Network Working Group
Request for Comments: 1504

A. Oppenheimer
Apple Computer
August 1993

Appletalk Update-Based Routing Protocol: Enhanced Appletalk Routing

Status of This Memo

This memo provides information for the Internet community. It does not specify an Internet standard. Distribution of this memo is unlimited.

Introduction

This memo is being distributed to members of the Internet community to fully document an Apple protocol that may be running over the Internet. While the issues discussed may not be directly relevant to the research problems of the Internet, they may be interesting to a number of researchers and implementers.

About This Document

This document provides detailed information about the AppleTalk Update-based Routing Protocol (AURP) and wide area routing. AURP provides wide area routing enhancements to the AppleTalk routing protocols and is fully compatible with AppleTalk Phase 2. The organization of this document has as its basis the three major components of AURP:

AppleTalk tunneling, which allows AppleTalk data to pass through foreign networks and over point-to-point links

the propagation of AppleTalk routing information between internet routers connected through foreign networks or over point-to-point links

the presentation of AppleTalk network information by an internet router to nodes and other Phase 2-compatible routers on its local internet

What This Document Contains

The chapters of this document contain the following information:

Chapter 1, "Introduction to the AppleTalk Update-Based Routing Protocol," introduces the three major components of AURP and the

key wide area routing enhancements that AURP provides to the AppleTalk routing protocols.

Chapter 2, "Wide Area AppleTalk Connectivity," provides information about AppleTalk tunneling through IP internets and over point-to-point links.

Chapter 3, "Propagating Routing Information With the AppleTalk Update-Based Routing Protocol," describes the essential elements of AURP, including the architectural model for update-based routing. This chapter provides detailed information about the methods that AURP uses to propagate routing information between internet routers connected through tunnels.

Chapter 4, "Representing Wide Area Network Information," describes optional features of AURP-some of which can also be implemented on routers that use RTMP rather than AURP for routing-information propagation. It gives detailed information about how an exterior router represents imported network information to its local internet and to other exterior routers. It describes network hiding, device hiding, network-number remapping, clustering, loop detection, hop-count reduction, hop-count weighting, and backup paths.

The Appendix, "Implementation Details," provides information about implementing AURP.

What You Need to Know

This document is intended for developers of AppleTalk wide area routing products. It assumes familiarity with the AppleTalk network system, internet routing, and wide area networking terms and concepts.

Format of This RFC Document

The text of this document has been quickly prepared for RFC format. However, the art is more complex and is not yet ready in this format. We plan to incorporate the art in the future. Consult the official APDA document, as indicated below, for the actual art.

For More Information

The following manuals and books from Apple Computer provide additional information about AppleTalk networks. You can obtain books published by Addison-Wesley at your local bookstore. Contact APDA, Apple’s source for developer tools, to obtain technical reference materials for developers:

APDA
Apple Computer, Inc.
20525 Mariani Avenue, M/S 33-G
Cupertino, CA 95014-6299

These manuals provide information about some AppleTalk network products:

The Apple Ethernet NB User's Guide explains how to install and use an Apple Ethernet NB Card and EtherTalk software on an AppleTalk network.

The Apple InteroPoll Network Administrator's Guide describes how to perform maintenance and troubleshooting on an AppleTalk network using InteroPoll, a network administrator's utility program.

The Apple Internet Router Administrator's Guide explains how to install the Apple Internet Router Basic Connectivity Package and how to use the Router Manager application program. It provides information about setting up the router, configuring ports to create local area and wide area internets, monitoring and troubleshooting router operation, and planning your internet.

Using the AppleTalk/IP Wide Area Extension explains how to install and use the AppleTalk/IP Wide Area Extension for the Apple Internet Router. It provides information about tunneling through TCP/IP networks, configuring an IP Tunnel access method for an Ethernet or Token Ring port on the Apple Internet Router, troubleshooting IP tunneling problems, and configuring MacTCP.

The AppleTalk Remote Access User's Guide explains how to use a Macintosh computer to communicate with another Macintosh computer over standard telephone lines to access information and resources at a remote location.

The Apple Token Ring 4/16 NB Card User's Guide explains how to install and operate the card and TokenTalk software on a Token Ring network.

The MacTCP Administrator's Guide, version 1.1, explains how to install and configure the MacTCP driver, which implements TCP/IP (Transmission Control Protocol/Internet Protocol) on a Macintosh computer.

The following books provide reference information about AppleTalk networks:

The Advantages of AppleTalk Phase 2 provides a detailed description of the enhanced internetworking capabilities of AppleTalk Phase 2, and a brief guide to upgrading an AppleTalk internet to AppleTalk Phase 2. Available from Apple Computer.

The AppleTalk Network System Overview provides a technical introduction to the AppleTalk network system and its protocol architecture. Published by Addison-Wesley Publishing Company.

The AppleTalk Phase 2 Introduction and Upgrade Guide is a detailed guide to upgrading AppleTalk network hardware, drivers, and application programs to AppleTalk Phase 2, and briefly describes extensions to the AppleTalk network system that enhance its support for large networks. Available from Apple Computer.

The AppleTalk Phase 2 Protocol Specification is an addendum to the first edition of Inside AppleTalk that defines AppleTalk Phase 2 extensions to AppleTalk protocols that provide enhanced AppleTalk addressing, routing, and naming services. Available from APDA.

Inside AppleTalk, second edition, is a technical reference that describes the AppleTalk protocols in detail and includes information about AppleTalk Phase 2. Published by Addison-Wesley Publishing Company.

The Local Area Network Cabling Guide provides information about network media, topologies, and network types. Available from Apple Computer.

Planning and Managing AppleTalk Networks provides in-depth information for network administrators about planning and managing AppleTalk networks-including AppleTalk terms and concepts, and information about network services, media, topologies, security, monitoring and optimizing network performance, and troubleshooting. Published by Addison-Wesley Publishing Company.

Understanding Computer Networks provides an overview of networking-including basic information about protocol architectures, network media, and topologies. Published by Addison-Wesley Publishing Company.

The AppleTalk Update-Based Routing Protocol Specification is the official Apple specification of AURP. It includes the artwork currently missing from this document. Available from APDA.

Table of Contents

1. INTRODUCTION TO THE APPLETALK UPDATE-BASED ROUTING PROTOCOL

The AppleTalk Update-based Routing Protocol (AURP) provides wide area routing enhancements to the AppleTalk routing protocols and is fully compatible with AppleTalk Phase 2. AURP consists of three major components:

- AppleTalk tunneling through foreign network systems-for example, TCP/IP (Transmission Control Protocol/Internet Protocol) and over point-to-point links

- the propagation of routing information between internet routers connected through foreign network systems or over point-to-point links

- the presentation of AppleTalk network information by an internet router to nodes or to other Phase 2-compatible routers on its local internet-in other words, on the AppleTalk internet connected directly to the router

Chapter 3, "Propagating Routing Information With the AppleTalk Update-Based Routing Protocol," describes the elements of AURP that are essential for a minimal implementation of AURP. AURP includes many optional features for the presentation of network information. You can implement many of these optional features on routers that use either AURP or RTMP (Routing Table Maintenance Protocol) for routing-information propagation.

Figure 1-1 shows how the three major components of AURP interact.

<<Figure 1-1 Major components of AURP>>

Wide Area Routing Enhancements Provided by AURP

AURP provides AppleTalk Phase 2-compatible routing for large wide area networks (WANs). Key wide area routing enhancements provided by AURP include:

- tunneling through TCP/IP internets and other foreign network systems

- point-to-point tunneling

- basic security-including device hiding and network hiding

- remapping of remote network numbers to resolve numbering conflicts

internet clustering to minimize routing traffic and routing-information storage requirements

hop-count reduction to allow the creation of larger internets improved use of alternate paths through hop-count weighting and the designation of backup paths

2. WIDE AREA APPLETALK CONNECTIVITY

This chapter describes the wide area connectivity capabilities provided by the AppleTalk Update-based Routing Protocol (AURP), including:

AppleTalk tunneling

tunneling through TCP/IP internets

tunneling over point-to-point links

AppleTalk Tunneling

Tunneling allows a network administrator to connect two or more native internets through a foreign network system to form a large wide area network (WAN). For example, an AppleTalk WAN might consist of two or more native AppleTalk internets connected through a tunnel built on a TCP/IP internet. In such an AppleTalk WAN, native internets use AppleTalk protocols, while the foreign network system uses a different protocol family.

A tunnel connecting AppleTalk internets functions as a single, virtual data link between the internets. A tunnel can be either a foreign network system or a point-to-point link. Figure 2-1 shows an AppleTalk tunnel.

<<Figure 2-1 AppleTalk tunnel>>

There are two types of tunnels:

dual-endpoint tunnels, which have only two routers on a tunnel-for example, point-to-point tunnels

multiple-endpoint tunnels-herein referred to as multipoint tunnels-which have two or more routers on a tunnel

AURP implements multipoint tunneling by providing mechanisms for data encapsulation and the propagation of routing information to specific routers.

Exterior Routers

An AppleTalk router with a port that connects an AppleTalk internet to a tunnel is an exterior router. An exterior router always sends split-horizoned routing information to the other exterior routers on a multipoint tunnel. That is, an exterior router on a multipoint tunnel sends routing information for only its local internet to other exterior routers on that tunnel. An exterior router never exports routing information obtained from other exterior routers on the tunnel, because the exterior routers communicate their own routing information to one another.

As shown in Figure 2-2, the absence or presence of redundant paths, or loops, across a tunnel changes the way an exterior router defines its local internet. For more information about redundant paths, see the section "Redundant Paths" in Chapter 4. If no loops exist across a tunnel, an exterior router's local internet comprises all networks connected directly or indirectly to other ports on the exterior router. When loops exist across a tunnel, an exterior router's local internet comprises only those networks for which the next internet router is not across a tunnel. Using this definition of a local internet, two exterior routers' local internets might overlap if loops existed across a tunnel. For more information about routing loops, see the section "Routing Loops" in Chapter 4.

<<Figure 2-2 An exterior router's local internet>>

An exterior router functions as an AppleTalk router within its local internet and as an end node in the foreign network system connecting AppleTalk internets. An exterior router uses RTMP to communicate routing information to its local internet, and uses AURP and the network-layer protocol of the tunnel's underlying foreign network system to communicate with other exterior routers connected to the tunnel. An exterior router encapsulates AppleTalk data packets using the headers required by the foreign network system, then forwards the packets to another exterior router connected to the tunnel.

FORWARDING DATA: When forwarding AppleTalk data packets across a multipoint tunnel, an exterior router

- encapsulates the AppleTalk data packets in the packets of the tunnel's underlying foreign network system by adding the headers required by that network system

- adds an AURP-specific header-called a domain header-immediately preceding each AppleTalk data packet

A domain header contains additional addressing information-including a source domain identifier and destination domain identifier. For more information about domain headers, see the sections "AppleTalk Data-Packet Format" and "AppleTalk Data-Packet Format for IP Tunneling" later in this chapter. For detailed information about domain identifiers, see the section "Domain Identifiers" later in this chapter.

Before forwarding a data packet to a network in another exterior router's local internet, an exterior router must obtain the foreign-protocol address of the exterior router that is the next internet router in the path to the packet's destination network. The exterior router then sends the packet to that exterior router's foreign-protocol address using the network-layer protocol of the foreign network system. The exterior router need not know anything further about how the packet traverses this virtual data link.

Once the destination exterior router receives the packet, it removes the headers required by the foreign network system and the domain header, then forwards the packet to its destination in the local AppleTalk internet.

If the length of an AppleTalk data packet in bytes is greater than that of the data field of a foreign-protocol packet, a forwarding exterior router must fragment the AppleTalk data packet into multiple foreign-protocol packets, then forward these packets to their destination. Once the destination exterior router receives all of the fragments that make up the AppleTalk data packet, it reassembles the packet.

CONNECTING MULTIPLE TUNNELS TO AN EXTERIOR ROUTER: An exterior router can also connect two or more multipoint tunnels. As shown in Figure 2-3, when an exterior router connects more than one multipoint tunnel, the tunnels can be built on any of the following:

the same foreign network system

different foreign network systems

similar, but distinct foreign network systems

<<Figure 2-3 Connecting multiple tunnels to an exterior router>>

Whether the tunnels connected to an exterior router are built on similar or different foreign network systems, each tunnel acts as an independent, virtual data link. As shown in Figure 2-4, an exterior router connected to multiple tunnels functions logically as though it were two or more exterior routers connected to the same AppleTalk

network, with each exterior router connected to a different tunnel.

<<Figure 2-4 An exterior router connected to multiple tunnels>>

Fully Connected and Partially Connected Tunnels

An AppleTalk multipoint tunnel functions as a virtual data link. AURP assumes full connectivity across a multipoint tunnel-that is, all exterior routers on such a tunnel can communicate with one another. An exterior router always sends split-horizoned routing information to other exterior routers on a multipoint tunnel. That is, an exterior router on a multipoint tunnel sends routing information for only its local internet to other exterior routers on that tunnel. An exterior router never exports routing information obtained from other exterior routers on the tunnel, because exterior routers communicate their routing information to one another.

If all exterior routers connected to a multipoint tunnel are aware of and can send packets to one another, that tunnel is fully connected. If some of the exterior routers on a multipoint tunnel are not aware of one another, the tunnel is only partially connected. Figure 2-5 shows examples of a fully connected tunnel, a partially connected tunnel, and two fully connected tunnels.

<<Figure 2-5 Fully connected and partially connected tunnels>>

In the second example shown in Figure 2-5, the network administrator may have connected the tunnel partially for one of these reasons:

to prevent the local internets connected to exterior routers A and C from communicating with one another, while providing full connectivity between the local internets connected to exterior router

B and the local internets connected to both exterior routers A and C

because local internets connected to exterior routers A and C need access only to local internets connected to exterior router B-not to each other’s local internets

because exterior routers A and C-which should be aware of one another-were misconfigured

Generally, an exterior router cannot determine whether a multipoint tunnel is fully connected or partially connected. In the second example in Figure 2-5, exterior router B does not know whether exterior routers A and C are aware of one another. However, exterior

router B must assume that the tunnel is fully connected, and that exterior routers A and C can exchange routing information. An exterior router should never forward routing information received from other exterior routers back across the tunnel. It should always send split-horizoned routing information to other exterior routers.

If connecting exterior routers A and C directly would be either expensive or slow, a network administrator could instead establish two independent multipoint tunnels-one connecting exterior routers A and B, another connecting exterior routers B and C-as shown in the third example in Figure 2-5. Exterior routers A and C could then establish connectivity by routing all data packets forwarded by one to the other through exterior router B.

Hiding Local Networks From Tunnels

When configuring a tunneling port on an exterior router, a network administrator can provide network-level security to a network in the exterior router's local internet by hiding that network. Hiding a specific network in the exterior router's local internet prevents internets across a multipoint tunnel from becoming aware of the presence of that network. When the exterior router exchanges routing information with other exterior routers connected to the tunnel, it exports no information about any hidden networks to the exterior routers from which the networks are hidden.

An administrator can specify that certain networks in the exterior router's local internet be hidden from a specific exterior router connected to the tunnel or from all exterior routers on the tunnel.

Nodes on the local internet of an exterior router from which a network is hidden cannot access that network. Neither the zones on a hidden network nor the names of devices in those zones appear in the Chooser on computers connected to such an internet. When a network is hidden, its nodes are also unable to access internets from which the network is hidden. If a node on a hidden network sends a packet across a tunnel to a node on an internet from which it is hidden, even if the packet arrives at its destination, the receiving node cannot respond. The exterior router connected to the receiving node's internet does not know the return path to the node on the hidden network. Thus, it appears to the node on the hidden network that the node to which it sent the packet is inaccessible.

ADVANTAGES AND DISADVANTAGES OF NETWORK HIDING: Network hiding provides the following advantages:

- On large, global WANs, a network administrator can configure network-level security for an organization's internets.

It reduces the amount of network traffic across both a tunnel and the internets connected to that tunnel.

Network hiding has the following disadvantages:

Nodes on hidden networks have limited access to internets across a tunnel.

AppleTalk networking software running on a node on a hidden network lists all of the AppleTalk zone names exported by exterior routers connected to a tunnel, but may list the names of only some or none of the devices in those zones. It cannot list the names of devices that are unable to respond to Name Binding Protocol (NBP) lookups originating from a node on a hidden network.

Domain Identifiers

Exterior routers assign a unique domain identifier to each AppleTalk internet, or domain. Domain identifiers enable exterior routers on a multipoint tunnel to distinguish individual AppleTalk internets in a wide area internet from one another.

