Lucent Technologies
Bell Labs Innovations

Stinger® T1000 Module

Routing and Tunneling Supplement

For software version 9.4-185.1
December 2002

Notice

Every effort was made to ensure that the information in this document was complete and accurate at the time of printing, but information is subject to change.

European Community (EC) RTTE compliance

CE Hereby, Lucent Technologies, declares that the equipment documented in this publication is in compliance with the essential requirements and other relevant provisions of the Radio and Telecommunications Technical Equipment (RTTE) Directive 1999/5/EC.

To view the official *Declaration of Conformity* certificate for this equipment, according to EN 45014, access the Lucent INS online documentation library at http://www.lucentdocs.com/ins.

Safety, compliance, and warranty information

Before handling any Lucent Access Networks hardware product, read the *Edge Access and Broadband Access Safety and Compliance Guide* included in your product package. See that guide also to determine how products comply with the electromagnetic interference (EMI) and network compatibility requirements of your country. See the warranty card included in your product package for the limited warranty that Lucent Technologies provides for its products.

Security statement

In rare instances, unauthorized individuals make connections to the telecommunications network through the use of access features.

Trademarks

Lucent, the Lucent logo, and all Lucent brand and product names are trademarks or registered trademarks of Lucent Technologies Inc. Other brand and product names are trademarks of their respective holders.

Ordering information

You can order the most up-to-date product information and computer-based training online at http://www.lucentdocs.com/bookstore.

Feedback

Lucent Technologies appreciates customer comments about this manual. Please send them to techcomm@lucent.com.

Lucent Technologies

Customer Service

Product and service information, and software upgrades, are available 24 hours a day. Technical assistance options accommodate varying levels of urgency.

Finding information and software

To obtain software upgrades, release notes, and addenda for this product, log in to Lucent OnLine Customer Support at `http://www.lucent.com/support`.

Lucent OnLine Customer Support also provides technical information, product information, and descriptions of available services. The center is open 24 hours a day, seven days a week. Log in and select a service.

Obtaining technical assistance

Lucent OnLine Customer Support at `http://www.lucent.com/support` provides easy access to technical support. You can obtain technical assistance through email or the Internet, or by telephone. If you need assistance, make sure that you have the following information available:

- Active service or maintenance contract number, entitlement ID, or site ID
- Product name, model, and serial number
- Software version or release number
- Software and hardware options
- If supplied by your carrier, service profile identifiers (SPIDs) associated with your line
- Your local telephone company's switch type and operating mode, such as AT&T 5ESS Custom or Northern Telecom National ISDN-1
- Whether you are routing or bridging with your Lucent product
- Type of computer you are using
- Description of the problem

Obtaining assistance through email or the Internet

If your services agreement allows, you can communicate directly with a technical engineer through Email Technical Support or a Live Chat. Select one of these sites when you log in to `http://www.lucent.com/support`.

Calling the technical assistance center (TAC)

If you cannot find an answer through the tools and information of Lucent OnLine Customer Support or if you have a very urgent need, contact TAC. Access Lucent OnLine Customer Support at `http://www.lucent.com/support` and click `Contact Us` for a list of telephone numbers inside and outside the United States.

Alternatively, call 1-866-LUCENT8 (1-866-582-3688) from any location in North America for a menu of Lucent services. Or call +1 510-769-6001 for an operator. If you do not have an active services agreement or contract, you will be charged for time and materials.

Contents

Figures

Tables

About This Supplement

The Stinger T1000 module includes two physical components:

- The T1000 line interface module (LIM) (STGR-LIM-T1000) is an integrated IP router module for Stinger units, with session termination and aggregation functionality.
- The (optional) T1000 line element module (LEM) (STGRFS-LEM-2 or STGRLS-LEM-2) provides two 10/100BaseT Ethernet interfaces for the T1000 LIM that can be used for direct IP over Ethernet egress from a Stinger unit.

Note The Stinger T1000 module is currently available only for Stinger FS and Stinger LS platforms. It is not currently available for the Stinger RT.

What is in this supplement

This supplement describes how to configure the T1000 IP routing and optional virtual private network (VPN) capabilities. The following manuals are also important sources of T1000 configuration information:

- *Stinger T1000 Module Configuration Guide*—to configure the module's interfaces, termination of Point-to-Point Protocol (PPP) sessions over Ethernet (PPPoE) or ATM (PPPoA), and aggregation of ATM-1483 PVCs
- *Stinger ATM Configuration Guide*—to configure ATM connections

For information about configuring frame relay, to support T1000 termination of PPP sessions over frame relay, see the *Stinger 32-Port IDSL Line Interface Module (LIM) Configuration Guide*.

For information about installing modules, see the Stinger unit's *Getting Started Guide*.

Warning Before installing your Stinger unit, be sure to read the safety instructions in the *Edge Access Safety and Compliance Guide*. For information specific to your Stinger unit, see the "Safety-Related Physical, Environmental, and Electrical Information" appendix in the *Getting Started Guide* for your Stinger unit.

