

Minshall & Ritter IETF Draft: SNMP over AppleTalk expires 22 January 1993  
SNMP over a Multi-protocol Internet Working Group  
Internet Draft: [draft-ietf-mpsnmp-appletalk](#)

Greg Minshall  
Novell, Inc.  
Mike Ritter  
Apple Computer, Inc.

**22 July 1992** (version 2.5)

SNMP over AppleTalk

## **1. 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. 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.

This draft document is being circulated for comment. The changes recommended by consensus in the IETF's "SNMP over a Multi-protocol Internet" working group are incorporated, and after review it will be submitted to the RFC editor as a Proposed Standard protocol specification. Distribution of this memo is unlimited.

## **2. Introduction**

This memo describes the method by which the Simple Network Management Protocol (SNMP) as specified in [1] can be used over AppleTalk protocols [2] instead of the Internet UDP/IP protocol stack. This specification is useful for network elements which have AppleTalk support but lack TCP/IP support. It should be noted that if a network element supports multiple protocol stacks, and UDP is available, it is the preferred network layer to use.

SNMP has been successful in managing Internet capable network elements which support the protocol stack at least through UDP, the connectionless Internet transport layer protocol. As originally designed, SNMP is capable of running over any reasonable transport mechanism (not necessarily a transport protocol) that supports bi-directional flow and addressability.

Many non-Internet capable network elements are present in networks. Some of these elements are equipped with the AppleTalk protocols. One method of using SNMP to manage these elements is to define a method of transmitting an SNMP message inside an AppleTalk protocol data unit.

### **3. Background**

The AppleTalk equivalent of UDP (and IP) is DDP (Datagram Delivery Protocol). The header field of a DDP datagram includes (at least conceptually) source and destination network numbers, source and destination node numbers, and source and destination socket numbers. Additionally, DDP datagrams include a "protocol type" in the header field which may be used to further demultiplex packets. The data portion of a DDP datagram may contain from zero to 586 octets.

AppleTalk's Name Binding Protocol (NBP) is a distributed name-to-address mapping protocol. NBP names are logically of the form "object:type@zone", where "zone" is determined, loosely, by the network on which the named entity resides; "type" is the kind of entity being named; and "object" is any string which causes "object:type@zone" to be unique in the AppleTalk internet. Generally, "object" also helps an end-user determine which instance of a specific type of service is being accessed. NBP names are not case sensitive. Each field of the NBP name ("object", "type", and "zone") is limited to 32 octets. The octets usually consist of human-readable ascii characters.

### **4. Specification**

SNMP REQUESTS encapsulated according to this standard will be sent to DDP socket number 8; they will contain a DDP protocol type of 8. The data octets of the DDP datagram will be a standard SNMP message as defined in [1].

SNMP RESPONSES encapsulated according to this standard will be sent to the DDP socket number which originated the corresponding SNMP request; they will contain a DDP protocol type of 8. The data octets of the DDP datagram will be a standard SNMP message as defined in [1]. (Note: as stated in [1], section 4.1, the \*source\* address of a RESPONSE PDU will be the same as the \*destination\* address of the corresponding REQUEST PDU.)

A network element which is capable of responding to SNMP REQUESTS over AppleTalk must advertise this capability via the AppleTalk Name Binding Protocol using an NBP type of "SNMP Agent" (hex 53, 4E, 4D, 50, 20, 41, 67, 65, 6E, 74).

A network management station which is capable of receiving an SNMP TRAP must advertise this capability via the AppleTalk Name Binding Protocol using an NBP type of "SNMP Trap Handler" (hex 53, 4E, 4D, 50, 20, 54, 72, 61, 70, 20, 48, 61, 6E, 64, 6C, 65, 72).

SNMP TRAPS encapsulated according to this standard will be sent to DDP socket number 9; they will contain a DDP protocol type of 8. The data octets of the DDP datagram will be a standard SNMP message as defined in [1]. The agent-addr

field of the Trap-PDU must be filled with a NetworkAddress of all zeros (the unknown IP address). Thus, to identify the trap sender, the name and value of the nbpObject and nbpZone corresponding to the nbpEntry with the nbpType equal to "SNMP Agent" should be included in the variable-bindings of any trap that is sent [3].

The NBP name for both an agent and a trap handler should be stable - it should not change any more often than the IP address of a typical TCP/IP end system changes. It is suggested that the NBP name be stored in some form of stable storage (PRAM, local disk, etc.).

## **6. Discussion of AppleTalk Addressing**

### **6.1 Introduction**

The AppleTalk protocol suite has certain features not manifest in the standard TCP/IP suite. Its unique naming strategy and the dynamic nature of address assignment can cause problems for SNMP management stations that wish to manage AppleTalk networks. TCP/IP end nodes, as of this writing, have an associated IP address which distinguishes each from the other. AppleTalk end nodes, in general, have no such characteristic. The network level address, while often relatively stable, can change at every reboot (or more frequently).

