

Internetworking Over NBMA
INTERNET-DRAFT
<[draft-ietf-ion-scsp-atmarp-01.txt](#)>

James V. Luciani
(Bay Networks)
Joel Halpern
(Newbridge Networks, Inc)
Barbara A. Fox
(Lucent Technologies)
Expires May 1999

A Distributed ATMARP Service Using SCSP

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 ``[l1d-abstracts.txt](#)'' listing contained in the Internet-Drafts Shadow Directories on [ftp.is.co.za](#) (Africa), [ftp.nordu.net](#) (Northern Europe), [ftp.nis.garr.it](#) (Southern Europe), [munari.oz.au](#) (Pacific Rim), [ftp.ietf.org](#) (US East Coast), or [ftp.isi.edu](#) (US West Coast).

Abstract

This document describes a method for distributing an ATMARP service within a LIS[1]. This method uses the Server Cache Synchronization Protocol (SCSP)[2] to synchronize the ATMARP server databases within a LIS. When SCSP is used to synchronize the caches of ATMARP servers in a LIS, the LIS defines the boundary of an SCSP Server Group (SG).

1. Introduction

An ATMARP Client implicitly registers (e.g., by sending an ATMARP Request for itself; see [1]) and refreshes its own ATMARP information with a single ATMARP server in its atm\$arp-req-list table. In addition, the ATMARP Client uses the ATMARP service to gain access to and re-validate ATMARP information about other ATMARP Clients in its

INTERNET-DRAFT

SCSP

Expires May 1999

Logical IP Subnet (LIS). Since, there MAY be multiple ATMARP servers in a given LIS, and since any ATMARP server within the LIS MUST be able to reply to ATMARP requests for ATMARP information about any ATMARP Clients within the LIS, there MUST be a method by which to synchronize ATMARP information across all ATMARP Servers within the LIS. The Server Cache Synchronization Protocol (SCSP) solves the generalized server synchronization/cache-replication problem for distributed databases, and thus SCSP MAY be applied to the ATMARP server database synchronization problem with the LIS. When SCSP is used to synchronize the caches of ATMARP servers in a LIS, the LIS defines the boundary of an SCSP Server Group (SG).

SCSP is defined in two parts: the protocol independent part and the client/server protocol specific part. The protocol independent part is specified in [2] whereas this document will specify the client/server protocol specific part where ATMARP is the client/server protocol.

[2.](#) Overview

All ATMARP servers belonging to a Logical IP Subnet (LIS)[1] are said to belong to a Server Group (SG). An SG is identified by, not surprisingly, its SGID which is contained in a field in all SCSP packets. All SCSP packets contain a Protocol ID (PID) field as well. This PID field is set to 0x0001 to signify that SCSP is synchronizing ATMARP server databases as opposed to synchronizing some other protocol's databases (see Section B.2.0.1 of [2] for more details). In general, PIDs for SCSP will be assigned by IANA upon request given that a client/server protocol specific specification has been written. In the case of ATMARP, the client/server protocol specific specification was initially written at the same time as SCSP, and thus a PID=0x0001 was assigned by the author.

SCSP places no topological requirements upon an ATMARP SG. Obviously, however, the resultant graph of ATMARP servers must span the set of ATMARP servers to be synchronized. For more information about the client/server protocol independent part of SCSP, the reader is encouraged to see [2].

When an ATMARP SG is using SCSP for synchronization, a given ATMARP Client will use only one ATMARP server and it will use that server for remainder of its participation in the SG. This server is said to

be the "serving ATMARP server." There needs to be some hysteresis on refreshes since every ATMARP Request may cause a cache update/refresh in the serving ATMARP Server, and such refreshes might cause excessive traffic if propagated to all ATMARP Servers in the SG. In the case of mere refreshes, where no change occurs to the ATMARP

Server's cache entry for the ATMARP Client, SCSP updates will occur at a maximum rate of once every $10 + \text{Random}(2)$ minutes.

When an ATMARP server receives database information via SCSP, it checks it against the locally registered clients. Information not related to locally registered clients is simply accepted. (Note to implementors: this may result in transient multiple possible resolutions for an IP->ATM address binding. The server may provide any one of these bindings to a client who sends in an ATMARP request.) If there is a collision with the locally registered client base, then the base must be checked for whether the information is associated with a connected client. If it is not, then the local information is presumed to be supersceded and must be purged and transmitted to peer ATMARP servers with 0 lifetime.

If the local information is associated with a connected client, then an effort needs to be made to determine if the remote information is also associated with a connected client. Therefore:

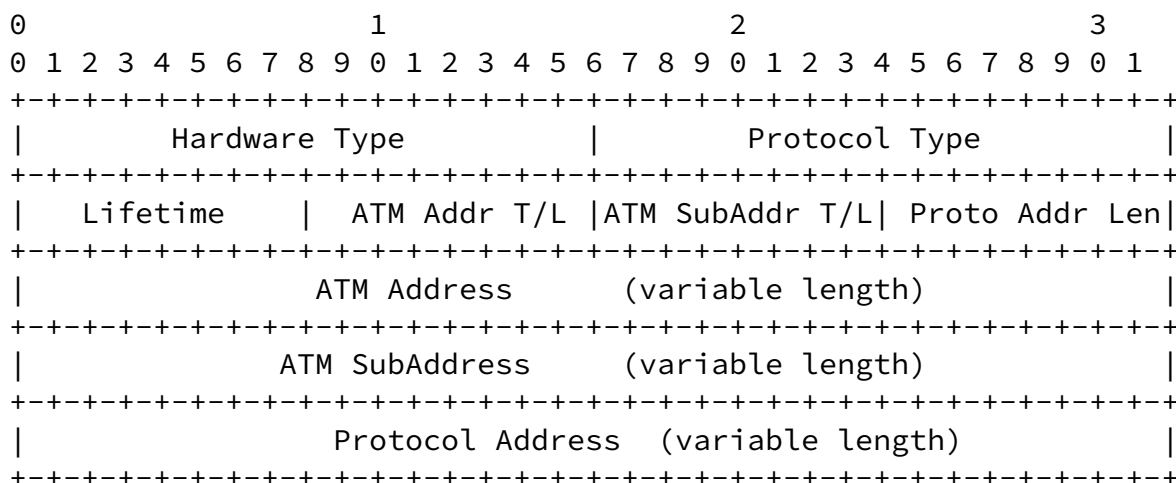
- 1) The ATMARP server will wait an interval T1 before taking any action. If the duplicative information is purged, no further action is necessary.
- 2) If the duplicate information is not purged, the ATMARP Server will then consider both sets of information (local and remote) to be invalid. A Network Management notification is generated.
- 3) The ATMARP Server will disconnect the local client. An ARPNAK should be sent to the client prior to disconnecting the VC. After waiting an interval T2, the ATMARP Server will purge the locally generated information from the system. (This avoids the potential for one of the two clients to randomly remain in the system when both are wrong.)

When an ATMARP server receives a registration from a client, it checks the database of known clients. If there is a known local collision, the procedures from Classic II are followed. If there is a collision with an entry in the remote database, then the registration is accepted and flooded out, and the timer from (1)

above is started so as to resolve the apparent collision.

3. Format of the CSA Record ATMARP Specific Part

CSA Records in SCSIP contain a "Client/Server Protocol Specific Part" which contains the non-protocol independent information for a given server's cache entry.



With the exception of the State and unused fields, these fields contain the values specified in the ATMARP Request and Reply packets defined in [1] which are used to create, update, and access ATMARP server cache entries.

Hardware Type

Defines the type of "link layer" addresses being carried. This value is the ATM Forum 'address family number' specified in [3] as 19 decimal (0x0013). This is the ar\$hrd field defined in [1].

Protocol Type

This field is the protocol type number for the protocol using ATMARP from [3]. (IP is 0x0800). This is the ar\$pro field from [1].

Lifetime

This field contains a value (in minutes) which represents the lifetime of the ATMARP entry. A value of 0 indicates that the entry should be deleted.

Note that a time-out of a cache entry does not cause a CSA Record to be sent because, if everything is working properly then all ATMARP servers have the cache entry timing out at the same time. Thus, the individual servers would take the appropriate actions necessary.

ATM Addr T/L

This field contains the type and length of the ATM Address field. The type and length encodings are described in Section 8.7.3 of [\[1\]](#).

ATM SubAddr T/L

This field contains the type and length of the ATM SubAddress field. The type and length encodings are described in Section

8.7.3 of [\[1\]](#).

Proto Addr Len

This field contains the length of the Protocol Address field. For IPv4, the value is 4.

ATM Address

This is the ATM address of an address binding in an ATMARP server cache entry. The address, if specified, is E.164 or ATM Forum NSAPA.

ATM Subaddress

This is the ATM subaddress of an address binding in an ATMARP server cache entry. The subaddress, if specified, is an ATM Forum NSAPA. If null, no storage will be allocated.

Protocol Address

This is the internetwork address of an address binding in an ATMARP server cache entry.

The following sections give values for fields of the SCSP Protocol Independent Part of the various SCSP messages.

[4.1](#) Values for the SCSP "Mandatory Common Part"

Protocol ID = 0x0001
Sender ID Len = 0x04
Recvr ID Len = 0x04

See Section B.2.0.1 of [\[2\]](#) for a detailed description of these fields.

[4.2](#) Values for the SCSP "CSAS Record"

Cache Key Len = 0x04
Orig ID Len = 0x04

See Section B.2.0.2 of [\[2\]](#) for a detailed description of these fields.

References

- [1] "Classical IP and ARP over ATM", Mark Laubach and Joel Halpern, [RFC 2225](#).

Luciani, Halpern, Fox

[Page 5]

INTERNET-DRAFT

SCSP

Expires May 1999

- [2] "Server Cache Synchronization Protocol (SCSP)", Luciani, Armitage, Halpern, [RFC 2334](#).
- [3] Assigned Numbers, J. Reynolds and J. Postel, [RFC 1700](#).

Acknowledgments

We would like to thank the members of the ION working group of the IETF, whose review and discussion of this document has been invaluable.

Author's Address

James V. Luciani

Bay Networks, Inc.
3 Federal Street, BL3-04
Billerica, MA 01821
phone: +1-508-916-4734
email: luciani@baynetworks.com

Joel Halpern
Newbridge Networks, Inc
593 Herndon Parkway
Herndon, VA 22070-5241
phone: +1-703-736-5954
email: jhalpern@Newbridge.com

Barbara A. Fox
Lucent Technologies
300 Baker Ave, Suite 100
Concord, MA 01742-2168
phone: +1-978-287-2843
email: barbarafox@lucent.com