Documentation conventions

Following are the special characters and typographical conventions that might be used in this manual:

Convention	Meaning
`Monospace text`	Represents text that appears on your computer's screen, or that could appear on your computer's screen.

Convention	Meaning
Boldface monospace text	Represents characters that you enter exactly as shown (unless the characters are also in ***italics***—see *Italics*, below). If you could enter the characters but are not specifically instructed to, they do not appear in boldface.
Italics	Represent variable information. Do not enter the words themselves in the command. Enter the information they represent. In ordinary text, italics are used for titles of publications, for some terms that would otherwise be in quotation marks, and to show emphasis.
[ ]	Square brackets indicate an optional argument you might add to a command. To include such an argument, type only the information inside the brackets. Do not type the brackets unless they appear in boldface.
|	Separates command choices that are mutually exclusive.
>	Points to the next level in the path to a parameter or menu item. The item that follows the angle bracket is one of the options that appear when you select the item that precedes the angle bracket.
Key1-Key2	Represents a combination keystroke. To enter a combination keystroke, press the first key and hold it down while you press one or more other keys. Release all the keys at the same time. (For example, Ctrl-H means hold down the Control key and press the H key.)
Press Enter	Means press the Enter, or Return, key or its equivalent on your computer.
Note	Introduces important additional information.
 **Caution:**	Warns that a failure to follow the recommended procedure could result in loss of data or damage to equipment.
 **Warning:**	Warns that a failure to take appropriate safety precautions could result in physical injury.
 **Warning:**	Warns of danger of electric shock.

Stinger documentation set

The Stinger documentation set consists of the following manuals, which can be found at `http://www.lucent.com/support` and `http://www.lucentdocs.com/ins`.

- **Read me first:**
  - *Edge Access and Broadband Access Safety and Compliance Guide*. Contains important safety instructions and country-specific information that you must read before installing a Stinger unit.

  - *TAOS Command-Line Interface Guide*. Introduces the TAOS command-line environment and shows you how to use the command-line interface effectively. This guide describes keyboard shortcuts and introduces commands, security levels, profile structure, and parameter types.

- **Installation and basic configuration:**
  - *Getting Started Guide* for your Stinger platform. Shows how to install your Stinger chassis and hardware. This guide also shows you how to use the command-line interface to configure and verify IP access and basic access security on the unit, and how to configure Stinger control module redundancy on units that support it.
  - Module guides. For each Stinger line interface module (LIM), trunk module, or other type of module, an individual guide describes the module's features and provides instructions for configuring the module and verifying its status.
- **Configuration:**
  - *Stinger ATM Configuration Guide*. Describes how to integrate the Stinger into the ATM and Digital Subscriber Line (DSL) access infrastructure. The guide explains how to configure PVCs, and shows how to use standard ATM features such as quality of service (QoS), connection admission control (CAC), and subtending.
  - *Stinger IP2000 Configuration Guide*. For Stinger systems with the T1000 control module, this guide describes how to integrate the system into the IP infrastructure. Topics include IP-routed switch-through ATM PVCs and RFC 1483 PVCs that terminate on the T1000, IEEE 802.1Q VLAN, and forwarding multicast video transmissions on DSL interfaces.
  - *Stinger Private Network-to-Network Interface (PNNI) Supplement*. For the optional PNNI software, this guide provides quick-start instructions for configuring PNNI and soft PVCs (SPVCs), and describes the related profiles and commands.
  - *Stinger SNMP Management of the ATM Stack Supplement.* Describes SNMP management of ATM ports, interfaces, and connections on a Stinger unit to provide guidelines for configuring and managing ATM circuits through any SNMP management utility.
  - *Stinger T1000 Module Routing and Tunneling Supplement.* For the optional T1000 module, this guide describes how to configure the Layer 3 routing and virtual private network (VPN) capabilities.
- **RADIUS:** *TAOS RADIUS Guide and Reference*. Describes how to set up a unit to use the Remote Authentication Dial-In User Service (RADIUS) server and contains a complete reference to RADIUS attributes.
- **Administration and troubleshooting:** *Stinger Administration Guide*. Describes how to administer the Stinger unit and manage its operations. Each chapter focuses on a particular aspect of Stinger administration and operations. The chapters describe tools for system management, network management, and Simple Network Management Protocol (SNMP) management.
- **Reference**:
  - *Stinger Reference*. An alphabetic reference to Stinger profiles, parameters, and commands.
  - *TAOS Glossary.* Defines terms used in documentation for Stinger units.