The definition of an AppleTalk domain identifier is extensible to allow for future use when many additional types of AppleTalk tunnels and tunneling topologies may exist:

Under the current version of AURP, each exterior router connected to a multipoint tunnel assigns a domain identifier to its local AppleTalk internet that uniquely identifies that internet on the tunnel. If redundant paths connect an AppleTalk internet through more than one exterior router on a tunnel, each exterior router can assign a different domain identifier to that internet, or AppleTalk domain, as shown in Figure 2-6.

Under future routing protocols, a domain identifier will define the boundaries of an AppleTalk domain globally-for all exterior routers. Thus, a domain identifier will be unique among all domains in a wide area internet. All exterior routers within a wide area internet will use the same domain identifier for a given AppleTalk internet, as shown in Figure 2-6.

<<Figure 2-6 Domain identifiers>>

To simplify an exterior router’s port configuration, a parameter that is already administrated-such as a node address-can serve as the basis for an exterior router’s domain identifier.

GENERAL DOMAIN-IDENTIFIER FORMAT: Figure 2-7 shows the general form of a domain identifier.

<<Figure 2-7 General domain-identifier format>>

The general domain identifier (DI) consists of the following fields:

Length: Byte 1 represents the length of the DI in bytes, not including the length byte. A DI must consist of an even number of bytes. Thus, the length byte is always an odd-numbered byte. The length field permits tunneling through foreign network systems that have addresses of any length-including the long addresses characteristic of X.25 and OSI. The value of the length byte varies, depending on the format of the DI.

Authority: Byte 2 indicates the authority that administrates the identifier bytes of the DI. At present, Apple has defined only two authority-byte values:

$01-indicates that the subsequent bytes correspond to a unique, centrally administrated IP address

$00-the null DI-indicates that no additional bytes follow

All other authority-byte values are reserved and should not be used.

Identifier: The identifier field starts at byte 3 and consists of a variable number of bytes of the type indicated by the authority byte.

NULL DOMAIN-IDENTIFIER FORMAT: The use of a null domain identifier is appropriate only when there is no need to distinguish the domains connected to a tunnel-for example, where a tunnel exists within a single internet-or for a point-to-point link. Figure 2-8 shows the null form of a domain identifier.

<<Figure 2-8 Null domain-identifier format>>

A null domain identifier consists of the following bytes:

Length: Byte 1 contains the value $01, defining the length of the null DI as one byte.

Authority: Byte 2 contains the value $00, indicating a null DI.

AppleTalk Data-Packet Format

Part of the format of an AppleTalk data packet sent across a multipoint tunnel or a point-to-point link depends on the underlying

foreign network system. The headers required by a foreign-network protocol always precede an AppleTalk data packet sent across a multipoint tunnel. A domain header generally immediately precedes the AppleTalk data packet. Figure 2-9 shows the format of an AppleTalk data packet preceded by a domain header.

<<Figure 2-9 AppleTalk data-packet format with a domain header>>

A domain header consists of the following fields:

Destination DI: The length of the destination DI field in bytes depends on the type of DI.

Source DI: The length of the source DI field in bytes depends on the type of DI.

Version number: The version number field is two bytes in length and currently contains the value 0001.

Reserved: The two-byte field that follows the version number field is reserved for future use and is set to 0000.

Packet type: The two-byte packet type field contains the value 0002 to identify the data that follows as AppleTalk data-distinguishing it from other data, such as routing data. In the future, Apple may define other values for this field.

An AppleTalk data packet does not require a domain header if

it is sent across a multipoint tunnel or point-to-point link that provides separate channels for data and routing packets

the domain header’s destination DI and source DI fields would both contain null DIs

Omitting a domain header reduces overhead associated with the exchange of routing information, without any loss of routing information. Figure 2-10 shows the format of an AppleTalk data packet without a domain header.

<<Figure 2-10 AppleTalk data-packet format without a domain header>>

IP Tunneling

The Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite is a widely used communications standard that provides interoperability among computers from various vendors, including Apple, IBM, Digital Equipment Corporation, Sun, and Hewlett-Packard.

Descriptions of three of the most important TCP/IP protocols follow:

The Transmission Control Protocol (TCP) is a transport-layer protocol that provides reliable data transmission between processes-that is, between programs that communicate with one another. This connection-oriented, byte-stream protocol ensures error-free, sequential data delivery, without loss or duplication.

The User Datagram Protocol (UDP) is a transport-layer protocol that provides best-effort, low-overhead interprocess data transmission. This datagram-oriented protocol allows higher-layer protocols that do not require reliability to transmit data without incurring the overhead associated with TCP. UDP does no error checking, does not acknowledge its successful receipt of data, and does not sequence incoming messages. UDP messages may be lost, duplicated, or improperly sequenced.

The Internet Protocol (IP) is a network-layer protocol that provides connectionless, best-effort datagram delivery across multiple networks. Each host on a TCP/IP network has a unique, centrally administrated internet address, called an IP address, that identifies the node. The header of an IP datagram contains its source and destination IP addresses, allowing any host to route a datagram to its destination. TCP/IP provides connectivity between many different network types that use data frames of various sizes. Therefore, IP can fragment a datagram before sending it across an internet. Datagram fragments can fit into data frames of any size. Once all of a datagram’s fragments reach their destination, IP reassembles the datagram.

Protocols in higher layers pass data to TCP or UDP for delivery to peer processes. TCP and UDP encapsulate the data in segments, using the appropriate headers, then pass the segments to IP. IP further encapsulates the data in IP datagrams, determines each datagram’s path to its destination, and sends the datagrams across the internet.

Figure 2-11 shows how the TCP/IP family of protocols conforms to the Open Systems Interconnection (OSI) model.

<<Figure 2-11 TCP/IP protocol stack and the OSI model>>

Exterior routers that connect AppleTalk internets through a TCP/IP tunnel are configured as nodes on both an AppleTalk internet and on the TCP/IP internet. Thus, an exterior router on a TCP/IP tunnel is also an IP end node in the TCP/IP network system. Exterior routers use the TCP/IP internet only to exchange AppleTalk routing information and AppleTalk data packets with one another. An exterior router encapsulates AppleTalk data packets in IP datagrams before

sending them across the TCP/IP internet to a forwarding exterior router, which decapsulates the packets, then forwards them to their destination AppleTalk networks.

IP Domain-Identifier Format

Under the current version of AURP, exterior routers on IP tunnels must use domain identifiers that are based on IP addresses. An exterior router on an IP tunnel derives its domain identifier from its IP address. Thus, a network administrator does not need to configure an exterior router's domain identifier. Figure 2-12 shows the IP form of a domain identifier.

<<Figure 2-12 IP domain-identifier format>>

An IP domain identifier consists of the following fields:

Length: Byte 1 contains the value $07, defining the length of the IP DI as seven bytes.

Authority: Byte 2 contains the value $01, indicating that the remainder of the DI is based on an IP address.

Distinguisher: Bytes 3 and 4 are reserved for future use and are set to 0 ($00).

IP address: Bytes 5 through 8 contain the four-byte IP address of either the sending or the receiving exterior router.

NOTE: Future versions of AURP will allow exterior routers to usealternative formats for domain identifiers, even on IP tunnels.

AppleTalk Data-Packet Format for IP Tunneling

The following protocol headers precede an AppleTalk data packet that is forwarded across an IP tunnel by an exterior router:

a data-link header

an IP header

a User Datagram Protocol (UDP) header

a domain header

An exterior router encapsulates AppleTalk data packets in UDP packets when forwarding them through its UDP port 387, across an IP tunnel, to UDP port 387 on another exterior router. When encapsulating data

packets, an exterior router should always use UDP checksums. When a destination exterior router receives the UDP packets at UDP port 387, it decapsulates the packets.

A domain header consists of the following fields:

Destination DI: This field contains the DI of the exterior router to which a packet is being forwarded.

Source DI: This field contains the DI of the exterior router that is forwarding a packet.

Version number: The version number field is two bytes in length and currently contains the value 0001.

Reserved: The two-byte field that follows the version number field is reserved for future use and is set to 0000.

Packet type: The two-byte packet type field contains the value 0002 to identify the data that follows as AppleTalk data-distinguishing it from other data, such as routing data.

An AppleTalk data packet consists of a domain header and AppleTalk data. Figure 2-13 shows the format of an AppleTalk data packet forwarded across an IP tunnel.

<<Figure 2-13 AppleTalk data packet forwarded across an IP tunnel>>

Point-to-Point Tunneling

In point-to-point tunneling, two remote AppleTalk local area networks (LANs) connected to half-routers communicate with one another over a point-to-point link. A point-to-point link may consist of modems communicating over a standard telephone line or a leased line, such as a T1 line. Figure 2-14 shows an example of point-to-point tunneling.

<<Figure 2-14 Point-to-point tunneling>>

Generally, exterior routers use null domain identifiers on point-to-point links, because there is no IP address to be administrated and the opposite end of the tunnel is already uniquely identified. However, an exterior router may use other domain-identifier formats.

Point-to-Point Protocol

The Point-to-Point Protocol (PPP) is a data-link-layer protocol that provides a standard method of encapsulating and decapsulating

network-layer protocol information, and transmitting that information over point-to-point links. PPP includes an extensible Link Control Protocol (LCP) and a suite of Network Control Protocols (NCPs) that configure, enable, and disable various network-layer protocols.

The AppleTalk Control Protocol (ATCP) is a PPP NCP for AppleTalk protocols. ATCP configures, enables, and disables the AppleTalk network-layer protocol DDP on the half-router at each end of a point-to-point link. ATCP also specifies the protocol that a half-router uses to propagate routing information-for example, AURP. When using AURP for routing-information propagation, a half-router uses a specific PPP protocol type to identify AURP routing-information packets-that is, packets preceded by a domain header. PPP provides separate channels for AppleTalk data packets and AppleTalk routing-information packets. Thus, a half-router can use DDP encapsulation to send AppleTalk data packets without including their domain headers. When using AURP, a half-router should accept both AppleTalk data packets that are preceded by domain headers and DDP-encapsulated packets.

NOTE: The Request for Comments (RFC) 1378, "The PPP AppleTalk Control Protocol (ATCP)," provides a detailed specification of ATCP, as well as information about using PPP to send AppleTalk data.

3. PROPAGATING ROUTING INFORMATION WITH THE APPLETALK UPDATE-BASED ROUTING PROTOCOL

This chapter describes the required elements of AURP. It provides detailed information about using the AppleTalk Update-based Routing Protocol (AURP) to propagate routing information between AppleTalk exterior routers connected through a foreign network or over a point-to-point link, and includes information about

- the AURP architectural model
- one-way connections
- exchanging routing information
- updating routing information
- notifying other exterior routers that an exterior router is going down
- obtaining zone information
- packet formats

error codes

AURP Architectural Model

AURP provides the functionality of the Routing Table Maintenance Protocol (RTMP) and the Zone Information Protocol (ZIP) while eliminating most of the routing traffic generated by these protocols. Figure 3-1 shows the architectural model for AURP.

<<Figure 3-1 AURP architectural model>>

Generally, an AppleTalk router uses RTMP and ZIP to maintain routing information, and sends RTMP data packets, ZIP Queries, and ZIP Replies out its ports. However, if one of the router's ports is connected to an AppleTalk tunnel, the architectural model for the router's central routing module becomes more complex. Logically, the central routing module in an exterior router communicates RTMP and ZIP information to an RTMP/ZIP-to-AURP conversion module, which sends AURP data packets out the tunneling port.

RTMP/ZIP-to-AURP Conversion Module

The RTMP/ZIP-to-AURP conversion module maintains split-horizoned routing-table information and network number-to-zone name mappings for each exterior router on the tunnel-that is, a copy of the routing information for each exterior router's local internet. Figure 3-2 shows the architectural components of the RTMP/ZIP-to-AURP conversion module.

<<Figure 3-2 RTMP/ZIP-to-AURP conversion module architecture>>

The AURP module of the conversion module obtains routing information from the other exterior routers on the tunnel, then periodically updates the routing-table information and the mappings in the conversion module. The RTMP module passes this routing-table information to the exterior router's central routing module. Logically, the RTMP module generates an RTMP data packet for each exterior router on the tunnel every ten seconds-the RTMP retransmission time-then passes the packet to the central routing module.

The RTMP/ZIP-to-AURP conversion module also maintains a split-horizoned copy of the routing information maintained by the exterior router in which it resides. Logically, the conversion module obtains the routing information from RTMP data packets and ZIP Replies sent by the exterior router's central routing module, then updates the routing information in the conversion module.

The AURP module exports routing information about its local AppleTalk internet to other exterior routers on the tunnel.

AURP Transport Layering

AURP can propagate routing information between exterior routers using

- a simple, reliable transport based on an underlying datagram service-such as the default transport-layer service for AURP, AURP-Tr. See the section "AURP-Tr," later in this chapter, for more information.

- a more complex transport-layer service-such as TCP

Figure 3-3 shows the AURP transport-layering model.

<<Figure 3-3 AURP transport-layering model>>

Maintaining Current Routing Information With AURP

AURP allows exterior routers to maintain current routing information for other exterior routers on a tunnel by supporting

- the reliable, initial exchange of split-horizoned routing information - that is, the routing information for an exterior router's local internet

- reliable updates to that information whenever it changes

If an internet topology does not change, AURP generates significantly less routing traffic than RTMP and ZIP. Thus, an administrator can connect very large AppleTalk internets through a tunnel, and the resulting internet generates little or no routing traffic on the tunnel.

When an exterior router discovers another exterior router on the tunnel-that is, a peer exterior router-it can request that exterior router to send its routing information. In a reliable, initial exchange of split-horizoned routing information, the peer exterior router returns its network-number list. The peer exterior router also returns each connected network's zone information in an unsequenced series of zone-information packets. If the exterior router requesting the routing information does not receive complete zone information for a network, it must retransmit requests for zone information until it receives the information.

Once an exterior router requesting routing information from a peer exterior router has received that exterior router's network-number

list and complete zone information, it typically requests the peer exterior router to notify it of any changes to that routing information. The peer exterior router then provides the requesting exterior router with reliable updates to its routing information- however, it sends no other routing information.

Notifying Other Exterior Routers of Events

If an exterior router has requested notification of changes in another exterior router's split-horizoned routing information, that exterior router must notify the requesting exterior router of any event that changes its routing information. Thus, an exterior router must send updated routing information to the requesting exterior router whenever any of the following events occur:

- the addition of a new, exported network-that is, a network that is not hidden-to the exterior router's local internet and, consequently, to its routing table
- a change in the path to an exported network that causes the exterior router to access that network through its local internet rather than through a tunneling port
- the removal of an exported network from the exterior router's routing table because a network in the exterior router's local internet has gone down
- a change in the path to an exported network that causes the exterior router to access that network through a tunneling port rather than through its local internet
- a change in the distance to an exported network
- a change to a zone name in the zone list of an exported network- an event not currently supported by ZIP or the current version of AURP
- the exterior router goes down or is shut down

Routing-information updates allow an exterior router to maintain accurate, split-horizoned routing information for a peer exterior router on a tunnel.

AURP-Tr

AURP-Tr, the default transport-layer service for AURP, provides a simple, reliable transport that is based on an underlying datagram service. When using AURP-Tr, only one sequenced transaction can be

outstanding, or unacknowledged, at a time-greatly simplifying the implementation of AURP, without limiting its functionality.

One-Way Connections

A one-way connection is an asymmetrical link between a data sender and a data receiver that are using AURP-Tr, in which an exterior router functioning as a data sender sends a sequenced, reliable, unidirectional data stream to an exterior router functioning as a data receiver. An exterior router can send routing information over a one-way connection as

sequenced data

transaction data

Sequenced data is data sent in sequence by the data sender and delivered reliably to the data receiver. Typically, the sending of sequenced data is unprovoked-that is, it is not requested by a data receiver. However, a data receiver can request sequenced data. Figure 3-4 shows sequenced data being sent across a one-way connection.

<<Figure 3-4 Sequenced data on a one-way connection>>

Transaction data-also referred to as out-of-band data-is data sent unsequenced by the data sender through a linked request/response transaction that is initiated by the data receiver.

The data receiver can use a one-way connection to request transaction data from the data sender. If the data receiver does not receive a response, it must retransmit its request. Figure 3-5 shows a one-way connection on which the data receiver requests transaction data from the data sender.

<<Figure 3-5 Request for transaction data on a one-way connection>>

Generally, communication between two exterior routers is bidirectional-that is, two one-way connections exist between the exterior routers, with each exterior router acting as the data sender on one connection and the data receiver on the other. Thus, each exterior router can send its routing information to the other.