Thus, a thrust of this proposal is that a "name" (as opposed to an "address") for an end system be used as the identifying attribute. This is the equivalent, when dealing with TCP/IP end nodes, of using the domain name. While the mapping (DNS name, IP address) is more stable than the mapping (NBP name, DDP address), the mapping (DNS name, IP address) is not required to exist (eg, hosts with no host name, only an IP address). In contrast, all AppleTalk nodes that implement this specification are required to respond to NBP lookups and confirms (eg., implement the NBP protocol stub), which guarantees that the mapping (NBP name, DDP address) will exist.

In determining the SNMP name to register for an agent, it is suggested that the SNMP name be a name which is associated with other network services offered by the machine. On a Macintosh system, for example, it is suggested that the system name (the "Macintosh Name" for System 7.0 which is used to advertise file sharing, program-to-program communication, and possibly other services) be used as the "object" field of the NBP name. This name has AppleTalk significance, and is tightly bound to the network's concept of a given system's identity.

NBP lookups, which are used to turn NBP names into DDP addresses, can cause large amounts of network traffic as well as consume CPU resources. It is also the case that the ability to perform an NBP lookup is sensitive to certain network disruptions (such as zone table inconsistencies, etc.) which would not prevent direct AppleTalk communications between a management station and an agent.

Thus, it is recommended that NBP lookups be used infrequently with the primary

purpose being to create a cache of name-to-address mappings. These cached mappings should then be used for any further SNMP requests. It is recommended that SNMP management stations maintain this cache between reboots. This caching can help minimize network traffic, reduce CPU load on the network, and allow for (some amount of) network trouble shooting when the basic name-to-address translation mechanism is broken.

## **6.2 How To Acquire NBP names:**

A management station may have a pre-configured list of names of agents to manage. A management station may allow for an interaction with an operator in which a list of manageable agents is acquired (via NBP) and presented for the operator to choose which agents should be managed by that management station. Finally, a management station may manage all manageable agents in a set of zones or networks.

An agent must be configured with the name of a specific management station or group of management stations before sending SNMP traps. In the absence of any such configured information, an agent is NOT to generate any SNMP traps. In particular, an agent is NEVER to initiate a wildcard NBP lookup to find a management station to receive a trap. All NBP lookups generated by an agent must be fully specified. Note, however, that this does not apply to a configuration utility that might be associated with such an agent. Such a utility may well allow a user to navigate around the network to select a management station or stations to receive SNMP traps from the agent.

## **6.3 When To Turn NBP Names Into Addresses:**

When SNMP agents or management stations use a cache entry to address an SNMP packet, they should attempt to confirm the mapping if it hasn't been confirmed in T1 seconds. This cache entry lifetime, T1, has a minimum, default value of **60 seconds**. This value should be configurable.

A management station may decide to prime its cache of names prior to actually sending any SNMP requests to any given agent. In general, it is expected that a management station may want to keep certain mappings "more current" than other mappings. In particular, those nodes which represent the network infrastructure (routers, etc.) may be deemed "more important" by the management station.

Note, however, that a long-running management station starting up and reading a configuration file containing a number of NBP names should not attempt to prime its cache all at once. It should, instead, issue the resolutions over an extended period of time (perhaps in some pre-determined or configured priority order). Each resolution might, in fact, be a wildcard lookup in a given zone.

An agent should NEVER prime its cache. It should do NBP lookups (or confirms) only when it needs to send an SNMP trap to a given management station. An agent does not need to confirm a cache entry to reply to a request.

## **6.4 How To Turn NBP Names Into Addresses:**

If the only piece of information available is the NBP name, then an NBP lookup should be performed to turn that name into a DDP address.

However, if there is a piece of stale information, it can be used as a hint to perform an NBP confirm (which sends a unicast to the network address which is presumed to be the target of the name lookup) to see if the stale information is, in fact, still valid.

An NBP name to DDP address mapping can also be confirmed implicitly using only SNMP transactions. If a management station is sending a get-request, it can add a request, in the same packet, for the destination's nbpObject and nbpZone corresponding to the nbpEntry with the nbpType equal to "SNMP Agent" [3]. The source DDP address can be gleaned from the reply and used with the nbpObject and nbpZone returned to confirm the cache entry.

The above notwithstanding, for set-requests, there is a race condition that can occur where an SNMP request may go to the wrong agent (because the old node went down and a new node came up with the same DDP address.) This is problematic because the wrong agent might generate a response packet that the management station could not distinguish from a reply from the intended agent. In the future, when SNMP security is implemented, each packet is authenticated at the destination, and the reply should implicitly confirm the validity of the cache entry used and prevent this race condition.

### **6.5 What if NBP is broken:**

Under some circumstances, there may be connectivity between a management station and an agent, but the NBP machinery required to turn an NBP name into a DDP address may be broken. Examples of failures which would cause this include: NBP FwdReq (forward NBP lookup onto locally attached network) broken at a router on the network containing the agent; NBP BrRq (NBP broadcast request) mechanism broken at a router on the network containing the management station (because of a zone table mis-configuration, for example); or NBP broken in the target node.