Related documents

The following requests for comments (RFCs) are relevant to the TAOS software described in this supplement:

- RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support.*
- RFC 2661, *Layer Two Tunneling Protocol "L2TP."*
- RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*
- RFC 2328, *OSPV Version 2*
- RFC 2107, *Ascend Tunnel Management Protocol - ATMP.*
- RFC 2030, *Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI.*
- RFC 1858, *Security Considerations for IP Fragment Filtering.*
- RFC 1812, *Requirements for IP Version 4 Routers.*
- RFC 1787, *Routing in a Multi-provider Internet.*
- RFC 1701, *Generic Routing Encapsulation (GRE).*
- RFC 1700, *Assigned Numbers.*
- RFC 1582, *Extensions to RIP to Support Demand Circuits.*
- RFC 1519, *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy.*
- RFC 1458, *Requirements for Multicast Protocols.*
- RFC 1433, *Directed ARP.*
- RFC 1281, *Guidelines for the Secure Operation of the Internet.*
- RFC 1256, *ICMP Router Discovery Messages.*
- RFC 1112, *Host Extensions for IP Multicasting.*

IP Routing

1

Introduction to IP routing

When you reset the system, the system creates an IP routing table containing all the routes it has learned, including the following:

- Routes for local active IP interfaces (configured `ip-interface` profiles)
- Routes for active WAN IP sessions
- Routes for inactive WAN IP sessions (configured `connection` profiles)
- Routes defined in `ip-route` profiles or RADIUS route profiles

If dynamic routing protocols are enabled on one or more interfaces, the system adds routes it learns from routing-update packets. In addition, the system is continuously updating its routing table by adding routes for links that become active and removing routes for inactive sessions. If a nailed connection goes down, the system removes the route from its routing table.

Routes and interfaces

An IP route specifies a destination address, a gateway to the network, and an interface that leads to the gateway. It can also specify metrics and other values associated with the route.

A route defined in a profile is a *static route*. A *dynamic route* is learned from Routing Information Protocol (RIP) updates sent by other routers. Dynamic updates provide access to many more routes than those actually configured in the system, and are

updated automatically as routes change. However, dynamic updates cause additional routing overhead, so they are disabled by default.

An *interface* is a point of ingress to or egress from the system. For example, a local interface is an Ethernet port and a WAN interface is a nailed or switched connection. An *IP interface* is the logical IP address that enables IP data to be sent and received.

Displaying the routing table

To view the routing table, use the `netstat` command. For example:

```
admin> netstat -r
Destination       Gateway    IF       Flg   Pref Met     Use       Age
0.0.0.0/0         10.32.8.1  ie0      SGP    60    1    31460      1986
0.0.0.0/0         20.1.1.8   wan9     *SGP   60    8        0         0
10.4.5.0/24       10.4.5.6   wan12    SG     120   7        0   1978086
10.4.5.6/32       10.4.5.6   wan12    S      120   7        1   1978086
10.56.1.0/24      -          ie0-1    C        0   0        0   4504466
10.56.1.1/32      -          local    CP       0   0        0   4504466
127.0.0.0/8       -          bh0      CP       0   0        0   450446
127.0.0.1/32      -          local    CP       0   0        0   4504466
127.0.0.2/32      -          rj0      CP       0   0        0   4504466
10.32.8.0/24      -          ie0      C        0   0     7820   4504466
10.32.8.0/24      10.32.8.21 wan11    *SG    120   7        0   1978086
10.32.8.21/32     10.32.8.21 wan11    S      120   7        1   1978086
10.32.8.25/32      -         local    CP       0   0    47039   4504466
224.0.0.0/4        -         mcast    CP       0   0        0   4504466
224.0.0.1/32       -         local    CP       0   0        0   4504466
224.0.0.2/32       -         local    CP       0   0        0   4504466
224.0.0.5/32       -         local    CP       0   0     3158   4504466
224.0.0.6/32       -         local    CP       0   0        0   4504466
224.0.0.9/32       -         local    CP       0   0    14194   4504466
255.255.255.255/32 -         ie0      CP       0   0        0   4504466
```

For each route in the table, the Destination and Gateway fields show the destination address and the address of the next-hop router used to reach that destination. The zero destination address is the default route. If the system does not find a route for a packet's destination, it forwards the packet to the default route rather than dropping the packet. Note that the system uses the most specific route (having the longest prefix) that matches a given destination. Direct routes do not show a gateway address.

An asterisk (*) in the flags column indicates a hidden route, which is not included in routing updates sent by the system and is not used for forwarding packets. Hidden routes are used only for display purposes.

The IF field shows the name of the interface through which a packet addressed to the entry's destination will be sent. The route to the `mcast` interface name encapsulates the multicast forwarder for the entire class D address space. (For more information, see Chapter 2, "IGMP Multicast Forwarding.")

Routes to the local machine display the `local` interface name. Packets to the 224.0.0.1 and 224.0.0.2 interfaces can be multicast and received like normal multicast packets, but upon receiving such a packet, the router does not forward it to

another link layer device. Effectively, these packets have a maximum transmission unit (MTU) of 1.