Initial Information Exchange

When an AppleTalk exterior router discovers another exterior router on the tunnel, it uses the underlying transport-layer service to open a connection with that exterior router. When using AURP-Tr, an exterior router opens this connection as a one-way connection.

Open Request Packet

Once the data receiver opens a connection using the underlying transport, the data receiver sends an Open Request packet, or Open-Req, to the data sender. An Open-Req packet includes the following information:

Send update information flags: The states of the four send update information (SUI) flags indicate whether the data sender should send various types of update information over the connection. Typically, the four SUI flags are set to 1.

Version number: The version number field indicates the version of AURP used by the data receiver. The current version number of AURP is 1.

Data field: The optional data field allows exterior routers with capabilities beyond those described in this document to notify other exterior routers about such options, by initiating option negotiation. An exterior router that has similar capabilities indicates that it accepts the options, completing option negotiation. An exterior router that lacks such options ignores the information in the data field.

Open Response Packet

When an exterior router receives an Open-Req, it becomes the data sender and responds with an Open Response packet, or Open-Rsp, as follows:

If the exterior router accepts the connection, it returns information about its setup in the Open-Rsp. An Open-Rsp also contains an optional data field. This data field indicates whether the exterior router accepts the options in the data field of the Open-Req to which it is responding.

If the exterior router cannot accept the connection-for example, because the Open-Req does not contain the correct version number-it returns an error in the Open-Rsp and closes the transport-layer connection.

Figure 3-6 shows a connection-opening dialog between a data sender and a data receiver.

<<Figure 3-6 Connection-opening dialog>>

Routing Information Request Packet

Under AURP, once two exterior routers establish a connection, the data receiver can request the data sender to send its routing information by sending it a Routing Information Request packet, or RI-Req.

Routing Information Response Packets

When the data sender receives an RI-Req, it reliably sends a sequence of Routing Information Response packets, or RI-Rsp, to the exterior router requesting the information.

The RI-Rsp packets provide a list of exported networks on the data sender's local internet and the distance of each network from the data sender. The data sender must finish sending RI-Rsp packets to the exterior router requesting routing information before it can send any other sequenced data over the connection. Figure 3-7 shows a routing-information request/response dialog between a data sender and a data receiver.

<<Figure 3-7 Routing-information request/response dialog>>

Zone Information Request Packet

The data receiver can obtain zone information for known networks on the data sender's local internet at any time, by sending it a Zone Information Request packet, or ZI-Req. A ZI-Req lists the numbers of networks for which the data receiver is requesting zone information.

IMPORTANT: To prevent other exterior routers on a tunnel from sending endless streams of ZI-Req packets across the tunnel-causing what is referred to as a ZIP storm-an exterior router must not export information about a network until it has a complete zone list for that network.

Zone Information Response Packets

When the data sender receives a ZI-Req, it responds by sending unsequenced Zone Information Response packets, or ZI-Rsp, to the data receiver. Zone information is transaction data-thus, its reliable delivery is not guaranteed. Figure 3-8 shows a zone-information request/response dialog between a data sender and a data receiver.

<<Figure 3-8 Zone-information request/response dialog>>

Recovering Lost Zone Information

A data receiver enters a network-to-zone list association in its routing table for each network for which it receives a ZI-Rsp packet. If a data receiver that requested zone information for a network does not receive a complete zone list for that network, it must retransmit ZI-Req packets, requesting zone information for that network, until it receives that network's complete zone information.

To determine if any ZI-Rsp packets were lost, the data receiver periodically scans its routing table for networks for which the associated zone lists are incomplete-that is, for zone lists that do not include all zones associated with the networks. The data receiver sends a ZI-Req to each data sender from which it received incomplete zone information, listing the numbers of networks for which it has incomplete zone lists. The data sender responds to zone information requests by sending ZI-Rsp packets containing the requested information to the data receiver.

Using AURP-Tr for Initial Information Exchange

The following sections describe the use of AURP-Tr-the default transport-layer service for AURP-for initial information exchange.

OPEN REQUEST PACKET: An exterior router sends an Open-Req packet to

- request that an AURP-Tr one-way connection with another exterior router be established
- specify the connection ID for that connection
- pass the AURP version number, SUI flags, and optional data to the other exterior router

If the exterior router does not receive an Open-Rsp from the exterior router to which it sent an Open-Req, it must retransmit the Open-Req.

OPEN RESPONSE PACKET: When using AURP-Tr, an exterior router sends an Open-Rsp to

- acknowledge that a one-way connection has been established
- reject a connection
- return information about its environment, as well as any optional data, to the exterior router from which it received an Open-Req

If an exterior router receives an Open-Req on a one-way connection that is already open-that is, if it receives an Open-Req with the same connection ID as an open one-way connection-an Open-Rsp sent previously may have been lost. The exterior router receiving the duplicate Open-Req should send a duplicate Open-Rsp to the sending exterior router, unless it has already received some other packet on the connection-such as an RI-Req-indicating the existence of a fully established connection.

ROUTING INFORMATION RESPONSE PACKETS: When responding to a request for routing information using AURP-Tr, an exterior router sends a sequence of RI-Rsp packets to the exterior router requesting the information. However, an exterior router's complete list of network numbers often fits in a single RI-Rsp packet. Each RI-Rsp packet contains the following information:

Connection ID: The connection ID identifies the specific one-way connection to which a packet belongs.

Sequence number: The sequence number identifies an individual packet on a connection. Packets on a connection are numbered starting with the number 1.

The data sender sending routing information must wait for the data receiver to acknowledge that it has received each RI-Rsp packet in the sequence-by sending an RI-Ack packet-before sending the next RI-Rsp packet. Each RI-Rsp contains a flag that indicates whether it is the last packet in the sequence. In the last RI-Rsp in the sequence, this flag is set to 1. If the data sender receives no acknowledgment of an RI-Rsp from the data receiver within a specified period of time, it must retransmit the RI-Rsp.

ROUTING INFORMATION RESPONSE PACKETS: When an exterior router receives an RI-Rsp, it verifies the packet's connection ID and sequence number. The connection ID must be the same as that in the Open-Req. The sequence number must be either

the last sequence number received, indicating that the previous acknowledgment was lost or delayed, and that this is a duplicate RI-Rsp the next number in the sequence, indicating that this RI-Rsp contains new routing information

If the connection ID or sequence number is invalid, the data receiver discards the packet. Figure 3-9 shows a dialog between a data sender and a data receiver in which the data receiver requests routing information, the data sender responds by sending its routing information, and the data receiver acknowledges the data sender's response. If the data sender receives no acknowledgment, it sends

duplicate RI-Rsp packets until the data receiver responds with an acknowledgment.

<<Figure 3-9 Routing-information request/response/acknowledgment dialog>>

Once the data receiver has verified the information in the RI-Rsp, it responds with a Routing Information Acknowledgment packet, or RI-Ack, which contains the following information:

Connection ID: The connection ID is the same as that in the RI-Rsp packet.

Sequence number: The sequence number is the same as that in the RI-Rsp packet.

Send zone information flag: The state of the send zone information (SZI) flag in an RI-Ack packet indicates whether the RI-Ack packet doubles as a ZI-Req packet. If the SZI flag is set to 1, the data receiver sends the zone information associated with the networks about which it sent routing information in the previous RI-Rsp.

Figure 3-10 shows a data receiver sending zone information to a data sender in response to a ZI-Req and in response to an RI-Ack, which optimizes the data flow.

When the data sender receives an RI-Ack, it verifies that the RI-Ack corresponds to the outstanding RI-Rsp-that is, both packets have the same connection ID and sequence number. Once the data sender has verified the information in the RI-Ack, it responds by sending the next RI-Rsp in the sequence, if any.

<<Figure 3-10 Nonoptimized and optimized flows of zone information>>

ZONE INFORMATION RESPONSE PACKETS: If the data sender receives an RI-Ack with its SZI flag set to 1, it responds by sending ZI-Rsp packets that contain the zone information associated with the networks about which it sent routing information in the RI-Rsp being acknowledged-just as it would if it received a ZI-Req for those networks.

The data sender sends RI-Rsp and ZI-Rsp packets as independent data streams. It sends RI-Rsp packets as sequenced data and ZI-Rsp packets as transaction data. If the data sender receives an RI-Ack with its SZI flag set to 1, it sends an unsequenced series of ZI-Rsp packets that contain the following information:

Connection ID: The connection ID is the same as that in the

associated RI-Req.

Network number and zone list tuples: The exterior router sends the zone information associated with each network number in the corresponding RI-Rsp.

Reobtaining Routing Information

An exterior router can reobtain another exterior router's complete routing information at any time, by sending an RI-Req packet. An exterior router might need to reobtain complete routing information for a one-way connection on which it is the data receiver under the following circumstances:

During the initial routing-information exchange, the exterior router set the SUI flags in the Open-Req to disable updates. The exterior router can subsequently poll the other exterior router on the connection by sending an RI-Req to that exterior router to determine whether any of its routing information has changed.

The exterior router set the SUI flags to request updates, but suspects that the routing information for the other exterior router on the connection is incorrect or obsolete. The exterior router should send an RI-Req to the other exterior router to obtain its complete, updated routing information.

Whenever an exterior router receives an RI-Req from an exterior router requesting updated routing information, it responds by sending RI-Rsp packets, just as it does when it first receives an RI-Req. The data sender also resets the SUI flags for that one-way connection, so they correspond to those in the RI-Req.

If the data sender is sending other sequenced update information when it receives an RI-Req, it cannot respond to the RI-Req until the data receiver acknowledges the last outstanding packet in the sequence. If AURP uses an underlying transport-layer service that does not provide reliable delivery, such as AURP-Tr, it may be necessary for the data receiver to retransmit an RI-Req.

Updating Routing Information

Once an exterior router receives the routing and zone information for another exterior router's local internet, if the receiving exterior router has set the SUI flags in the Open-Req to request updates, the data sender notifies the data receiver of any subsequent changes to that information.

Informed-Routers List

An exterior router maintains an informed-routers list containing the network address of each exterior router that has requested dynamic updating of routing information. Once an exterior router has sent routing information for its local internet to other exterior routers on the tunnel, it must reliably send updated routing information to all accessible exterior routers in its informed-routers list whenever its routing information changes.

Sending Routing Information Update Packets

An exterior router communicates changes in its routing information by sending Routing Information Update, or RI-Upd, packets to another exterior router. When the routing information for an exterior router's local internet changes, the exterior router need not send an RI-Upd immediately. Generally, an exterior router buffers the update information, then sends updates periodically. The exterior router must wait at least an update interval between sending updates. The value of this update interval

cannot be less than ten seconds

should be specifiable by a network administrator

It is possible that more than one update event for a particular network might occur within one update interval. One of these events might supercede another-for example, a Network Added event followed by a Network Deleted event for the same network. In this case, the exterior router can represent the two events logically as one event. Under AURP, an exterior router can have only one event pending for a given network. An exterior router can combine any series of events for a network into a single pending event. In Figure 3-11, a state diagram shows the update event that an exterior router should have pending for a network, based on the other events that have occurred during the update interval.

<<Figure 3-11 A state diagram showing pending update events>>

Four of the states correspond to four pending update events. Two states indicate that no update event is pending:

Net Up-indicates that no update event is pending for a network in the exterior router's local internet

Net Down-indicates that no update event is pending for a network in another exterior router's local internet or the network does not exist

A single RI-Upd packet may contain different types of update events-for example, several Network Added events and several Network Deleted events. For information about update events, see the section "Routing-Information Update Events" later in this chapter.

A data sender should send an RI-Upd packet to an exterior router in its informed-routers list only if the packet contains one or more update events of a type indicated by the SUI flags of the last Open-Req or RI-Req received from that exterior router. Because an RI-Upd that contains one or more events of a type requested by an exterior router may also contain events of types not requested, an exterior router must be able to handle events of all types. Thus, a data sender can send an RI-Upd that contains various types of update events to all exterior routers that have requested update events of any of those types.

Sending Updates Following the Initial Exchange of Routing Information

While a data sender has update events pending-that is, when update events have occurred but the data sender has not yet sent RI-Upd packets for those events-another exterior router may establish a new connection with the data sender. The data sender must present consistent routing information to all exterior routers on the tunnel, on both existing connections and any new connections. For example, if a pending update event indicated that a new network had become available, the newly connected exterior router could be informed of that network's presence on the internet either by

- sending it an RI-Rsp packet including routing information for the new network

- sending it an RI-Rsp packet that does not include routing information for the new network, then sending it the RI-Upd packet that includes the pending update event

AURP does not specify a scheme for sending update information following the initial exchange of routing information on a new connection. However, the Appendix, "Implementation Details," describes one possible method of doing this.

Using AURP-Tr to Update Routing Information

The following sections describe the use of AURP-Tr for sending routing-information updates.

ROUTING INFORMATION UPDATE PACKETS: Each RI-Upd packet contains the following information:

Connection ID: The connection ID identifies the specific one-way connection to which the RI-Upd belongs.

Sequence number: The sequence number identifies an individual RI-Upd on a connection.

If an update cannot be contained in one RI-Upd packet, the data sender must send a sequence of RI-Upd packets. While the data sender need not wait for the duration of an update interval before sending each RI-Upd packet in a sequence, it must wait for the data receiver to acknowledge that it has received the RI-Upd packet that is currently outstanding before sending the next RI-Upd packet in the sequence.

If the data sender sending an RI-Upd does not receive an acknowledgment, or RI-Ack, from the data receiver within a specified period of time, the data sender should periodically retransmit the RI-Upd until it receives an acknowledgment from the data receiver. Once the data sender retransmits the RI-Upd a specified number of times, if it does not receive an RI-Ack, it should assume that the one-way connection on which it is the data sender is down. For more information about routers going down, see the section "Using AURP-Tr to Detect Routers Going Down" later in this chapter.

ROUTING INFORMATION ACKNOWLEDGMENT PACKET: When a data receiver receives an RI-Upd, it verifies the packet's connection ID and sequence number. The connection ID must be the same as that in the Open-Req for the connection. The sequence number must be either:

- the last sequence number received, indicating that the previous acknowledgment was lost or delayed, and that this is a duplicate RI-Upd

- the next number in the sequence, indicating that the RI-Upd contains new routing information

If the sequence number has any other value, the data receiver ignores the RI-Upd. Once the data receiver has verified the RI-Upd packet's connection ID and sequence number, it responds by sending a Routing Information Acknowledgment packet, or RI-Ack, which contains the following information:

Connection ID: The connection ID is the same as that in the RI-Upd packet.

Sequence number: The sequence number is the same as that in the RI-Upd packet.

Figure 3-12 shows a data receiver responding to an RI-Upd by sending an RI-Ack.

<<Figure 3-12 A routing-information update/acknowledgment dialog>>

When a data sender receives an RI-Ack, it verifies that the RI-Ack corresponds to the outstanding RI-Upd-that is, both packets have the same connection ID and sequence number. Once the data sender has verified the information in the RI-Ack, it responds by sending the next RI-Upd in the sequence, if any.

Routing-Information Update Events

An RI-Upd packet may contain any of five different types of routing-information update events. The following sections describe these events.

NETWORK ADDED EVENT: An exterior router sends a Network Added (NA) event under the following circumstances:

- A new network that appears in the exterior router's routing table is in the exterior router's local internet and is not hidden-that is, it is an exported network.

- The port through which an exterior router accesses a network changes from a tunneling port to another port on the router and the network is not hidden.

If a network in an exterior router's routing table becomes accessible across the tunnel, the exterior router does not send an NA event. An exterior router sends only split-horizoned routing information to other exterior routers on the tunnel.

An NA event lists the network numbers associated with the new network and the network's distance in hops. Another exterior router can request the zone information associated with the new network at any time by sending a ZI-Req, once it receives an RI-Upd containing an NA event for the network.

When using AURP-Tr, an exterior router can request zone information for new networks by setting the SZI bit in an RI-Ack that it sends in response to an RI-Upd. If a data sender receives an RI-Ack with its SZI flag set to 1, the data sender sends the zone information associated with each new network for which it sent an NA event in the RI-Upd.

Figure 3-13 shows a data receiver responding to an RI-Upd by sending an RI-Ack in which the SZI bit is set to 1, optimizing the flow of

zone information by causing the data sender to respond with a ZI-Rsp.

<<Figure 3-13 An optimized flow of zone information>>

NETWORK DELETED EVENT: An exterior router sends a Network Deleted (ND) event if an exported network that was formerly accessible through its local internet no longer appears in its routing table. An ND event lists the network numbers associated with the deleted network.

