Moy. OSPFv2, RFC 1583. Internet Engineering Task Force, March 1994.

[Moy95] J. Moy. Extending OSPF to Support Demand Circuits, RFC 1793. Internet Engineering Task Force, April 1995.

[Moy97] J. Moy. OSPFv2, RFC 2178. Internet Engineering Task Force, July 1997.

[PB97] T. Przygienda and C. Bullard. Baseline Text for PNNI Peer Authentication and Cryptographic Data Integrity. ATM Forum 97-0472, July 1997.

Authors' Addresses

Tony Przygienda
Bell Labs, Lucent Technologies
101 **Crawfords Corner Road**
Holmdel, NJ 07733-3030
prz@dnrc.bell-labs.com

Patrick Droz
IBM Research Division
Saumerstrasse 4
8803 **Ruschlikon**
Switzerland
dro@zurich.ibm.com