Displaying the interface table

To display the interface table, use the `-i` option on the `netstat` command line:

```
admin> netstat -i
Name       MTU   Net/Dest        Address        Ipkts  Ierr Opkts    Oerr
ie0        1500 10.32.8.0/24  10.32.8.25     1018339    1   622450    1
ie0-1      1500 10.56.1.0/24  10.56.1.1            0    0        0    0
lo0        1500 127.0.0.1/32  127.0.0.1        26622    0    26622    0
rj0        1500 127.0.0.2/32  127.0.0.2            0    0        0    0
bh0        1500 127.0.0.3/32  127.0.0.3            1    0        1    0
wanabe     1500 127.0.0.3/32  127.0.0.3            0    0        0    0
local     65535 127.0.0.1/32  127.0.0.1       233371    0   233371    0
mcast     65535 224.0.0.0/4   224.0.0.0            0    0        0    0
tunnel8    1500 10.32.8.0/24  10.32.8.25           0    0        0    0
vr0_main   1500 10.32.8.25/32 10.32.8.25           0    0        0    0
sip0       65535 -              -                  0    0        0    0
wan11      1500 10.32.8.21    10.32.8.25           0    0        0    0
wan12      1500 10.4.5.6      10.32.8.25           0    0        0    0
wan13      1500 -              -                   0    0        0    0
wan14      1500 -              -                   0    0        0    0
ie1-15-1   1500 -              -                   0    0        0    0
ie1-15-2   1500 -              -                   0    0        0    0
ie1-15-3   1500 -              -                   0    0        0    0
ie1-15-4   1500 -              -                   0    0        0    0
ie1-15-1-1 1500 -              -                   0    0        0    0
ie1-15-1-2 1500 -              -                   0    0        0    0
ie1-15-1-3 1500 -              -                   0    0        0    0
```

The entries named `ie0` or `ie`*N-N-N*`[-`*N*`]` represent Ethernet interfaces. *N-N-N-N* represents the shelf number, slot number, item number, and logical-item number of the interface. When the logical-item number is zero (the physical interface), it does not appear in the interface name. The same sequence of numbers forms the address used to index the `ip-interface` profile. For example, the default profile for 1-4-1 is indexed as follows:

```
IP-INTERFACE { { 1 4 1 } 0 }
```

When the logical-item number is *not* zero, it does appear in the interface name. Again, the sequence of numbers is identical to the profile index. For example, suppose an `ip-interface` profile has the following index:

```
IP-INTERFACE { { 1 4 1 } 3 }
```

This profile has the following interface name:

```
ie1-4-1-3
```

The other names in the interface table have the following significance:

- The `lo0` (loopback) interface is the local loopback.
- The `rj0` (reject) and `bh0` (blackhole) interfaces are used in the pool-summary feature.

- The `wanabe` interface is an inactive RADIUS dial-out profile.
- The `local` interface is the local machine.
- The `mcast` interface is the multicast interface, which represents the multicast forwarder for the entire class D address space. For details, see Chapter 2, "IGMP Multicast Forwarding."
- The `tunnel` interface is a pseudo-interface that is used only when the system is configured as an ATMP router home agent. In that configuration, the system creates a route for each registered mobile client. Regardless of how many tunnels the home agent might terminate, there is always a single tunnel interface. (The number terminating the tunnel interface name is an internal number that can change from one software version to the next.)
- The `vr0_main` interface represents the router itself. For details, see Chapter 4, "Virtual Routing."
- The `sip0` interface is a soft IP interface. For details, see "Setting a system source IP address" on page 1-36.
- The numbered WAN (`wan`*N*) interfaces are WAN connections, which are entered in the interface table as they become active.

IP address syntax

TAOS uses dotted decimal format (not hexadecimal) for IP addresses. If no subnet mask is specified, the system uses a default mask based on the address class. Table 1-1 shows address classes and the number of network bits in the default mask for each class.

Table 1-1. IP address classes and number of network bits

Class	Address range	Default network bits
Class A	0.0.0.0–127.255.255.255	8
Class B	128.0.0.0–191.255.255.255	16
Class C	192.0.0.0–223.255.255.255	24

For example, a class C address, such as 198.5.248.40, has 24 network bits, leaving 8 bits for the host portion of the address. If no subnet mask is specified for a class C address, TAOS uses the default mask of 24 bits, as shown in Figure 1-1.

Figure 1-1. Default subnet mask for class C IP address

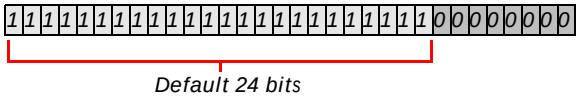

A subnet address includes a prefix length, which specifies the number of network bits in the address. For example, the following address specifies a 29-bit subnet:

```
ip-address = 198.5.248.40/29
```

In this address, 29 bits of the address are used to specify the network. The three remaining bits are used to specify unique hosts on the subnet. With three bits used to specify hosts on a 29-bit subnet, eight different bit combinations are possible. Of those eight possible host addresses, two are reserved:

000 — Reserved for the network (base address)
001
010
100
110
101
011
111 — Reserved for the broadcast address of the subnet

Note Be careful with zero subnets (subnets with the same base address as a class A, B, or C network). Early implementations of TCP/IP did not allow them. For example, the subnet 192.32.8.0/30 was illegal because it had the same base address as the class C network 192.32.8.0/24, while the subnet 192.32.8.4/30 was legal. Modern implementations of TCP/IP support zero subnets, and the TAOS implementation of RIP treats these subnets the same as any other network. However, you must treat zero subnets consistently throughout your network. Otherwise, you will encounter routing problems.

Table 1-2 shows subnet masks and prefix lengths for a class C network number.