NETWORK ROUTE CHANGE EVENT: An exterior router sends a Network Route Change (NRC) event if the path to an exported network through its local internet changes to a path through a tunneling port, causing split-horizoned processing to eliminate that network's routing information. An NRC event lists the network numbers associated with the network to which the path changed.

NETWORK DISTANCE CHANGE EVENT: An exterior router sends a Network Distance Change (NDC) event if the distance to an exported network accessible through its local internet changes. An NDC event indicates the network to which the distance changed and the network's distance in hops. An exterior router must send an NDC event even if the distance to a network changes to 15 hops. The exterior router that receives an NDC event with a hop count of 15 should process that event just as it would an ND event.

ZONE NAME CHANGE EVENT: This event is reserved for future use.

Processing Update Events

According to the architectural model, a data receiver that is processing an event contained in an RI-Upd packet updates the corresponding information in its central routing table. For example, if a data receiver receives an RI-Upd containing an ND event or an NRC event, it sets the corresponding network's routing-table entry to BAD. The data receiver then initiates a notify-neighbor process, by sending RTMP data packets that identify bad entries in its routing table to routers on its local internet.

Processing Inconsistent Update Events

If the data receiver's copy of the data sender's routing table does not match that in the data sender's current routing table, it is possible that the data receiver might receive an RI-Upd containing an event that is incongruous with its current routing-table information. For example, this might occur if the information in the data sender's routing table were changing during its initial exchange of routing information with the data receiver, as described in the section

"Sending Updates Following the Initial Exchange of Routing Information" earlier in this chapter. The data receiver might receive an RI-Upd that contains an ND, NRC, or NDC event for a network not known to be in the data sender's routing table; or an NA event for a network already known to be in its routing table. The data receiver should

ignore ND and NRC events for unknown networks

process an NDC event for an unknown network as an NA event

process an NA event for a known network as an NDC event

Maintaining a Central Routing Table

According to the architectural model, an exterior router maintains a separate routing table for each other exterior router on a tunnel. In a typical implementation, however, an exterior router maintains a central routing table that contains information about each path to each network known to that exterior router-including its port, next internet router (IR), and distance in hops.

If no loops exist across a tunnel, an exterior router can reach a network that is accessible through that tunnel through only one exterior router, as shown in Figure 3-14. Such a network is accessible neither through the exterior router's local internet nor through any other exterior router on the tunnel. Thus, the central routing table would contain only one path for that network.

If a loop exists across a tunnel, an exterior router may be able to access a network through two or more exterior routers on the tunnel, or through both its local internet and an exterior router. Thus, when a loop exists across a tunnel, the central routing table may contain more than one path for each network. Figure 3-14 shows two examples of internets on which loops exist.

<<Figure 3-14 Internets with and without loops>>

Maintaining an Alternative-Paths List

If a loop exists across a tunnel and an exterior router maintains a single central routing table, that table must include an alternative-paths list for each network known to the exterior router. This alternative-paths list contains the routing information that an exterior router might otherwise maintain in separate routing tables for the other exterior routers on a tunnel. An entry for each alternative path to a network consists of the address of the alternative next IR for that network and the network's distance

through that next IR.

Because RTMP periodically retransmits information about alternative paths, the exterior router's alternative-paths list needs to provide information only about alternative paths to networks across tunneling ports. Thus, the alternative-paths list for a network provides complete information about all paths to that network across tunnels- but not necessarily about all paths through the exterior router's local internet.

An exterior router must maintain an alternative-paths list, because once a data sender has reliably sent routing information to a data receiver, the data sender does not retransmit that information. Even though a path may not currently be the optimal path to a network, an exterior router must maintain information about that path, in the event that it later becomes the optimal path.

NOTE: Zone information is unaffected by the path taken to a network. Therefore, an exterior router need not maintain duplicate zone information in the alternative-paths list.

Using the Alternative-Paths List in Event Processing

An exterior router uses its alternative-paths list when processing events.

PROCESSING A NETWORK ADDED EVENT: If an exterior router receives an NA event, it searches its central routing table for the network indicated in the event.

If the exterior router finds no entry for that network in its central routing table, it creates a new entry using the routing information contained in the NA event.

If the exterior router finds an existing entry for that network in its central routing table and the next IR for that entry is not the exterior router that sent the event, it determines whether the NA event provides a better path to that network.

If the NA event provides a better path to the network or the state of the routing-table entry for that network is BAD, the exterior router replaces the current entry with the routing information contained in the NA event. In the current entry, if the path to the network is through a tunnel, as indicated by the next IR, the exterior router transfers the current entry to the network's alternative-paths list.

If the NA event does not provide a better path to the network,

the exterior router adds the routing information contained in the NA event to the alternative-paths list for the network.

If the exterior router finds an existing entry for that network, in which the next IR is the exterior router that sent the event, the exterior router should process the NA event just as it would an NDC event.

PROCESSING A NETWORK DELETED EVENT: If an exterior router receives an ND event, it searches its central routing table for the network indicated in the event.

If the exterior router finds no entry for that network in its central routing table, it ignores the event. See the section "Processing Inconsistent Update Events" earlier in this chapter.

If the exterior router that is the data receiver determines that the exterior router that sent the ND event is the next IR for that network and there is an alternative-paths list for the network, the data receiver replaces the network's current routing information with the entry in the network's alternative-paths list that provides the shortest distance to that network and removes that entry from the network's alternative-paths list. If the network's alternative-paths list contains more than one entry providing the distance that constitutes the shortest distance to the network, the data receiver can use any of those entries.

If the exterior router that is the data receiver determines that the exterior router that sent the ND event is the next IR for that network and there is no alternative-paths list for the network, the data receiver sets the network's routing-table entry to BAD, then initiates a notify-neighbor process.

If the exterior router that is the data receiver determines that the exterior router that sent the ND event is not the next IR for that network, the data receiver searches that network's alternative-paths list for an entry in which the next IR is the data sender and removes that entry from the list.

PROCESSING A NETWORK ROUTE CHANGE EVENT: If an exterior router receives an NRC event, it processes that event as an ND event. Generally, an NRC event should not cause an exterior router to set the state of a network's routing-table entry to BAD. An NRC event indicates that the data sender has an alternative path to the network through the tunnel. The data receiver either is already aware of or will soon discover this alternative path.

PROCESSING A NETWORK DISTANCE CHANGE EVENT: If an exterior router receives an NDC event with a hop count of 15, it processes that event just as it would an ND event. Otherwise, it searches its central routing table for the network indicated in the event.

If the exterior router finds no entry for that network in its central routing table, it processes that event as an NA event.

If the exterior router that is the data receiver determines that the exterior router that sent the NDC event is the next IR for the network, the data receiver replaces the distance to that network that is currently in its central routing table with the distance indicated in the NDC event.

If the exterior router that is the data receiver determines that the exterior router that sent the NDC event is not the next IR for the network, the data receiver

replaces the distance in the corresponding entry in the network's alternative-paths list with the distance indicated in the NDC event creates an entry in the alternative-paths list that contains the routing information in the NDC event, if it finds no entry for that network in the alternative-paths list

Finally, regardless of whether the central routing table indicates that the exterior router that sent the NDC event is the network's next IR, the data receiver compares the distances in entries in the network's alternative-paths list to the distance in its central routing table. If an entry in the alternative-paths list contains a shorter path to the network, the exterior router transfers that entry to the central routing table. This ensures that the exterior router's central routing table contains the shortest path to the network.

If the data receiver replaces the entry currently in its central routing table with that in the NDC event and the current entry provides a path to the network through a tunnel, the data receiver transfers the current entry to the network's alternative-paths list.

If the data receiver transfers an entry in the network's alternative-paths list to its central routing table, it removes that entry from the alternative-paths list.

RESPONDING TO EVENTS IN THE LOCAL INTERNET: An exterior router that uses AURP must respond appropriately to events that originate in its local internet. Such events occur when the routing information for a network in the exterior router's local internet changes and another path to that network exists through the tunnel. An exterior router

handles such events as follows:

If the exterior router replaces the current routing-table entry for a network with routing information provided by an event originating in its local internet-that is, provided by RTMP-and the current path to the network is through a tunnel, the exterior router transfers the current entry to the network's alternative-paths list.

If the exterior router sets the state of a routing-table entry to BAD or removes an entry from its central routing table, the exterior router replaces that entry with the entry in the alternative-paths list that provides the shortest distance to the network in the entry being replaced.

If the distance to a network in the exterior router's local internet changes, the exterior router compares the distances in entries in the network's alternative-paths list to the distance in its central routing table. If an entry in the alternative-paths list provides a shorter distance to the network, the exterior router transfers that entry to its central routing table. This ensures that the exterior router's central routing table contains the shortest path to the network.

Router-Down Notification

Prior to going down, or becoming inactive, an exterior router must notify all other exterior routers in its informed-routers list that it is going down. An exterior router does this by using the underlying transport-layer service to close its connection with each exterior router.

Sending a Router Down Packet

Optionally, an exterior router can send a Router Down packet, or RD packet, to each exterior router before it goes down. An RD packet contains an error code that indicates the exterior router's reason for terminating its connection with each exterior router.

Generally, only the exterior router functioning as the data sender on a one-way connection sends RD packets. However, if just a single one-way connection exists between two exterior routers, the exterior router functioning as the data receiver on that connection can send an RD packet.

Using AURP-Tr to Notify Other Routers That a Router Is Going Down

When using AURP-Tr, an exterior router sends an RD packet to

notify another exterior router that it is terminating a connection

pass an error code that indicates its reason for terminating the connection

As shown in Figure 3-15, once the data receiver verifies the RD packet's connection ID, it acknowledges that it received the RD packet by sending an RI-Ack. Then, the data sender terminates the connection.

<<Figure 3-15 Acknowledging an RD packet>>

If a Router Goes Down Without Notifying Other Routers

If an exterior router crashes or goes down without sending an RD packet, or becomes inaccessible due to a network problem, other exterior routers on the tunnel must be able to discover that the exterior router is down. Generally, the underlying transport-layer service provides a mechanism for informing an exterior router that an exterior router in its informed-routers list has gone down or become inaccessible.

If an exterior router determines that another exterior router is down, it must

remove that exterior router from its informed-routers list

remove that exterior router's routing information from all of its routing tables

close any one-way connections with that exterior router

If an exterior router rediscovers an exterior router that had previously gone down, it must again exchange initial routing information with that exterior router.

Using AURP-Tr to Detect Routers Going Down

An exterior router using AURP-Tr associates a last-heard-from timer with each exterior router from which it has received routing information-that is, with each one-way connection on which it is the data receiver. Each time the exterior router receives an RI-Rsp, RI-Upd, or ZI-Rsp over a connection-verifying that its connection with the data sender is still active-it resets the last-heard-from timer for that connection.

For each one-way connection on which it is the data receiver, the exterior router has a last-heard-from timeout value. If a

connection's last-heard-from timer reaches that timeout value, the data receiver sends a Tickle packet over that connection. If the data sender on the connection is still accessible, it responds with a Tickle-Ack, as shown in Figure 3-16. When the data receiver receives the Tickle-Ack, it resets the last-heard-from timer for that connection. If the data receiver receives no Tickle-Ack-even after retransmitting the Tickle several times-it assumes that the connection is down.

<<Figure 3-16 Acknowledging a Tickle packet>>

If the exterior router determines that the connection is down and an associated one-way connection exists on which it is the data sender, it should send a null RI-Upd over that connection to determine whether that one-way connection is still active.

If the data receiver on the connection is still accessible, it responds with an RI-Ack, as shown in Figure 3-17. If the data sender receives no RI-Ack-even after retransmitting the null RI-Upd several times-it determines that the one-way connection on which it is the data sender is also down.

<<Figure 3-17 Acknowledging an RI-Upd packet>>

The value of the last-heard-from timeout should be configurable. The minimum last-heard-from timeout should be 30 seconds. If a connection's last-heard-from timeout is greater than two minutes-the tickle-before-data time-and the data receiver has not reset the connection's last-heard-from timer for at least this tickle-before-data time, the data receiver must send a Tickle to the data sender before forwarding an AppleTalk data packet to it. If the data sender on the connection is still accessible, it responds with a Tickle-Ack. When the data receiver receives the Tickle-Ack, it resets the last-heard-from timer for that connection. If the data receiver receives no Tickle-Ack, even after retransmitting the Tickle, it assumes that the data sender is no longer accessible and closes the connection.

Obtaining Zone Information

AURP supports two commands that allow an exterior router to obtain routing information for zones rather than for networks-the Get Domain Zone List (GDZL) command and the Get Zone Nets (GZN) command. These commands constitute request/response transactions, and are similar to ZI-Req and ZI-Rsp. An exterior router sends these commands unsequenced over a connection.

NOTE: Under AURP, the implementation of the Get Domain Zone List command and the Get Zone Nets command in an exterior router is

optional. However, an exterior router must at least be able to return an error to a GDZL-Req or a GZN-Req.

Get Domain Zone List Command

The Get Domain Zone List command, or GDZL, allows an exterior router to obtain a zone list for an internet. As shown in Figure 3-18, GDZL functions similarly to the ZIP GetZoneList command. However, a GDZL-Rsp returns a split-horizoned zone list-that is, it returns only the zones in the exterior router's local internet, rather than the exterior router's entire zone list. A GDZL-Rsp does not return zones in networks that are accessible through the tunnel, unless those zones are also in networks that are accessible through the exterior router's local internet.

<<Figure 3-18 Get Domain Zone List request/response dialog>>

Get Zone Nets Command

The Get Zone Nets command, or GZN, allows an exterior router to obtain a list of the networks in an exterior router's local internet that are associated with a particular zone name. As shown in Figure 3-19, GZN functions similarly to ZI-Req and ZI-Rsp, but a GZN-Req packet contains a single zone name and GZN-Rsp packets contain network tuples that have the same format as the tuples in an RI-Rsp. A GZN-Rsp returns network tuples only for networks that are accessible through the exterior router's local internet.

<<Figure 3-19 Get Zone Nets request/response dialog>>

Using AURP-Tr to Process Sequence Numbers

When an exterior router acting as a data receiver sends an Open-Req to establish a one-way connection, it expects the data sender to respond by sending sequenced data packets, starting with the sequence number 1. The data receiver's response to each packet that it receives depends on the packet's sequence number:

Whenever the data receiver receives an RI-Rsp, RI-Upd, or RD packet that has the expected sequence number and connection ID, it sends an RI-Ack packet having that sequence number, then increases the sequence number that it expects by one, until the sequence number reaches 65,535. Sequence numbers wrap around and the sequence number 0 is reserved, so the sequence number 1 follows 65,535. Thus, when comparing sequence numbers, an exterior router interprets the sequence number 65,535 as one less than the sequence number 1.

If the data receiver expects sequence number n and receives a packet with the sequence number n-1, that packet was delayed and is a duplicate of another packet already received. The data receiver must retransmit an RI-Ack packet, because the data sender may not have received the RI-Ack packet previously sent-that is, the RI-Ack may have been lost.

If the data receiver expects sequence number n and receives a packet with the sequence number n+1, it should discard the packet and terminate the one-way connection on which it is the data receiver. Because AURP-Tr supports only one outstanding transaction at a time, the receipt of such a packet indicates that the connection is out of sync.

If the data receiver expects sequence number n and receives a packet with a sequence number other than n-1, n, or n+1, the packet was delayed and is a duplicate of another packet already received. The data receiver need not send an RI-Ack, because the data sender must have received an RI-Ack for that sequence number prior to sending a packet with the sequence number n-1. The data receiver should discard the packet.

NOTE: If the sequence numbers have not wrapped around, a sequence number greater than n+1 indicates that the connection is out of sync.

Using AURP-Tr to Process Connection IDs

If an exterior router acting as either a data receiver or a data sender on a one-way connection receives a packet from an exterior router with which it has a one-way connection, it checks the connection ID in the packet to verify that the packet was sent on that connection. If the packet contains a connection ID that does not match that expected for the connection, the exterior router discards the packet.

If a data sender receives an Open-Req from an exterior router with which it already has a connection and the connection ID does not match that for the connection already established, it should not discard the packet without verifying whether the connection is still active. The receipt of such a packet may indicate that the data receiver on the connection has been restarted and has opened a new one-way connection, without first terminating its original connection. The exterior router acting as the data sender should send a null RI-Upd over the connection to determine whether it is still active. If the data sender receives an RI-Ack in response to the null RI-Upd, it discards the Open-Req and the original connection remains active. If the data sender receives no RI-Ack after retransmitting the null RI-Upd, it closes the original connection, then sends an

Open-Rsp to the next Open-Req received.