A management station which is dedicated to AppleTalk management might choose to alleviate some of these failures by implementing the router portion of NBP within the management station itself. For example, the management station might already know all the zones on the AppleTalk internet and the networks on which each zone appears. Given an NBP lookup which fails, the management station could send an NBP FwdReq to the network in which the agent was last located. If that failed, the station could then send an NBP LkUp (an NBP lookup packet) as a directed (DDP) multicast to each network number on that network. Of the above (single) failures, this combined approach will solve the case where either the local router's BrRq to NBP FwdReq mechanism is broken or the remote router's NBP FwdReq to NBP LkUp mechanism is broken.

### **6.6 How to Represent NBP Names in MIBs**

There are occasions when it is necessary to represent a transport

service address in a MIB. For instance, the SNMP party MIB [6] uses an OBJECT IDENTIFIER to define the transport domain (IP, AppleTalk, etc.) and an OCTET STRING to represent an address within that domain. The following definitions are provided for use in such a scheme.

RFC-yyyy DEFINITIONS ::= BEGIN

IMPORTS

experimental

FROM RFC1155-SMI;

snmpOverAppleTalk

OBJECT IDENTIFIER ::= { experimental xx } -- xx TBD --

snmpOverAppleTalkDomain

OBJECT IDENTIFIER ::= { snmpOverAppleTalk 1}

-- An ATAddress is an octet string with a maximum length of 99 octets. The  
-- layout of an NBP name as an ATAddress is similar to a 'packed Pascal string'  
-- format: It consists of six fields: three length fields a single octet  
-- long, and three fields, up to thirty-two octets in length each, in the  
-- following order: length field, object field, length field, type field,  
-- length field, zone field. The layout of the octet string for the ATAddress  
-- of an SNMP Agent on the machine named 'theDragon' in the zone  
-- 'DA3/Nets-R-Us' would be:  
-- 0x09'theDragon'0x0A'SNMP Agent'0x0D'DA3/Nets-R-Us',  
-- where 0xNN implies a length octet in hexadecimal and the quoted strings  
-- represent ascii characters. The object, type, and zone fields may contain  
-- any octet except 0xFF and are case-insensitive [3].

--  
-- When devices are installed, they need to be configured with an  
-- initial set of SNMP parties. The configuration of SNMP parties  
-- requires (among other things) the assignment of several OBJECT  
-- IDENTIFIERS. Any local network administration can obtain the  
-- delegated authority necessary to assign its own OBJECT IDENTIFIERS.  
-- However, to cater for those administrations who have not obtained  
-- the necessary authority, this document allocates a branch of the  
-- naming tree for use with the following conventions.

initialAppleTalkPartyId OBJECT IDENTIFIER ::=  
{ snmpOverAppleTalk 2 }

-- This prefix is used in an analogous fashion as

--

-- initialPartyId

--

-- as defined in [6]. For an SNMP protocol entity residing at  
-- AppleTalk transport address 'N', the identities of its six initial  
-- parties are formed by appending a subidentifier for each octet in the  
-- AppleTalk TransportAddress, to

--

```
--      initialAppleTalkPartyId
--
END
```

## **7. Acknowledgements**

Some of the boilerplate in this memo is copied from [\[4\]](#), [\[5\]](#), and [\[7\]](#). The Apple-IP Working Group was instrumental in defining this document. Their support and work was greatly appreciated.

## **8. References**

- [1] Case, J., M. Fedor, M. Schoffstall, M., and J. Davin, "A Simple Network Management Protocol (SNMP)", [RFC 1157](#), May 1990.
- [2] Sidhu, G., R. Andrews, and A. Oppenheimer, "Inside AppleTalk (Second Edition)", Addison-Wesley, 1990.
- [3] Waldbusser, Steven, "AppleTalk Management Information Base", [RFC 1243](#), August 1991.
- [4] Schoffstall, M., C. Davin, M. Fedor, and J. Case, "SNMP over Ethernet", RFC 1089, February 1989.
- [5] Bostock , Steve, "SNMP over IPX" (draft to replace [RFC 1298](#))
- [6] McCloghrie, K., Davin, J., and J. Galvin, "Definitions of Managed Objects for Administration of SNMP Parties", RFC in preparation.
- [7] Piscitello, Dave, "Guidelines for the Specification of Protocol Support of the SNMP", RFC in preparation.

## **10. Authors' Addresses**

Greg Minshall  
Novell, Inc.  
**1340 Treat Blvd, ste. 500**  
Walnut Creek, CA 94596

Phone: 510 947-0998  
Fax: 510 947-1238  
EMail: minshall@wc.novell.com

Mike Ritter  
Apple Computer, Inc.  
**10500 North De Anza Boulevard, MS: 35-K**  
Cupertino, California 95014

Phone: 408 862-8088

Fax: 408 862-1159

E-Mail: MWRITTER@applelink.apple.com

Minshall & Ritter IETF Draft: SNMP over AppleTalk expires 22 January 1993