Table 1-2. Decimal subnet masks and prefix lengths

Subnet mask	Number of host addresses	Prefix length
255.255.255.0	254 hosts + 1 broadcast, 1 network base	/24
255.255.255.128	126 hosts + 1 broadcast, 1 network base	/25
255.255.255.192	62 hosts + 1 broadcast, 1 network base	/26
255.255.255.224	30 hosts + 1 broadcast, 1 network base	/27
255.255.255.240	14 hosts + 1 broadcast, 1 network base	/28
255.255.255.248	6 hosts + 1 broadcast, 1 network base	/29
255.255.255.252	2 hosts + 1 broadcast, 1 network base	/30
255.255.255.254	Invalid mask (no hosts)	/31
255.255.255.255	1 host—a host route	/32

The broadcast address of any subnet has the host portion of the IP address set to all ones. The network address (or base address) represents the network itself, because the host portion of the IP address is all zeros. For example, supposing the IP configuration assigns the following address to a remote router:

```
198.5.248.120/29
```

The Ethernet network attached to that router has the following address range:

`198.5.248.120 – 198.5.248.127`

A host route is a special-case IP address with a prefix length of /32. For example:

`198.5.248.40/32`

Host routes are to a single host, rather than to a router or subnet.

Configuring LAN IP interfaces

A LAN IP interface is an Ethernet port configured for IP. TAOS creates an `ip-interface` profile for an Ethernet port when it first detects the presence of the port. For example, the following output shows the default `ip-interface` profiles for the shelf controller and a T1000 module installed in slot 3:

```
admin> dir ip-interface
    6  09/14/2001  10:13:24  { { any-shelf any-slot 0 } 0 }
    8  09/14/2001  10:13:24  { { shelf-1 first-control-module 1 } 0 }
   19  09/14/2001  10:14:02  { { shelf-1 second-control-module 1 } 0 }
    8  09/14/2001  11:36:32  { { shelf-1 slot-3 2 } 0 }
   64  09/14/2001  11:53:12  { { shelf-1 slot-3 1 } 0 }
```

The profile for the first Ethernet port on a T1000 module in shelf 1, slot 3, uses the following index:

`{ { 1 3 1 } 0 }`

This index consists of a physical address and a logical-item number in the following format:

`{ { shelf-num slot-num item-num } logical-item-num }`

The logical item addresses a specific logical interface. It is zero except when multiple (virtual) interfaces have been configured on the physical port. For more details, see “Typical definition of virtual LAN interfaces” on page 1-9.

Overview of LAN interface settings

Following are the parameters in an `ip-interface` profile, shown with default settings:

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 } ]
interface-address* = { { any-shelf any-slot 0 } 0 }
ip-address = 0.0.0.0/0
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only = no
```

Parameter	Setting
interface-address	Physical address of the Ethernet interface in the Stinger unit, or, if the item number is not zero, the virtual interface address.
ip-address	IP address of the LAN interface.
proxy-mode	Enable/disable proxy ARP responses for remote devices that are assigned local addresses.
rip-mode	Enable/disable RIP updates on the interface. RIP is disabled by default on LAN interfaces.
route-filter	Filter for RIP update packets. For details, see Chapter 7, "Packet Filters."
rip2-use-multicast	Enable/disable use of the multicast address (224.0.0.9) rather than the broadcast address for RIP updates. By default, RIP updates use the multicast address.
multicast-allowed	Multicast forwarding option. See Chapter 2, "IGMP Multicast Forwarding."
multicast-rate-limit	Multicast forwarding option. See Chapter 2, "IGMP Multicast Forwarding."
multicast-group-leave-delay	Multicast forwarding option. See Chapter 2, "IGMP Multicast Forwarding."
directed-broadcast-allowed	Enable/disable forwarding of directed broadcast traffic onto the interface and its network.
vrouter	Name of a virtual router. See Chapter 4, "Virtual Routing."
management-only-interface	Enable/disable management-only on the IP interface.

Typical LAN IP interface configuration

The following commands set the IP address of the Ethernet interface of the first control module:

```
admin> dir ip-interface
    6  09/14/2001  10:13:24  { { any-shelf any-slot 0 } 0 }
    8  09/14/2001  10:13:24  { { shelf-1 first-control-module 1 } 0 }
   19  09/14/2001  10:14:02  { { shelf-1 second-control-module 1 } 0 }

admin> read ip-interface { { 1 8 1 } 0 }
IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 } read

admin> set ip-address = 10.1.2.65/24

admin> write
IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 } written
```

In this example, the Ethernet interface is connected to the 10.1.2 subnet. To enable the interface to communicate with routers on other local subnets, either the system must have a static route configuration to another router in its own subnet, or the interface must enable RIP. (For an example of configuring a route to a local router, see "Typical default route configurations" on page 1-26.)

After you assign an IP address, you can verify that the Stinger unit is a valid IP host on that network segment by pinging another host, as shown in the following example:

```
admin> ping 10.65.212.19
PING 10.65.212.19: 56 Data bytes
64 bytes from 10.65.212.19: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 10.65.212.19: icmp_seq=3 ttl=255 time=0 ms
^C
--- 10.65.212.19: Ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

Enabling proxy ARP

When you enable proxy ARP, hosts on the local network can ARP for hosts or subnets that reside across the WAN but have an IP address on the local network. The router responds to the ARP requests, and then routes the packets for those connections across the WAN.