NOTE: An exterior router can act as the data sender on only a single one-way connection between itself and a given exterior router. That is, multiple one-way connections in the same direction cannot exist between two exterior routers.

When establishing a one-way connection with a given data sender, a data receiver using AURP-Tr must send an Open-Req that has a different connection ID from that used in its last connection with the data sender. Otherwise, if the last connection to the data sender had terminated abnormally and the new connection used the same connection ID, the data sender might determine that the last connection was still active and interpret the Open-Req as a retransmission of the Open-Req for the last connection. The data sender might respond to the Open-Req by sending an Open-Rsp or ignore the Open-Req, but would not open a new connection.

If a data receiver's implementation of AURP-Tr cannot guarantee the use of different connection IDs on successive connections with a given data sender, the data receiver must send an RI-Req immediately after it establishes a connection with a data sender. If the data sender already has a connection with the data receiver, it will send an RI-Rsp with a sequence number other than 1. The data receiver should then terminate that connection and open a new connection using a different connection ID.

Using Retransmission Timers Under AURP-Tr

When an AppleTalk tunnel exists through a foreign network's internet, the delay and loss characteristics of the tunnel's underlying foreign network system complicate the setting of retransmission timers. A physical connection can be built between two exterior routers using different media-for example, a single Ethernet LAN, a fast point-to-point link, an IP internet, or a slow link over an asynchronous modem. It is important to minimize performance degradation due to

- packets being dropped or delayed by the underlying foreign network system

- the inefficient use of the underlying foreign network system's resources due to excessive retransmissions

Most higher-level transport-layer services provide guaranteed packet delivery. It is not necessary to retransmit AURP packets when using such transport-layer services. When using AURP-Tr, an exterior router should employ an adaptive retransmission algorithm whenever possible. An adaptive retransmission strategy like that used in TCP

maintains the estimated times required to send a packet and receive an acknowledgment-that is, average round-trip times

maintains standard deviations from the average round-trip times

derives retransmission timers from the average round-trip times While AURP does not specify an adaptive retransmission algorithm, the use of such an algorithm is recommended.

NOTE: Often, long intervals exist between AURP packets sent successively on a connection by an exterior router-for example, between RI-Upd packets. Therefore, an adaptive retransmission algorithm used with AURP should give more weight to packets sent recently over a connection than would be appropriate for a general data-stream protocol like TCP.

When an exterior router initially opens a connection, no transaction history is available. It is recommended that the retransmission algorithm use a truncated, exponential backoff scheme for the initial Open-Req sequence, because the exterior router with which the data receiver is establishing a connection may be inaccessible or down. An exterior router should not retransmit an Open-Req at a rate faster than once every two seconds.

Hiding Local Networks From Remote Networks

As described in the section "Hiding Local Networks From Tunnels" in Chapter 2, a network administrator can configure an exterior router to hide specific networks in its local internet from networks connected to other exterior routers on the tunnel. When exchanging routing information with other exterior routers on the tunnel, the exterior router exports no routing information for hidden networks in its local internet to exterior routers from which those networks are hidden.

An exterior router using AURP does not include routing information for hidden networks in RI-Rsp, RI-Upd, or GZN-Rsp packets sent to exterior routers from which those networks are hidden. The exterior router also excludes from GDZL-Rsp packets any zones that appear only in the zone lists of hidden networks.

To maintain network-level security, an exterior router should discard any AppleTalk data packet sent to a network in its local internet by an exterior router from which that network is hidden.

NOTE: An exterior router hides a network by excluding the routing information for that network from RI-Rsp, RI-Upd, GZN-Rsp, and GDZL-Rsp packets. However, network management packets-such as RTMP Route

Data Response (RDR) packets that are not split horizoned, and Simple Network Management Protocol (SNMP) packets-should include the routing information for hidden networks. For detailed information about the effects of AURP on network management, see the section "Network Management" in Chapter 4.

AURP Packet Format

An exterior router encapsulates both AURP packets and AppleTalk data packets using the same headers. Before forwarding AURP packets across a tunnel, an exterior router encapsulates the AURP packets in packets of the tunnel's underlying foreign network system-by adding the headers required by that network system. For more information about these headers, see the sections "Forwarding Data," "AppleTalk Data-Packet Format," and "AppleTalk Data-Packet Format for IP Tunneling" in Chapter 2.

When using AURP-Tr in conjunction with TCP/IP, an exterior router encapsulates AURP packets in UDP packets prior to forwarding them across an IP tunnel through UDP port 387. When another exterior router on the tunnel receives the UDP packets at UDP port 387, it decapsulates the packets.

Domain Headers in AURP Packets

When forwarding AURP packets across a tunnel, an exterior router adds a domain header immediately preceding each packet. A domain header contains additional addressing information, including its source domain identifier and destination domain identifier (DI). The last two bytes of the domain header are set to 0003, indicating that the packet is an AURP packet rather than an AppleTalk packet. AURP data follows the domain header. Figure 3-20 shows the protocol headers, the domain header, and the routing data header that encapsulate a routing data packet sent across an IP tunnel.

<<Figure 3-20 A routing data packet on an IP tunnel>>

An exterior router interprets the domain identifiers in the domain header of an AURP packet differently from those in the domain headers of an AppleTalk data packet. Only network entities with AppleTalk addresses have domain identifiers associated with them. Exterior routers do not have AppleTalk addresses on the tunnel-thus, they do not have true domain identifiers.

DESTINATION DOMAIN IDENTIFIER: The destination DI in an AURP packet's domain header is the DI that is associated with any network numbers corresponding to networks that reside in the receiving exterior router's domain. Only ZI-Req packets include such network numbers.

Whenever possible, a domain header should specify a destination DI-that is, the DI for the networks that reside in the domain of the exterior router that is to receive the packet. When an exterior router sends an Open-Req to open a connection, the destination DI is not yet known. However, under the current version of AURP, the exterior router can either derive the destination DI from the destination's IP address or, on point-to-point links, include the null DI.

SOURCE DOMAIN IDENTIFIER: The source DI in an AURP packet's domain header is the DI that is associated with any network numbers corresponding to networks that reside in the sending exterior router's domain. RI-Rsp, RI-Upd, ZI-Rsp, and GZN-Rsp packets include such network numbers. A domain header should always specify a source DI-that is, the DI for the networks that reside in the domain of the exterior router that is sending the packet.

Routing Data Headers in AURP Packets

The routing data header that immediately precedes the AURP data in a routing data packet consists of an AURP-Tr header and an AURP header. The AURP-Tr header consists of the following fields:

Connection ID: The contents of this two-byte field identify the specific one-way connection to which a packet belongs.

Sequence number: The contents of this two-byte field identify an individual packet on a connection.

The AURP header consists of these fields:

Command code: This two-byte field identifies the command type. For information about command types, see the next section, "Command Types."

Flags: This two-byte field may contain different flags, depending on the command code. For information about flags, see the section "Routing Flags" later in this chapter.

Command Types

AURP defines the command types shown in Table 3-1:

Table 3-1 Command types

Command type	Abbreviation	Command code	Subcode
Routing Information Request	RI-Req	1	-
Routing Information Response	RI-Rsp	2	-
Routing Information Acknowledgment	RI-Ack	3	-
Routing Information Update	RI-Upd	4	-
Router Dow	RD	5	-
Zone Information Request	ZI-Req	6	1
Zone Information Response	ZI-Rsp	7	1 and 2
Get Zones Net Request	GZN-Req	6	3
Get Zones Net Response	GZN-Rsp	7	3
Get Domain Zone List Request	GDZL-Req	6	4
Get Domain Zone List Response	GDZL-Rsp	7	4
Open Request	Open-Req	8	-
Open Response	Open-Rsp	9	-
Tickle	-	14	-
Tickle Acknowledgment	Tickle-Ack	15	-

Routing Flags

AURP defines the flags shown in Table 3-2. All other flags are reserved. A data sender should set reserved flags to 0. A data receiver should ignore reserved flags.

Table 3-2 Flags

Flag	Event	Command types	Bit
Send update information (SUI) flag	NA	Open-Req and RI-Req	14
Send update information (SUI) flag	ND and NRC	Open-Req and RI-Req	13
Send update information (SUI) flag	NDC	Open-Req and RI-Req	12
Send update information (SUI) flag	ZC	Open-Req and RI-Req	11
Last flag	-	RI-Rsp and GDZL-Rsp	15
Remapping active flag	-	Open-Rsp	14
Hop-count reduction active flag	-	Open-Rsp	13
Reserved environment flags	-	-	12 and 11
Send zone information (SZI) flag	-	RI-Ack	14

Figure 3-21 shows the routing flags in Open-Req and RI-Req packets.

<<Figure 3-21 Routing flags in Open-Req and RI-Req packets>>

Figure 3-22 shows the routing flags in all packets other than Open-Req and RI-Req packets.

<<Figure 3-22 Routing flags in other packets>>

Open Request Packet

An Open-Req packet initiates the establishment of a one-way connection with a data sender. Figure 3-23 shows the format of an Open-Req packet. When sending an Open-Req packet, an exterior router inserts the next available connection ID in the packet's AURP-Tr header and sets its sequence number to 0. The AURP header of an Open-Req contains the command code 8. Its flag bytes contain send update information (SUI) flags. For the current version of AURP, the version number is 1.

An Open-Req packet's option data field contains

- an option count-indicating the number of option tuples to follow
- the option tuples

When the data sender receives an Open-Req, it can discard the option tuples for any options it does not implement. For information about option tuples, see the section "Option Tuples" later in this chapter.

<<Figure 3-23 Open-Req packet format>>

Open Response Packet

When the data sender receives an Open-Req, it responds by sending an Open-Rsp packet to establish a one-way connection with the data receiver. Figure 3-24 shows the format of an Open-Rsp packet. In its AURP-Tr header, an Open-Rsp packet contains the connection ID from the associated Open-Req packet and the sequence number 0. The AURP header of an Open-Rsp contains the command code 9 and its flag bytes contain environment flags that provide information about the data sender's environment-such as whether network-number remapping or hop-count reduction is active. For information about network-number remapping and hop-count reduction, see the sections "Network-Number Remapping" and "Hop-Count Reduction," respectively, in Chapter 4.

<<Figure 3-24 Open-Rsp packet format>>

An Open-Rsp packet's option data field contains

- a two-byte field that indicates either
  - the nominal rate at which the data sender sends updates-in multiples of ten seconds
  - an error code-which is a negative number-if the data sender cannot accept the connection

an option count-indicating the number of option tuples to follow

the option tuples

For information about error codes, see the section "Error Codes" later in this chapter. For information about option tuples, see the next section, "Option Tuples."

Option Tuples

Both Open-Req and Open-Rsp packets contain option tuples. An option tuple contains a one-byte length field that indicates the length of the remainder of the tuple, a one-byte type code, and an optional data field, as shown in Figure 3-25.

<<Figure 3-25 Option tuples>>

AURP currently defines the option-type codes shown in Table 3-3:

Table 3-3 Option-type codes

Option types	Type codes
Authentication	1
Reserved for future use	2-255

Routing Information Request Packet

An RI-Req packet requests the data sender to send RI-Rsp packets. Figure 3-26 shows the format for an RI-Req packet. When sending an RI-Req packet, an exterior router inserts the connection ID for the connection on which it is the data receiver in the packet's AURP-Tr header and sets the packet's sequence number to 0. The AURP header of an RI-Req contains the command code 1 and its flag bytes contain the send update information (SUI) flags.

<<Figure 3-26 RI-Req packet format>>

Routing Information Response Packet

When the data sender receives an RI-Req, it responds by sending a sequence of RI-Rsp packets. Figure 3-27 shows the format of an RI-Rsp packet. When sending an RI-Rsp packet, a data sender inserts the connection ID from the associated RI-Req in the RI-Rsp packet's AURP-Tr header and sets its sequence number to the next number in the sequence. The AURP header of an RI-Rsp packet contains the command code 2. In the last packet in a sequence of RI-Rsp packets, the

last-flag bit is set to 1.

<<Figure 3-27 RI-Rsp packet format>>

An RI-Rsp packet's routing data field contains zero or more routing tuples, which have a format similar to those in RTMP packets. An AURP tuple for a nonextended network is different from an RTMP tuple for an extended network in one respect-the range flag, or the sixth byte, in an AURP tuple for a nonextended network is set to 0. Figure 3-28 shows nonextended and extended network tuples in an RI-Rsp packet.

<<Figure 3-28 Nonextended and extended network tuples>>

Routing Information Acknowledgment Packet

When a data receiver receives an RI-Rsp, RI-Upd, or RD packet, it responds by sending an RI-Ack packet. Figure 3-29 shows the format of an RI-Ack packet. When sending an RI-Ack packet, a data receiver inserts the connection ID and sequence number from the associated RI-Rsp, RI-Upd, or RD packet in the RI-Ack packet's AURP-Tr header. The AURP header of an RI-Ack contains the command code 3. If the data receiver sends an RI-Ack using AURP-Tr, in response to an RI-Rsp or RI-Upd packet that contains an NA event, its flag bytes contain the send zone information flag. An RI-Ack packet contains no data.

<<Figure 3-29 RI-Ack packet format>>

Routing Information Update Packet

The occurrence of specified events requires the data sender to send an RI-Upd packet. Figure 3-30 shows the format of an RI-Upd packet. When sending an RI-Upd packet, a data sender inserts the connection ID for the current connection in the RI-Upd packet's AURP-Tr header and sets its sequence number to the next number in the sequence. The AURP header of an RI-Upd contains the command code 4 and its flag bytes are set to 0.

<<Figure 3-30 RI-Upd packet format>>

An RI-Upd packet's data field contains one or more event tuples. An event tuple for a nonextended network consists of a one-byte event code, the network number, and the distance to that network. An event tuple for an extended network consists of a one-byte event code, the first network number in the range of network numbers, the distance to the network, and the last network number in the range of network numbers. Figure 3-31 shows nonextended and extended network tuples in an RI-Upd packet.

<<Figure 3-31 Nonextended and extended network event tuples>>

AURP currently defines the event codes shown in Table 3-4:

Table 3-4 Event codes

Event	Abbreviation	Event code
Null event		0
Network Added event	NA	1
Network Deleted event	ND	2
Network Route Change event	NRC	3
Network Distance Change event	NDC	4
Zone Change event	ZC	5

A null event tuple contains no event data. The format of NA, ND, NRC, and NDC event tuples differs, depending on whether the event pertains to a nonextended or an extended network. The distance field does not apply to ND or NRC event tuples and should be set to 0. The ZC event tuple is not yet defined.

An RI-Upd packet should never contain two events that pertain to the same network. However, to ensure consistent behavior in the event that an exterior router receives a packet containing multiple events for one network, an exterior router should always process events in the order in which they occur in the RI-Upd packet. Thus, if an exterior router were to receive an RI-Upd that contained an NA event, then an ND event for the same network, the exterior router would delete the network from its routing table.

Router Down Packet

An exterior router should send an RD packet before it goes down. Figure 3-32 shows the format of an RD packet. When sending an RD packet, an exterior router inserts the connection ID for the current connection in the RD packet's AURP-Tr header. If the data sender sends an RD packet, it sets its sequence number to the next number in the sequence. If the data receiver sends an RD packet, it sets its sequence number to 0. The AURP header of an RD packet contains the command code 5 and its flag bytes are set to 0.

<<Figure 3-32 RD packet format>>

An RD packet's data field contains a two-byte error code that indicates the exterior router's reason for going down. For information about the error codes, see the section "Error Codes" later in this chapter.

Zone Information Request/Response Transactions

An exterior router returns information about its zones through request/response transactions. Three types of zone requests-ZI-Req, GDZL-Req, and GZN-Req-share the same command code and have subcodes that indicate the actual request type. Three types of zone responses-ZI-Rsp, GDZL-Rsp, and GZN-Rsp-share another command code and have subcodes that indicate the actual response type.

ZONE INFORMATION REQUEST PACKET: A ZI-Req packet causes the data sender to send ZI-Rsp packets. Figure 3-33 shows the format of a ZI-Req packet. When sending a ZI-Req packet, an exterior router inserts the connection ID for the connection on which it is the data receiver in the packet's AURP-Tr header and sets the packet's sequence number to 0. The AURP header of a ZI-Req contains the command code 6 and its flag bytes are set to 0.

<<Figure 3-33 ZI-Req packet format>>

A ZI-Req packet's data field contains the subcode 1 and a two-byte network number for each network about which the exterior router is requesting zone information. The network number for an extended network is the first network number in its range of network numbers.