You can enable `proxy-mode` by setting it to `active` (respond for active WAN connections only), `inactive` (respond only for inactive WAN connections), or `always` (respond for all pool addresses, including those for inactive connections). If the interface is set to respond to ARP requests for inactive sessions, the system attempts to establish the required session.

The following commands configure both LAN interfaces of a T1000 module in slot 3 to respond as proxies for ARP requests for active WAN connections:

```
admin> read ip-interface { { 1 3 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } read

admin> set proxy-mode = active

admin> write
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } written

admin> read ip-interface { { 1 3 2 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 2 } 0 } read

admin> set proxy-mode = active

admin> write
IP-INTERFACE/{ { shelf-1 slot-3 2 } 0 } written
```

Enabling RIP

RIP is off by default, so the router does not send out its routing table or receive routing information from other routers on the interface. Therefore, local hosts on other subnets cannot access remote hosts without static route configurations, and remote hosts do not have access to other routes maintained locally.

You can enable RIP to receive routing table updates, send them, or both. Receiving updates from other routers increases the size of the system's routing table. The table then provides access to more networks, but route searches are not as fast. Sending updates propagates information about remote networks to local routers.

Note Running RIP-2 and RIP-v1 on the same network in such a way that the routers receive each other's advertisements is not recommended. RIP-v1 guesses subnet masks, while RIP-2 handles them explicitly. Running the two versions on the

same network can result in RIP-v1 guesses overriding accurate subnet information obtained via RIP-2.

The following commands configure a T1000 LAN interface to receive RIP-2 updates on the multicast address:

```
admin> read ip-interface { { 1 3 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } read

admin> set rip-mode = routing-recv-v2

admin> write
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } written
```

Typical definition of virtual LAN interfaces

You can configure up to 16 `ip-interface` profiles for each T1000 module as a whole, with each profile specifying one IP address. The system creates the default profile for an interface and assigns it the zero logical-item number. To configure another IP address on a LAN interface, create an `ip-interface` profile with a nonzero logical-item number in its interface address. For example, the following commands create a virtual interface for an Ethernet port installed in shelf 1, slot 3:

```
admin> new ip-interface { { 1 3 1 } 1 }
IP-INTERFACE/{ { shelf-1 slot-3 1 } 1 } read

admin> set ip-addr = 10.9.1.212/24

admin> write
IP-INTERFACE/{ { shelf-1 slot-3 1 } 1 } written
```

The logical-item numbers do not have to be consecutive, but they must each be unique. The following restrictions apply to virtual LAN interfaces:

- The default `ip-interface` profile (with the zero logical-item number) must have an IP address configured. Otherwise, none of the other `ip-interface` profiles for the same port can function. (Do not delete the default profile and expect your other configurations to work.)
- If `proxy-mode` is enabled in any of the `ip-interface` profiles for a given Ethernet port, it is enabled for all ARP requests coming into the physical port.

Typical definition of the soft interface

TAOS supports a soft IP interface, which is an internal interface that is always active. As long as one of the system's IP interfaces is up, the soft interface address is reachable.

Note Do not use the IP address of a physical LAN interface for the soft interface address.

The `ip-interface` profile with the zero index is reserved for the soft interface. If RIP is enabled, the system advertises the interface address as a host route (with a prefix length of 32 bits) using the loopback interface. If RIP is not enabled, routers one hop away from the unit must have a static route to the soft interface address.

The following commands set the soft interface IP address to 1.1.1.128/24:

```
admin> read ip-interface { 0 0 0 }
IP-INTERFACE/{ { any-shelf any-slot 0 } 0 } read
```

```
admin> set ip-addr = 1.1.1.128/24

admin> write
IP-INTERFACE/{ { any-shelf any-slot 0 } 0 } written
```

To create an interface-independent address for a virtual router, create a new `ip-interface` profile with the logical-item value greater than zero. For example:

```
admin> new ip-interface
IP-INTERFACE/{ { any-shelf any-slot 0 } 0 } read

admin> set interface-address = { { 0 0 0 } 1 }
(New index value; will save profile as IP-INTERFACE/{ { any-shelf any-
slot 0 } 1 }.)

admin> set ip-addr = 10.10.1.1

admin> write
IP-INTERFACE/{ { any-shelf any-slot 0 } 1 } written
```

TAOS adds the soft address to its interface table with the name `sip#`, where # is the logical-item number from the `ip-interface` profile index. For more details about virtual routers, see Chapter 4, "Virtual Routing."

If RIP updates are enabled, the system advertises the interface address as a host route with a mask of /32, using the loopback interface. If RIP is not enabled, routers one hop away must have a static route to the soft address. To verify that other hosts in your network have a route to the soft address, run `ping` or `traceroute` from the other hosts. For example:

```
host1% ping 11.168.7.100
PING 11.168.7.100 (11.168.7.100): 56 Data bytes
64 bytes from 11.168.7.100: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 11.168.7.100: icmp_seq=7 ttl=255 time=0 ms
^C
--- 11.168.7.100 Ping statistics ---
8 packets transmitted, 8 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

Example of disabling directed broadcasts

Denial-of-service attacks known as *smurf* attacks typically use ICMP Echo Request packets with a spoofed source address and packets directed to IP broadcast addresses. These attacks are intended to degrade network performance, possibly to the point that the network becomes unusable.

To prevent the IP router from being used as an intermediary in this type of denial-of-service attack launched from another network, you must disable the router from forwarding directed broadcasts it receives from another network. The following example shows how to disable directed broadcasts that are not generated locally. All IP interfaces in the system must disable the feature explicitly. The sample commands configure both control module interfaces (so the broadcasts are still disabled if controller switchover occurs) and the IP interfaces of a T1000 module in shelf 1, slot 3.

```
admin> read ip-interface { { 1 8 1 } 0 }
IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 } read

admin> set directed-broadcast-allowed = no
```

```
admin> write
IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 } written
admin> read ip-interface { { 1 9 1 } 0 }
IP-INTERFACE/{ { shelf-1 second-control-module 1 } 0 } read
admin> set directed-broadcast-allowed = no
admin> write
IP-INTERFACE/{ { shelf-1 second-control-module 1 } 0 } written
admin> read ip-interface { { 1 3 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } read
admin> set directed-broadcast-allowed = no
admin> write
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } written
admin> read ip-interface { { 1 3 2 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 2 } 0 } read
admin> set directed-broadcast-allowed = no
admin> write
IP-INTERFACE/{ { shelf-1 slot-3 2 } 0 } written
```

Typical management-only interface configuration

Management-only means that incoming traffic on the interface terminates in the system itself. The traffic is not forwarded on any other interface. In addition, only traffic generated by the system is forwarded on the management-only interface. Traffic generated externally is dropped on the interface. Setting the `management-only` parameter to `yes` enforces these conditions on the port.

The following commands specify that the first Ethernet port on a T1000 module is a management-only interface:

```
admin> read ip-interface { { 1 3 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } read
admin> set management-only = yes
admin> write
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } written
```

The `ifmgr -d` command displays a `Management Only` field to reflect the port's status.

Configuring WAN IP interfaces

A WAN IP interface is a connection configured for IP. TAOS creates a routing interface for local `connection` profiles (if they do not use pool addresses) when the system starts up. For interfaces that use pool addresses or are defined in RADIUS user profiles, the system creates a routing interface when a session becomes active.

Overview of WAN interface settings

You configure WAN IP interfaces in `connection` profiles or RADIUS profiles. At a minimum, each profile specifies the IP address of the far-end device or a pool from which the system assigns an address. You can also set a variety of routing and service parameters.

Settings in connection profiles

Following are the IP options (shown with default settings) in a connection profile:

```
[in CONNECTION/"":ip-options]
ip-routing-enabled = yes
vj-header-prediction = yes
remote-address = 0.0.0.0/0
local-address = 0.0.0.0/0
routing-metric = 1
preference = 60
down-preference = 120
private-route = no
multicast-allowed = no
address-pool = 0
ip-direct = 0.0.0.0
rip = routing-off
route-filter = ""
source-ip-check = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
client-dns-primary-addr = 0.0.0.0
client-dns-secondary-addr = 0.0.0.0
client-dns-addr-assign = yes
client-default-gateway = 0.0.0.0
tos-filter = ""
client-wins-primary-addr = 0.0.0.0
client-wins-secondary-addr = 0.0.0.0
client-wins-addr-assign = yes
private-route-table = ""
private-route-profile-required = no
```