ZONE INFORMATION RESPONSE PACKET: There are two types of ZI-Rsp packets-nonextended ZI-Rsp packets and extended ZI-Rsp packets. The format of a nonextended ZI-Rsp packet is similar to that of a nonextended AppleTalk ZIP Reply packet. When the data sender receives a ZI-Req and the zone list for the network or networks for which that ZI-Req requested zone information fits in one ZI-Rsp packet, it sends a nonextended ZI-Rsp.

An extended ZI-Rsp packet is similar to an extended AppleTalk ZIP Reply packet. When the data sender receives a ZI-Req and the zone list for a network about which that ZI-Req requested zone information does not fit in a single ZI-Rsp packet, it sends a sequence of extended ZI-Rsp packets.

Figure 3-34 shows the format of a ZI-Rsp packet. When sending a ZI-Rsp packet, a data sender inserts the connection ID from the associated ZI-Req packet in the packet's AURP-Tr header and sets the packet's sequence number to 0. A ZI-Rsp packet's AURP header contains the command code 7 and its flag bytes are set to 0. The subcode 1 indicates a nonextended ZI-Rsp packet, while the subcode 2 indicates an extended ZI-Rsp packet.

<<Figure 3-34 ZI-Rsp packet format>>

A ZI-Rsp packet's data field contains the requested zone information. Its format is similar to that of a ZIP Reply packet.

In a nonextended ZI-Rsp packet, the first two bytes of the data field should indicate the number of tuples contained in the packet, while the remaining bytes constitute network number/zone name tuples. Within the packet, all of the tuples for a given network must be contiguous. NOTE: When sending a nonextended ZI-Rsp packet, an exterior router should attempt to specify the correct number of zone tuples. However, an exterior router receiving a nonextended ZI-Rsp packet should process all tuples contained in the packet, regardless of the number indicated in the header.

Network number/zone name tuples in a nonextended ZI-Rsp packet can use either the long tuple format or the optimized tuple format. A long network number/zone name tuple contains a network number, followed by the length of the zone name, and the zone name.

Using the optimized tuple format, an exterior router can compress a nonextended ZI-Rsp packet in which more than one network contains the same zone name in its zone list. If the high-order bit of the length byte for a given zone name is set to 1, the following 15 bits represent an offset from the length byte of the first zone name in the packet's data field to the actual location of the zone name length and the zone name. Whenever possible, it is recommended that an exterior router send optimized ZI-Rsp packets. All exterior routers must be able to receive optimized ZI-Rsp packets.

In an extended ZI-Rsp packet, the first two bytes of the data field indicate the total number of tuples in the zone list for the network or networks for which the corresponding ZI-Req requested zone information. The remaining bytes in the data field of an extended ZI-Rsp packet consist of network number/zone name tuples. All tuples in a single extended ZI-Rsp packet must contain the same network number. However, for consistency with the format of network number/zone name tuples in nonextended ZI-Rsp packets, the network number precedes each zone name in an extended ZI-Rsp packet. Duplicate zone names never exist in extended ZI-Rsp packets- therefore, extended ZI-Rsp packets use the long tuple format, rather than the optimized tuple format.

Figure 3-35 shows the long tuple and optimized tuple formats for a ZI-Rsp packet.

<<Figure 3-35 Long and optimized tuple formats>>

GET DOMAIN ZONE LIST REQUEST PACKET: A Get Domain Zone List Request packet, or GDZL-Req, requests the data sender to send GDZL-Rsp

packets. Figure 3-36 shows the format for a GDZL-Req packet. When sending a GDZL-Req packet, an exterior router inserts the connection ID for the connection on which it is the data receiver in the packet's AURP-Tr header and sets its sequence number to 0. The AURP header of a GDZL-Req contains the command code 6 and its flag bytes are set to 0.

<<Figure 3-36 GDZL-Req packet format>>

A GDZL-Req packet's data field contains the subcode 4 and the start index in the data sender's zone list at which to begin returning GDZL-Rsp packets.

GET DOMAIN ZONE LIST RESPONSE PACKET: When the data sender receives a GDZL-Req, it responds by sending a GDZL-Rsp packet. Figure 3-37 shows the format of a GDZL-Rsp packet. When sending a GDZL-Rsp packet, a data sender inserts the connection ID from the associated GDZL-Req packet in the packet's AURP-Tr header and sets its sequence number to 0. The AURP header of a GDZL-Rsp contains the command code 7 and its flag bytes are set to 0, except in the last packet containing zone information, which has its last flag set to 1.

<<Figure 3-37 GDZL-Rsp packet format>>

A GDZL-Rsp packet's data field contains the subcode 4, the start index from the associated GDZL-Req, and the zone list. If the data sender does not support the GDZL-Req, it should set the start index to -1.

GET ZONES NET REQUEST PACKET: A Get Zones Net Request packet, or GZN-Req, requests the data sender to send zone information for one specific zone. Figure 3-38 shows the format of a GZN-Req packet. When sending a GZN-Req packet, an exterior router inserts the connection ID for the connection on which it is the data receiver in the packet's AURP-Tr header and sets its sequence number to 0. The AURP header of a GZN-Req contains the command code 6 and its flag bytes are set to 0.

<<Figure 3-38 GZN-Req packet format>>

A GZN-Req packet's data field contains the subcode 3 and the name of the zone about which the GZN-Req is requesting zone information.

GET ZONES NET RESPONSE PACKET: When the data sender receives a GZN-Req, it responds by sending a GZN-Rsp packet, containing the requested zone information. Figure 3-39 shows the format of a GZN-Rsp packet. When sending a GZN-Rsp packet, a data sender inserts the connection ID from the associated GZN-Req packet in the GZN-Rsp

packet's AURP-Tr header and sets the GZN-Rsp packet's sequence number to 0. The AURP header of a GZN-Rsp contains the command code 7 and its flag bytes are set to 0.

<<Figure 3-39 GZN-Rsp packet format>>

A GZN-Rsp packet's data field contains the subcode 3, the zone name from the associated GZN-Req, the total number of network tuples for that zone, and as many network tuples as can fit in the packet. These tuples have the same format as those in RI-Rsp packets. If the data sender has no information about the zone, it returns a GZN-Rsp in which the number of network tuples is 0. If the data sender does not support the GZN-Req, it should set the number of network tuples to -1.

TICKLE PACKET: The data receiver sends a Tickle packet to verify that the data received from the data sender is still valid. Figure 3-40 shows the format of a Tickle packet. When sending a Tickle packet, an exterior router inserts the connection ID for the connection on which it is the data receiver in the packet's AURP-Tr header and sets its sequence number to 0. The AURP header of a Tickle contains the command code 14 and its flag bytes are set to 0. A Tickle packet contains no data.

<<Figure 3-40 Tickle packet format>>

TICKLE ACKNOWLEDGMENT PACKET: When the data sender receives a Tickle, it responds by sending a Tickle-Ack packet. Figure 3-41 shows the format of a Tickle-Ack. When sending a Tickle-Ack, a data sender inserts the connection ID from the associated Tickle in the Tickle-Ack packet's AURP-Tr header and sets its sequence number to 0. The AURP header of a Tickle-Ack packet contains the command code 15 and its flag bytes are set to 0. A Tickle-Ack packet contains no data.

<<Figure 3-41 Tickle-Ack packet format>>

Error Codes

Open-Rsp and RD packets contain error codes. AURP currently defines the error codes listed in Table 3-5.

Table 3-5 Error codes

Error code	Error
-1	Normal connection close
-2	Routing loop detected
-3	Connection out of sync

-4	Option-negotiation error
-5	Invalid version number
-6	Insufficient resources for connection
-7	Authentication error

4. REPRESENTING WIDE AREA NETWORK INFORMATION

This chapter describes optional features of AURP-some of which can also be implemented on routers that use RTMP rather than AURP for routing-information propagation. It provides detailed information about the presentation of wide area network information by exterior routers to nodes on their local internets or to other exterior routers, including:

- basic security-both network hiding and device hiding
- remapping of remote network numbers
- internet clustering
- loop detection
- hop-count reduction
- hop-count weighting
- backup paths
- network management

Network Hiding

An exterior router can hide networks by importing or exporting routing information only about specific networks.

Importing Routing Information About Specific Networks

A network administrator can configure a tunneling port on an exterior router to import only a subset of the routing information that it receives through the tunnel. To do so, the administrator hides specific networks connected to other exterior routers on the tunnel from the exterior router's local internet. For example, an exterior router can import only that routing information received from specific exterior routers, or routing information for networks in a specific network range or zone. By importing routing information only about specific networks, an exterior router can greatly reduce

the amount of routing information maintained by routers on its local internet

the number of zones and devices that are visible to devices on its local internet

Exporting Routing Information About Specific Networks

A network administrator can configure a tunneling port on an exterior router to export only a subset of its local internet's routing information-by hiding from other exterior routers on the tunnel specific networks in its local internet. For more information about hiding networks from other exterior routers, see the section "Hiding Local Networks From Tunnels" in Chapter 2.

Device Hiding

A router can prevent a device in its local internet from being visible to other nodes on a specific part or all other parts of the internet by not forwarding Name Binding Protocol (NBP) LkUp-Reply packets from that device. Hiding a device prevents nodes on the part of the internet from which it is hidden from knowing the name of the hidden device, making it more difficult for those nodes to access the hidden device. Any AppleTalk Phase 2 router can hide devices.

Advantages and Disadvantages

Device hiding is a flexible security mechanism that is appropriate for organizations that do not require true device-specific security. It is not a substitute for device-specific security. Device hiding can provide a degree of security on devices for which no other form of security exists-such as LaserWriter printers.

A user can write a program that can obtain access to a hidden device using its AppleTalk address. Device hiding cannot secure a device from a user that is not using NBP to access the device.

Device hiding does not provide true device-specific security. Many devices require device-specific security-for example, AppleShare file servers. Device-specific security can provide various levels of security, and may allow a network administrator to grant access privileges based on registered users and groups.

Configuring Device Hiding on a Port

When configuring a port on a router that implements device hiding, a network administrator should be able to hide any device that is accessible through that port from the other ports on the router. The

device being hidden need not reside on the network connected directly to the port being configured.

An administrator should be able to specify the ports from which to hide a device-either specific ports or all other ports.

When hiding devices, an administrator should be able to specify that a list of devices either be hidden or visible. The device list should include device names and device types-for example, We-B-Nets:AFPServer. An administrator should also be able to hide all devices of a given type-for example, all LaserWriter printers-or all devices of all types.

Filtering NBP LkUp-Reply Packets

To implement device hiding, a router selectively filters NBP LkUp-Reply packets. When a port's configuration specifies that devices accessible through the port be hidden, the router

- monitors all NBP LkUp-Reply packets received through that port-called the incoming port
- determines the port through which it is to forward such a packet-called the outgoing port
- obtains-from the port configuration for the incoming port-the list of devices to be hidden from the outgoing port
- determines whether it should filter all or part of an NBP LkUp-Reply packet
  - If a port's configuration does not specify that devices be hidden from the outgoing port, the router forwards the packet.
  - If a port's configuration specifies that devices be hidden from the outgoing port, the router checks each tuple in the NBP LkUp-Reply packet to determine whether it is from a device in the port's list of hidden devices. It marks tuples from hidden devices for deletion. Once the router scans the entire packet, it forwards the packet if no tuples were marked for deletion; it discards the packet if all tuples were marked for deletion; or, if only some tuples were marked for deletion, it rebuilds the packet without the tuples marked for deletion, then forwards the packet.

When the router rebuilds a packet, it adjusts the tuple count in the packet's NBP header to reflect the number of tuples remaining. If a rebuilt packet's DDP header contains a nonzero checksum, the router

verifies the original checksum, then sets it to 0.

This device-hiding scheme can handle both NBP Lookups and NBP Confirms, because a node responds to requests of either type with a LkUp-Reply packet.

LkUp-Reply packets do not contain the names of zones in which devices reside. Thus, if two devices having the same name and type are accessible through a port, a network administrator can hide both devices or neither device, but not just one of the devices.

When configuring ports on routers through which redundant paths to a device exist, a network administrator must hide that device on at least one port on each path to that device. Otherwise, only a router on which such a port was configured to hide the device would filter LkUp-Reply packets from the device. A router on which such a port was not configured to hide the device would not filter its LkUp-Reply packets. Figure 4-1 shows the proper configuration of device hiding when a loop exists on the internet.

<<Figure 4-1 Device hiding when a loop exists on the internet>>

Resolving Network-Numbering Conflicts

In addition to interconnecting different parts of one organization's internet, tunnels can interconnect the internets of multiple organizations. Each organization administrates its internet independently. Therefore, conflicting network numbers may exist on the internets, especially when many internets are interconnected. The following sections describe the methods that AURP uses to resolve various problems due to conflicting network numbers.

Network-Number Remapping

Network-number remapping resolves network-numbering conflicts, allowing network administrators to build very large internets. When configuring a port on an exterior router, an administrator can specify a range of AppleTalk network numbers to be used for imported networks-that is, networks that are accessible through half-routing or tunneling ports, for which the exterior router imports routing information from other exterior routers. The remapping range-the range of network numbers reserved for network-number remapping-must not conflict with any network numbers already in use on the exterior router's local internet.

The exterior router maps the network numbers in incoming packets into the remapping range. It converts remapped network numbers back to their actual network numbers for outgoing packets. To nodes and

routers within the exterior router's local internet, packets containing remapped network numbers apparently originate from or are being sent to networks having numbers in the remapping range.

UNIQUE IDENTIFIERS: In a tunneling environment, many different internets may include AppleTalk networks that have the same network numbers. Therefore, each exterior router on an internet must associate a unique identifier (UI) with each network that it exports across the tunnel-that is, each network in its local internet that is not hidden. Generally, some type of global administration of UIs is necessary.

On a given tunnel, each exterior router on which network-number remapping is active must have a unique domain identifier (DI). An exterior router using AURP derives a network's UI by concatenating the exterior router's DI-which is unique on the tunnel-with the packet's network number or range-which is unique within the exterior router's domain. For more information about domain identifiers, see the section "Domain Identifiers" in Chapter 2.

On a tunneling port, an exterior router refers to AppleTalk network numbers and network ranges using UIs. Whenever an exterior router sends or receives AppleTalk data packets across the tunnel, it refers to any network numbers or ranges in the packets-for example, in a packet's DDP header-by their UIs. For example, when an exterior router sends an RI- Rsp, which provides a list of network ranges for its local internet to other exterior routers on the tunnel, it lists the UIs corresponding to those network ranges. When an exterior router receives RI-Rsp packets from other exterior routers on the tunnel, it interprets the data in each packet as a list of UIs.

Network-number remapping should be an optional component of any tunneling scheme. An administrator should be able to configure a tunneling port with or without specifying network-number remapping. When network-number remapping is inactive on all of the exterior routers on a tunnel, each AppleTalk network number and range associated with the exterior routers must be unique.

MAPPINGS: An exterior router uses the following process to map AppleTalk network numbers and ranges to UIs, and vice versa:

The exterior router logically maps network numbers in the exterior router's local internet to the corresponding UIs before sending a packet out the tunneling port, as shown in Figure 4-2. The UI consists of the source DI in the domain header and the network number from the packet. Therefore, the exterior router changes no data in the packet to perform this mapping.

The exterior router logically maps UIs corresponding to local networks in packets received through the tunneling port back to their local network numbers before forwarding the packets to the exterior router's local internet, as shown in Figure 4-2. The exterior router changes no data in the packet. This mapping is the inverse of the previous mapping.

The exterior router maps UIs corresponding to network numbers for remote networks-that is, networks connected to other exterior routers on the tunnel-that are in packets received through the tunneling port to network numbers in the remapping range configured for the local internet, as shown in Figure 4-2. An exterior router remaps network numbers from the following fields in this way:

the source network number field in the DDP header of an AppleTalk data packet

the NBP entity address field in an AppleTalk data packet

the routing data field in an AURP routing-information packet

The exterior router maps network numbers in the remapping range configured for the local internet back to the corresponding UIs before sending packets out the tunneling port, as shown in Figure 4-2. This type of remapping applies only to network numbers that reside in a destination network-number field of a DDP header in an AppleTalk data packet. This mapping is the inverse of the previous mapping.

<<Figure 4-2 Mappings between local and remote internets' network numbers and UIs>>

NOTE: Network-number remapping changes an AppleTalk data packet's DDP header and may also change its data. Thus, if a packet contains a DDP checksum, when the exterior router remaps network numbers contained in the packet, it must verify that the checksum is correct, then set the checksum to 0. If the checksum is incorrect, the exterior router should discard the packet.

An exterior router can perform network-number remapping either statically or dynamically. Static remapping reserves specific network numbers in the remapping range for mapping specific UIs. Dynamic remapping assigns network numbers in the remapping range to networks as they become known to an exterior router.

Static remapping is simpler to implement and provides a known mapping for use in network management. However, it may limit the number of UIs that an exterior router can import into its local internet.

Dynamic mapping requires a scheme for network number reuse, but may provide connectivity to a greater number of networks across a tunnel.

To avoid having the same UI refer to two different networks when remapping network numbers dynamically, an exterior router should reuse network numbers in its remapping range only when no other network numbers are available. If a network goes down, an exterior router should not immediately reassign the UI that referred to that network to another network that just came up on the internet.

An exterior router connected to more than one tunnel should function as though it were two exterior routers-each connected to one tunnel and both connected to one AppleTalk internet. Thus, such an exterior router must use remapped network numbers when sending routing information across a tunnel about networks that are accessible through another tunnel.

Network Numbers in Data

To remap network numbers properly, an exterior router must be aware of their presence within AppleTalk data packets. It is difficult to detect network numbers in data packets, because they could be anywhere within a data packet. For example, NBP includes network addresses as part of its data-in entity addresses. However, the data packets for very few protocols contain any network numbers. Some third-party protocols may contain network addresses in their data. Protocols that contain network addresses in their data may not function properly across remapping exterior routers.

Packets used for network management-such as RTMP Route Data Response (RDR) and Simple Network Management Protocol (SNMP) packets-contain network numbers in their data. For detailed information about handling network numbers in SNMP packets, see the section "Network Management" later in this chapter.

Problems With Loops

Network-number remapping introduces some problems on an internet when loops exist across a tunnel. If network-number remapping is active, two AppleTalk internets connected by a tunnel should not be interconnected in any other way. If a redundant path to an internet exists, a remapped network range can loop back through that path to the exterior router that originally remapped the network range. When this occurs, two different network ranges-the network range actually configured and the remapping of the configured range-refer to one network.

The remapped network range apparently refers to a new network in the exterior router's local internet. Such a network is referred to as a shadow network. The exterior router cannot determine that it has received a network range that it had previously remapped, because there is no apparent difference between a remapped network range and an actual network range. Thus, unless an administrator configures an exterior router with an explicit list of networks to export, the exterior router again remaps the network range, then exports the remapped network range, sending it around the loop. The network range is remapped repeatedly until the apparent distance to the network exceeds the hop-count limit. Exterior routers that implement network-number remapping should avoid establishing such infinite loops. For information about preventing such loops, see the section "Routing Loops" later in this chapter.

Redundant Paths

Under certain circumstances, it might be desirable to create a redundant path, which is a special type of loop. Redundant paths connect an internet to a tunnel through two or more exterior routers. If network-number remapping is active, all redundant exterior routers must use the same DI to represent the local internet-and must map UIs representing remote networks in incoming packets to the same local network numbers.

To allow redundant exterior routers to achieve such cooperation, a network administrator might configure all redundant exterior routers with the same DI and complete remapping information for all imported networks. Alternatively, a network administrator might configure one exterior router with this information and all redundant exterior routers could obtain the information from the configured exterior router. AURP does not currently support this functionality, but may do so in the future.

Tunnels With Partial Network-Number Remapping

When network-number remapping is active on a tunneling port, an exterior router maps network numbers in packets received through the tunnel into the remapping range for its local internet. Because a network administrator configures network-number remapping on individual exterior routers, network-number remapping may be configured on some exterior routers on a tunnel, but not on others-potentially causing network-numbering conflicts due to partial network-number remapping. Whenever possible, an administrator should configure network-number remapping either on all exterior routers on a tunnel or on none of them. Otherwise, network-numbering conflicts are likely to occur on some of the exterior routers-especially on large, interorganizational internets.

In addition to potential network-numbering conflicts, partial network-number remapping and the lack of loop detection between nonremapping exterior routers may cause shadow copies of networks connected to more than one nonremapping exterior router to appear in the routing tables on remapping exterior routers.

An exterior router on which network-number remapping is active performs loop detection. Therefore, when network-number remapping is active on all of the exterior routers on a tunnel, no loops can exist across the tunnel. However, exterior routers on which network-number remapping is not active do not perform loop detection. Thus, when network-number remapping is not active on some of the exterior routers on a tunnel, any loops that exist between nonremapping exterior routers are not detected.

In the example shown in Figure 4-3, shadow copies of all networks that are in the local internets of both exterior router B and exterior router C, on which network-number remapping is not active, appear in the routing table of exterior router A, on which network-number remapping is active.

<<Figure 4-3 A tunnel with partial network-number remapping>>

Clustering Remapped Networks

Because a remapping range is a range of sequential network numbers, an exterior router can represent multiple remapped networks as a single extended network within its local internet-that is, it can cluster remapped networks. Clustering greatly reduces the size of the routing tables that are maintained and sent by routers within an internet, as well as the amount of RTMP traffic on the internet. Clustering may also reduce the amount of NBP traffic on an internet.

For example, as shown in Figure 4-4, if networks in an internet have the numbers 1, 100, and 1000, and an exterior router connected to a different part of the internet receives these network numbers across the tunnel, that exterior router might remap the network numbers to 21, 22, and 23. When sending RTMP packets within its local internet, the remapping exterior router can represent the three networks as a single extended network with a network range from 21 to 23. The zones associated with the extended network include all of the zones associated with the three imported network numbers.

<<Figure 4-4 Clustering remapped network numbers>>

An exterior router determines which remapped network numbers it should cluster. For example, an exterior router might create one cluster for each other exterior router on the tunnel. However, an

exterior router can include no more than 255 zones in one cluster.

An exterior router that implements clustering must maintain the actual network range and zone list for each network in a cluster. The exterior router monitors all NBP FwdReq packets to be forwarded across the tunnel-including those it generates in response to BrRq packets. It examines the DDP destination network number in each FwdReq packet to determine the cluster to which it is addressed. The exterior router then generates one FwdReq packet for each clustered network for which the FwdReq packet contains a zone name, and sends that packet to the next internet router for the network. The DDP destination network number in such a FwdReq packet corresponds to the starting network number of a network's actual network range.

A disadvantage of clustering is that clusters are static. An exterior router cannot notify its local internet that a specific network or zone in a cluster has gone down. An exterior router's implementation of clustering could allow a network administrator to initiate reclustering-in which the exterior router notifies the internet that an entire cluster has gone down, then creates a new cluster that does not include the networks that have gone down. However, such reclustering would cause a temporary loss of connectivity to those networks in the cluster that are still accessible. Therefore, an exterior router should not automatically recluster network numbers.

REUSING NETWORK NUMBERS WITHIN A CLUSTER: Under certain conditions, an exterior router that implements clustering might reuse network numbers within a cluster. If a network went down, then came back up with the same zone list, an exterior router could map its network range into the same remapping range and include it in the same cluster. Otherwise, an exterior router should not reuse network numbers within a cluster, unless no other network numbers within the remapping range are available. In any case, an exterior router can reuse network numbers within a cluster only if a new network has a network range that fits in an unused range of network numbers within the cluster and a zone list that is a subset of the cluster's zone list.

The implementation of clustering in an exterior router is complex. See the Appendix, "Implementation Details," for some ways in which clustering could be implemented.

Zone-Name Management

To enhance zone-name management within an AppleTalk internet, AURP provides Get Domain Zone List and Get Zone Nets requests-which function similarly to the ZIP GetZoneList command and ZI-Req command, respectively. However, as when using RTMP and ZIP, if two networks in

an internet include zones that have the same zone name in their zone lists, exterior routers merge the zones into one zone-regardless of whether network-number remapping is active on one or more of the exterior routers.

Because AppleTalk data packets often contain zone names, AURP provides no means of remapping zone names. When importing or exporting zone names, an exterior router should not modify them in any way.

On a very large internet, zone names may become unmanageable. Therefore, an administrator should use domain-specific prefixes-such as Engineering or Sales-for zone names on such an internet. The use of a third-party hierarchical Chooser also might simplify zone-name management.

Hop-Count Reduction

Generally, an exterior router increases the hop count in the DDP header of an AppleTalk data packet by at least one when it forwards the packet across a tunnel. Once a packet traverses 15 routers-either local routers or exterior routers-its hop count exceeds the maximum. Thus, when an exterior router receives a packet through its tunneling port, it should examine that packet's DDP hop count before forwarding the packet. If the exterior router receives a packet with a hop count of 15 hops, it does not forward the packet to another router, but discards the packet.

When a tunnel or point-to-point link connects AppleTalk internets, the distance that a packet must traverse can easily exceed 15 hops. A network administrator might need full connectivity between two internets at a distance exceeding 15 hops. If the distance across an exterior router's local internet is already at or near the 15-hop limit, the exterior router must reduce the perceived distance that a packet must traverse to allow the packet to reach a destination at a distance that exceeds 15 hops. To overcome DDP's 15-hop limit, an exterior router reduces the hop count in the DDP header of an Apple data packet received through a tunnel before forwarding the packet into its local AppleTalk internet. An exterior router should reduce the hop count only by the number of hops necessary to allow the packet to reach its destination without exceeding the hop-count limit.

When an exterior router receives a packet through the tunnel, it examines the routing-table entry for that packet's destination network to determine the remaining distance to that network. If the distance already traversed by the packet-the packet's current hop count-plus the distance to the destination network is less than 15

hops, the exterior router simply forwards the packet. If adding the destination network's distance to the packet's current hop count causes the hop count to exceed 15 hops, the exterior router sets the hop count to the following value: 15 minus the distance in hops to the destination network. The exterior router then forwards the packet.

Using hop-count reduction, an exterior router must overcome the 15-hop limits imposed by both DDP and RTMP. To overcome RTMP's 15-hop limit, an exterior router should represent all networks accessible through the tunnel to routers in its local internet as one hop away when hop-count reduction is active on a tunneling port. This allows routers to maintain and send routing information about networks beyond the 15-hop limit and achieve full connectivity.

Constraints on Hop-Count Reduction

An interdomain loop exists when a redundant path connects two parts of an internet that are connected through two exterior routers on a tunnel. The proper operation of hop-count reduction requires that no interdomain loops exist across a tunnel. For detailed information about interdomain loops see the next section, "Routing Loops."

Because network-number remapping requires that no interdomain loops exist on the internet, an exterior router can perform hop-count reduction whenever network-number remapping is active, without any risk of a packet being forwarded in an infinite routing loop. Generally, an exterior router should not perform loop detection when network-number remapping is inactive.

Routing Loops

A routing loop exists when more than one path connects two exterior routers-both the path through the tunnel and a path through the exterior routers' local internets. When network-number remapping is not active on an exterior router, a routing loop can provide an alternative path to a network. However, when network-number remapping or hop-count reduction is active on an exterior router, all exterior routers must avoid establishing loops across the tunnel. Otherwise, if a routing loop went undetected, multiple routing-table entries that referred to the same actual AppleTalk networks using different remapping ranges might fill the routing tables of all of the exterior routers on a tunnel.

First-Order Loops

In a first-order loop, a pair of exterior routers that are performing network-number remapping across a tunnel are also connected through

another path, on which there are no remapping exterior routers. In Figure 4-5, exterior routers A and B are remapping network numbers across an AppleTalk tunnel, and exterior router C-which is not remapping network numbers-creates a first-order routing loop. Exterior router A's network range, 1 through 4, loops back to it through the tunnel and may be remapped again.

<<Figure 4-5 A first-order loop>>

Second-Order Loops

In a second-order loop, one or more additional pairs of remapping exterior routers are in the loop. In Figure 4-6, exterior routers A and B are remapping network numbers across the AppleTalk tunnel that connects them, and another pair of exterior routers, C1 and C2-which are also performing remapping across the tunnel that connects them-creates a second-order routing loop. Exterior router A's network range, 1 through 4, is remapped by exterior router C2 to the network range 101 through 104, then loops back to exterior router A through the tunnel.

<<Figure 4-6 A second-order loop>>

Self-Caused and Externally Caused Loops

Routing loops can be either self-caused or externally caused. A self-caused loop results when the detecting exterior router itself comes on line. An externally caused loop results when another router comes on line somewhere on the internet, after the detecting router has been running for some time.

Loop-Detection Process

The following sections describe the phases of the minimal loop-detection process that an exterior router must employ when either network-number remapping or hop-count reduction is active. An exterior router can implement an enhanced loop-detection scheme.

LOOP-INDICATIVE ROUTING INFORMATION: A remapping exterior router should always examine routing information received through a tunnel for indications that a routing loop may exist. Loop-indicative routing information appears to refer to networks across the tunnel. However, it may actually refer to networks in the exterior router's own local internet if the networks' routing information has looped back through the tunnel.

In the following definition of loop-indicative information,

the network range for the network connected to a given port of an exterior router is referred to as ns through ne

the zone list for that network is referred to as z1 through zn

The routing information that a remapping exterior router receives through a tunneling port is loop indicative if both of the following conditions are true for some port on the router:

The size of the network range in the routing information is ne-ns+1.

The zone list in the routing information consists precisely of z1 through zn.

Thus, the routing information could represent a remapping of the network range for a network connected directly to one of the exterior router's ports.

An exterior router most commonly receives loop-indicative information at startup when the process of bringing up the tunnel may create a self-caused loop. An exterior router may also receive loop-indicative information if another router connects two AppleTalk domains that are already connected through the tunnel and creates an externally caused loop.

If a remapping exterior router receives loop-indicative routing information through a tunnel, it should start a loop-investigation process. For information about the loop-investigation process, see the next section, "Loop-Investigation Process."

LOOP-INVESTIGATION PROCESS: To confirm or deny the existence of a suspected loop, an exterior router performs a loop-investigation process, in which it sends an AppleTalk data packet out the tunneling port, then observes whether that packet loops back through a port connected to its local internet. The exterior router sends the packet to the address corresponding to its own address on the network that it suspects may actually be a shadow copy of a network connected directly to one of its ports.

LOOP PROBE PACKET: A Loop Probe packet is an AppleTalk data packet that an exterior router sends out a tunneling port to confirm or deny the existence of a loop. It is a new type of RTMP packet and has the function code 4. Figure 4-7 shows the format of a Loop Probe packet.

<<Figure 4-7 Loop Probe packet format>>

The source node ID and source network number in a Loop Probe packet

should be those of the port for which the exterior router received loop-indicative information. An exterior router can send a Loop Probe packet through any socket.

A Loop Probe packet's destination network number is the network number to which that port's network number would be remapped if the loop-indicative information were actually a shadow copy of that port's routing information. Refer to the port's actual network number as nu(ns<=nu<=ne). If the network range in the loop-indicative information were rs through re, the packet's destination network number would be rs+nu-ns.

A Loop Probe packet's destination node ID is that of the exterior router on the port for which the exterior router received loop-indicative information. The packet's destination socket is socket 1-the RTMP socket.

A Loop Probe packet's data field always begins with a long word that has the value 0. The remainder of the data field should contain information that the exterior router that sends the packet can use to identify that packet if it receives the packet through its local internet. An exterior router might receive a Loop Probe packet sent by another exterior router if a loop did not actually exist and the other exterior router sent a Loop Probe packet to a random node on the internet rather than to itself. The node receiving the Loop Probe packet might be an exterior router that also sent a Loop Probe packet. To prevent an exterior router that receives such a Loop Probe packet from falsely concluding that a loop exists, the exterior router sending the packet must insert sufficient data in that packet's data field to allow it to recognize the packet as the one it sent.

An exterior router initiating a loop-investigation process should forward a Loop Probe packet through the tunnel to the next internet router for the packet's destination network-just as it would any other AppleTalk data packet. This next internet router should always be the exterior router that sent the loop-indicative information.

A remapping exterior router forwarding a Loop Probe packet into its local internet must process that packet differently from other AppleTalk data packets in one way. If the exterior router's remapping database does not include the source network number in the packet's DDP header, the exterior router should forward the packet without remapping the source network number. At startup, remapping information is generally unavailable. However, the absence of remapping information should not affect the loop-detection process.

If a loop exists, the exterior router that originally sent the Loop

Probe packet receives that packet through its local internet. The data in the packet remains unchanged. The exterior router can use that data to confirm the existence of a loop on the internet.

If a Loop Probe packet returns to the exterior router through the tunnel out which it was sent, a loop exists between two other exterior routers on the tunnel, but does not involve the exterior router that sent the packet. The sending router need take no action.

An exterior router should send a Loop Probe packet at least four times. The retransmission timeout should be no less than two seconds. Once the exterior router has retransmitted a Loop Probe packet four times and that packet has not returned to the exterior router through its local internet, the exterior router determines that no loop exists.

If the exterior router receives a Loop Probe packet containing the correct data field through its local internet, this confirms the existence of a loop. The exterior router should deactivate the tunneling port, log an error, and set the state of all routing-table entries for exterior routers connected to that tunnel to BAD.

NOTE: The exterior router need not deactivate a tunneling port on which it detects a loop. However, the exterior router must disconnect with the exterior router that sent the loop-indicative information. However, disconnecting from only that exterior router might inadvertently result in a partially connected tunnel or in a lack of connectivity through the tunnel that would be difficult to detect.

LIMITATIONS OF LOOP DETECTION: This loop-detection process becomes ineffective if, at some point in the loop, another exterior router

- hides networks connected directly to the ports of the exterior router that sent the Loop Probe packet
- clusters the network ranges of networks connected directly to the exterior router’s ports
- is not remapping network numbers-resulting in partial network-number remapping

In such cases, the exterior router that initiated the loop-detection process may never receive loop-indicative information, even though a loop exists.

Using Alternative Paths

AURP provides two mechanisms that allow a network administrator to

configure a port on an exterior router to forward packets over an alternative path to a network only when the primary path to that network is unavailable:

- hop-count weighting
- backup paths

By configuring hop-count weighting on a port or configuring a port as a backup path, an administrator can reduce the amount of traffic on a slow point-to-point link or tunnel. These mechanisms are also available on links using RTMP.

Hop-Count Weighting

A network administrator can configure hop-count weighting on a port to increase the routing distance through a port by counting a link to another exterior router as more than one hop. Increasing the routing distance through a port may cause traffic to traverse an alternative path. The routers on an internet forward packets over an alternative path to a network if

- an alternative path is available
- the perceived distance to that network is shorter over the alternative path

However, a network administrator should not set the hop-count weight for a link so high that distances between networks across that link exceed the limit of 15 hops. Otherwise, if the link on which hop-count weighting was active were the only available path, the exterior router would be unable to provide full connectivity to all networks on the internet.

To implement hop-count weighting, an exterior router should make the following changes to RTMP and the DDP routing process:

- When an exterior router uses RTMP or AURP to broadcast the networks that are accessible through a link on which hop-count weighting is active, the distance attributed to each network should equal its actual distance plus the hop-count weight specified.
- Before an exterior router forwards a DDP data packet to a network across that link, it should add the specified hop-count weight to the value in the hop-count field of the packet’s DDP header.

Backup Paths

A network administrator can configure a port on an exterior router as a backup path. The routers on an internet forward AppleTalk data packets across a backup path only when an exterior router on which a port is configured as a backup path determines that no other path to a specific network or networks is available.

Regardless of the distance that routing packets must traverse across a primary path to a network, routers on the internet use the primary path as long as it remains available. When the exterior router on which a port is configured as a backup path determines that the primary path to a network is no longer available and that network is accessible across the backup path, the exterior router broadcasts routing information about networks accessible across the backup path to its local internet.

NOTE: An exterior router at each end of the backup path maintains a complete routing table for the entire internet, and sends AURP or RTMP routing packets across the backup path, regardless of whether the backup path is in use.

If an exterior router is currently providing access to a network through a backup path and the primary path to that network again becomes available, the exterior router starts broadcasting routing information that indicates the primary path to the network, rather than the backup path. The routers on the exterior router's local internet can again use the primary path to that network.

PROBLEMS REACTIVATING THE PRIMARY PATH: When an exterior router is providing access to a network through a backup path and the primary path to that network again becomes available, it is possible that the exterior router may not become aware that the primary path is available. This can occur when other routers in the exterior router's local internet use the backup path, rather than a newly available primary path, because the backup path traverses a shorter distance. The other routers have no way of knowing that an active path is a backup path. They do not notify the exterior router connected to the shorter backup path about the primary path's availability.

Once the primary path becomes unavailable and routers on the internet use the backup path, reconfiguring the exterior router so it will again use the primary path may be necessary.

Network Management

A Simple Network Management Protocol (SNMP) Management Information

Base (MIB) allows the remote management of tunneling, routing-information propagation, and the representation of wide area routing information. Refer to the "IETF Draft: Macintosh System MIB" on E.T.O. for detailed information about the structure and content of AURP's many remotely manageable parameters.

Network-Number Remapping and Network Management

The packets of network-management protocols-regardless of whether SNMP forms their basis-often contain information about specific AppleTalk network numbers. An exterior router cannot remap network numbers in data. Therefore, when querying devices across a tunnel, network-management protocols always return network numbers that have not been remapped. However, a remote network-management station using SNMP could use the AURP MIB to query a remapping exterior router to obtain remapped network numbers from the exterior router's remapping database.

Network Hiding and Network Management

Even though an exterior router is hiding a network from a particular port, that network's routing information should be available to a network-management station across that port. Network hiding should not affect network management. Thus, an exterior router should still return routing information for hidden networks in responses to network-management queries. A network-management station using SNMP could use the AURP MIB to query an exterior router to obtain information about hidden networks.

Unaffected Network-Management Packets

Network-management packets that network-number remapping and network hiding should not affect include:

- SNMP requests received through an AURP port
- SNMP responses sent through an AURP port
- RTMP responses sent through an AURP port
- Route Data responses sent through an AURP port
- ZIP queries received through an AURP port
- ZIP requests received through an AURP port
- ZIP replies sent through an AURP port

APPENDIX: IMPLEMENTATION DETAILS

This appendix provides information that may assist you in implementing AURP. It does not specify protocol requirements.

Developers implementing AURP routers may want to purchase the Apple Internet Router, a product of Apple Computer. The Apple Internet Router provides many additional examples of how you might implement the various features of AURP.

State Diagrams

Figure A-1 shows the state diagram for the AURP data receiver.

<<Figure A-1 AURP data receiver state diagram>>

Figure A-2 shows the state diagram for the AURP data sender.

<<Figure A-2 AURP data sender state diagram>>

AURP Table Overflow

It is possible for an AURP data receiver to have insufficient storage capacity to maintain all of the routing information sent to it by a peer data sender. Because the data sender does not retransmit routing information, the data receiver should set a flag indicating that a table-overflow condition exists. If additional storage later becomes available, the data receiver should try to obtain the missing information. If zone information is lost, the data receiver can obtain complete zone information by sending the appropriate ZI-Req packets. If network information is lost, the data receiver should send an RI-Req to obtain the complete routing table.

A Scheme for Updates Following Initial Information Exchange

As described in the section "Sending Updates Following the Initial Exchange of Routing Information" in Chapter 3, an exterior router must present complete and accurate routing information to all exterior routers, even if a new connection is established with that exterior router when the exterior router has update events pending-that is, update events not yet sent in RI-Upd packets. This section details one scheme for presenting routing information to both new and old connections correctly, even if multiple update events occur for a given network in an update period during which the exterior router establishes new connections. More complex schemes could provide more up-to-date information, at the cost of greater implementational complexity.

Assume that an exterior router has a number of AURP connections established with other routers and that a series of update events for a given network occur in the exterior router's local internet. Once these events have occurred, but before the update interval expires-that is, before the exterior router sends RI-Upd packets over its connections-the exterior router establishes a new AURP connection with another exterior router and receives an RI-Req packet from that exterior router. This section describes the information about the network that the RI-Rsp packet should contain. It also describes the update event that the exterior router should send in the next RI-Upd packet, assuming that it receives no additional update events for the network.

Two scenarios are possible. In the first scenario, a network for which the exterior router is not exporting information at the beginning of an update interval either comes up in the exterior router's local internet, or a new path to the network that is shorter than the path through the tunnel comes up in the exterior router's local internet. In either case, the RI-Rsp packet should not include the new network.

By not including the new network in the RI-Rsp, the implementation can simply continue to follow the state diagram provided in the section "Sending Routing Information Update Packets" in Chapter 3. If only an NDC event or no additional update event occurs for the network, the next RI-Upd packet that the exterior router sends on both old and new connections should contain an NA event for the network. If an NRC or ND event occurs for the network, the exterior router should not include an event tuple for the network in the RI-Upd. This sequence matches the state diagram precisely. If the RI-Rsp did contain information about the network, new connections would require a different state diagram.

In the second scenario, the exterior router initially exports information for a network, then an update event occurs for that network. In all cases, the RI-Rsp packet should contain up-to-date information about the network from the exterior router's central routing table, and the next RI-Upd packet should contain the specific event that the state table indicates for that network. For example, if an ND or NRC event occurs for the network, the network should not be included in the RI-Rsp, while if an NDC event occurs, it should be included in the RI-Rsp.

This scheme may result in some exterior routers receiving unexpected update events, which they must process as specified in the section "Processing Inconsistent Update Events" in Chapter 3. For example, another exterior router with which the exterior router establishes a new connection might receive an ND or NRC event for a network of

which it was unaware. The receiving exterior router would ignore the event.

In an alternative way of evaluating and possibly implementing this scheme, the information for a given network that is sent in the initial RI-Rsp packet depends on the particular update event that is pending for that network when the exterior router sends the RI-Rsp. Specifically, an exterior router should include a network for which it has an update event pending in the RI-Rsp packet only if the pending update event is an NDC. Otherwise, the exterior router should not include the network in the RI-Rsp. Following this RI-Rsp, the exterior router sends RI-Upd packets as usual, which include other pending events, as necessary.

Implementation Effort for Different Components of AURP

AURP contains various enhancements to AppleTalk routing. The only components of AURP that are required are those specified in Chapter 3. The required components of AURP provide the functionality needed to replace RTMP and ZIP, completely and compatibly, on tunnels and point-to-point links, without losing any functionality and with greatly reduced routing traffic. Optional features of AURP provide functionality beyond that of RTMP and ZIP. This functionality is especially useful in a wide area network environment.

The chart shown in Figure A-3 provides rough estimates of the percentage of development time needed to implement, debug, and test the various components of a complete AURP implementation. It can provide developers with some idea of the implementational complexity of these components and help developers make tradeoffs between features and development time.

<<Figure A-3 Implementation effort for AURP>>

Creating Free-Trade Zones

A useful feature of AURP is that it allows a network administrator to create free-trade zones. A free-trade zone is a part of an internet that is accessible by two other parts of the internet, neither of which can access the other. An administrator might create a free-trade zone to provide some form of interchange between two organizations that otherwise want to keep their internets isolated from each other, or between two organizations that otherwise do not have physical connectivity with one another.

AURP allows the creation of free-trade zones in two ways. In one method, described in the section "Fully Connected and Partially Connected Tunnels" in Chapter 2, an administrator intentionally

creates a partially connected tunnel. The administrator configures the exterior router to connect with two exterior routers between which a free-trade zone is to be established, but does not configure those exterior routers to connect with one another.

The second method of using AURP to create a free-trade zone involves the use of network hiding. An administrator can configure a single router to create a free-trade zone. No AURP tunnel need exist. As shown in Figure A-4, three ports are configured on a router. One port connects to the free-trade zone, while the other two ports connect to the parts of the internets that are otherwise isolated from one another.

<<Figure A-4 Creating free-trade zones>>

On the port connected to the free-trade zone, the administrator does not configure the router to hide any networks. The exterior router exports all networks from both organizations to the free-trade zone. On each port connected to an organization's internet, the administrator configures the router to export only the networks from the free-trade zone. The exterior router hides all the networks from the other organization's internet. In this way, each organization has access to the networks in the free-trade zone, and vice versa, but not to the networks in the other organization's internet.

Implementation Details for Clustering

The data structures that an exterior router uses to maintain information about clustering are key to the implementation of clustering. An exterior router should

- maintain mappings between the actual domain identifier and network range; the remapped network range; and the associated cluster
- maintain zone lists for each actual network and for the cluster as a whole
- use data structures that allow parts of the information to be marked for deletion, while maintaining that information for possible later reuse-for example, if a network goes down, then comes back up
- use data structures that are bidirectional-supporting both the conversion of a single FwdReq into multiple FwdReq packets and the manipulation of individual networks within the cluster

An exterior router can cluster any network numbers that is has remapped into an available range of contiguous network numbers. From both an implementation and a management point of view, it is

generally best for an exterior router to cluster all network numbers that it receives from a particular exterior router at a given time. For example, it may be desirable to cluster all of the network numbers included in the initial information exchange with a particular exterior router, then later, to cluster all of the network numbers received in NA events in a given RI- Upd packet.

Maintaining compatibility with AppleTalk Phase 2 complicates the implementation of clustering. An exterior router can include a maximum of 255 zones in a cluster. This limit may prevent the exterior router from clustering all of the network numbers that it receives at one time. When an exterior router receives a list of networks from another exterior router, it does not know how many different zone names the networks use. The exterior router does not have this information until it receives the associated ZI-Rsp packets. Therefore, an exterior router should not build a cluster until it has received a complete zone list for the network numbers being clustered. Once the exterior router has complete zone information for the network numbers, it can cluster the maximum number of network numbers allowed by the 255 zone limit.

AURP does not specify the method by which an exterior router, when forming a cluster, should determine the hop count for that cluster- that is, the apparent distance in hops to the single extended network that represents the cluster. Possible implementation options include

always setting the hop count to a constant value

setting the hop count to the minimum, average, or maximum of the hop counts for the networks within the cluster

In a large internet, setting the hop count for a cluster too high may make the networks in that cluster unreachable from some networks in the local internet of the exterior router that is clustering the network numbers.

Modified RTMP Algorithms for a Backup Path

In the following RTMP maintenance algorithms defined in Inside AppleTalk, the backup path is an RTMP link. These algorithms can be adapted to AURP according to the architectural model described in the section "AURP Architectural Model" in Chapter 3. Proposed modifications to these algorithms appear in boldface Courier font.

On Receiving an RTMP Data Packet Through a Port

IF P is connected to an AppleTalk network AND P’s network number range = 0

```
   THEN BEGIN
        P's network number range := packet's sender network
             number range;
        IF there is an entry for this network number range
        THEN delete it;
        Create a new entry for this network number range with
             Entry's network number range := packet's sender
                  network number range;
             Entry's distance := 0;
             Entry's next IR := 0;
             Entry's status := Good;
             Entry's port := P;
        END;
   FOR each routing tuple in the RTMP Data packet DO
        IF there is a table entry corresponding to the tuple's
             network number range
             THEN Update-the-Entry
        ELSE IF there is a table entry overlapping with the
             tuple's network number range
             THEN ignore the tuple
        ELSE IF P is not a backup path
             THEN Create-New-Entry
        ELSE     Create-New-Tentative Entry;

   Update-the-Entry

   IF (Entry's port is not a backup port AND P is a
        backup port)
   THEN Return; {Ignore tuple}
   IF (Entry's state = Bad) AND (tuple distance <15)
   THEN Replace-Entry
   ELSE
        IF Entry's distance >= (tuple distance +1) AND (tuple
             distance <15)
             OR  (Entry's port is a backup port and P is not a
                  backup port)
        THEN Replace-Entry
        ELSE IF Entry's next IR = RTMP Data packet's sender node
             address AND Entry's port = P
        THEN IF tuple distance <> 31 THEN BEGIN
             Entry's distance := tuple distance + 1;
             IF Entry's distance < 16
             THEN Entry's state := Good
             ELSE Delete the entry
        END
        Else Entry's state := Bad;
```

An exterior router uses the Create-New-Tentative-Entry algorithm when it discovers a previously unknown network across a backup path. An exterior router should not add an entry to the routing table being broadcast to its local internet until it determines definitely that no alternative path to a network is available. While waiting for another path to a network to become available, the exterior router temporarily stores the routing-table entry in a tentative routing table, as defined by the following algorithm:

```
Create-New-Tentative-Entry

IF tentative entry for tuple's network number range does not
     already exist
     THEN BEGIN
          Tentative entry's network number range =
               tuple's network number range;
          Tentative entry's distance := tuple's distance;
          Tentative entry's next IR = packet's node address;
          Tentative entry's port := P;
          Start a TBD-minute timer for this entry;
     END;
WHEN timer for this entry expires
     IF there is a table entry corresponding to or
          overlapping with the tentative entry's network
          number range
          THEN ignore the entry
     ELSE Create-New-Entry; {using data from the tentative
          entry}
     Delete tentative entry;
```

Security Considerations

This memo discusses a weak form of security called network hiding or device hiding. More general concerns about security are not addressed.

Author's Address

Alan B. Oppenheimer
Apple Computer, M/S 35-K
20525 Mariani Avenue
Cupertino, California 95014

Phone: 408-974-4744
EMail: Oppenheime1@applelink.apple.com

Note: The author would like to acknowledge the contribution of Pabini Gabriel-Petit here at Apple, who translated the engineering specification into human-readable form.