```
[in CONNECTION/"":ip-options:tos-options]
active = no
precedence = 000
type-of-service = normal
apply-to = incoming
marking-type = precedence-tos
dscp = 00
```

Parameter	Setting
ip-routing-enabled	Enable/disable IP routing for the interface. IP routing is enabled by default.
vj-header-prediction	Enable/disable Van Jacobson prediction for TCP packets on incoming calls using encapsulation protocols that support Van Jacobson compression. The default setting is yes.
remote-address	IP address of the remote device, which can include a subnet specification. If the address does not include a subnet mask, the router assumes the default subnet mask based on address class.

Parameter	Setting
local-address	IP address assigned to the local side of a numbered-interface connection. (For details, see "Typical numbered-interface configuration" on page 1-18.)
routing-metric	RIP metric for the specified route (a number from 1 to 15, default 1). If preference values are equal, the higher the metric, the less likely that the router will use the route.
preference	Preference value for the route. Valid values are from 0 to 255.
down-preference	Preference value for the route when the interface is down.
private-route	Enable/disable advertisement of the route when the router sends RIP or OSPF updates. With the yes setting, the route is excluded from update packets.
multicast-allowed	See Chapter 2, "IGMP Multicast Forwarding.".
address-pool	Number of the address pool from which to acquire an address (see "Configuring and using address pools" on page 1-60).
ip-direct	IP address of a host to which all IP packets received across the link will be directed (see "Typical IP-direct configuration" on page 1-20).
rip	Enable/disable RIP updates on the interface. RIP is disabled by default.
route-filter	Filter for RIP update packets. For details, see Chapter 7, "Packet Filters."
source-ip-check	Enable/disable antispoofing for the session. With the yes setting, the system does not accept packets that do not originate on the subnet to which the remote device is attached. The system determines the subnet during IPCP negotiation. If remote-address specifies a subnet, packets that originate on that subnet are accepted. If remote-address specifies a 32-bit mask, only packets from that host are accepted. Packets sent from an address that does not match are discarded.
multicast-rate-limit	Multicast forwarding option. (See Chapter 2, "IGMP Multicast Forwarding.")
multicast-group-leave-delay	Multicast forwarding option. (See Chapter 2, "IGMP Multicast Forwarding.")
client-dns-primary-addr	Client DNS option. (See "Using client DNS" on page 1-56.)
client-dns-secondary-addr	Client DNS option. (See "Using client DNS" on page 1-56.)
client-dns-addr-assign	Client DNS option. (See "Using client DNS" on page 1-56.)

Parameter	Setting
client-default-gateway	Default route for traffic from this connection. For details, see "Typical configuration of client default gateways" on page 1-21.
tos-filter	Name of a filter profile defining a type of service (TOS) filter. TOS filters are used to enable proxy-QoS for all packets that match the filter specification. For details, see Chapter 7, "Packet Filters."
client-wins-primary-addr	Windows Internet Name Service (WINS) option. See "Configuring Microsoft WINS assignment" on page 1-58.
client-wins-secondary-addr	WINS option. See "Configuring Microsoft WINS assignment" on page 1-58.
client-wins-addr-assign	WINS option. See "Configuring Microsoft WINS assignment" on page 1-58.
private-route-table	Private routing table option. See "Configuring private routing tables" on page 1-29.
private-route-profile-required	Private routing table option. See "Configuring private routing tables" on page 1-29.
tos-options:active	Enable/disable proxy-QoS and TOS for this connection (see "Typical QoS and TOS policy configuration" on page 1-22).
tos-options:precedence	Priority level of the data stream. The three most significant bits of the TOS byte are priority bits used to set precedence for priority queuing. When TOS is enabled, you can set the bits to one of the following values (most significant bit first): ■ 000—Normal priority ■ 001—Priority level 1 ■ 010—Priority level 2 ■ 011—Priority level 3 ■ 100—Priority level 4 ■ 101—Priority level 5 ■ 110—Priority level 6 ■ 111—Priority level 7 (the highest priority)
tos-options:type-of-service	Type of service of the data stream. The next four bits of the TOS byte are used to choose a link according to the type of service. When TOS is enabled, specify one of the following values: ■ normal—Normal service ■ cost—Minimize monetary cost ■ reliability—Maximize reliability ■ throughput—Maximize throughput ■ latency—Minimize delay

Parameter	Setting
`tos-options:apply-to`	Direction in which TOS is enabled. Specify one of the following values: ■ `incoming` (the default)—Bits are set in packets received on the interface. ■ `outgoing`—Bits are set in outbound packets only. ■ `both`—Incoming and outgoing packets are tagged.
`tos-options:marking-type`	Enable/disable marking of packets to provide information that can be utilized by other network elements in a network domain utilizing differentiated services code point (DSCP). Specify one of the following values: ■ `precedence-tos` (the default)—The system marks packets in a manner consistent with RFC 791, in which the first six bits in the second octet indicate the precedence and type of service (TOS) of the packet, as specified in the `precedence` and `type-of-service` settings. ■ `dscp`—The system marks packets as specified RFC 2474, making use of the DSCP value specified in the `dscp` setting. Changing the value of this setting takes effect for new connections when the profile is written. This setting does not apply if IP routing or TOS is disabled.
`tos-options:dscp`	DSCP value if DSCP is specified in the `marking-type` parameter. Values can range from 00 to FF hexadecimal specifying different classes of service, as defined for IPv4 headers in RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*. Changing the value of this setting takes effect for new connections when the profile is written. This setting does not apply if IP routing or TOS is disabled, or if `marking-type` is set to `precedence-tos`.

Settings in RADIUS profiles

The following attribute-value pairs configure IP options in a RADIUS profile:

RADIUS attribute	Value
Ascend-Route-IP (228)	Enable/disable IP routing for the interface. IP routing is enabled by default.
Framed-Compression (13)	Enable/disable Van Jacobson prediction. You can specify Van-Jacobson-TCP-IP to turn on TCP/IP header compression. If you do not specify this value, RADIUS uses the default of no header compression.

RADIUS attribute	Value
Framed-IP-Address (8)	IP address of the calling device.
Framed-IP-Netmask (9)	Subnet mask of the caller's address. If you do not specify a subnet mask, the router uses the default subnet mask based on address class.
Ascend-PPP-Address (253)	IP address assigned to the local side of a numbered-interface connection. For details, see "Typical numbered-interface configuration" on page 1-18.
Ascend-IF-Netmask (153)	Subnet mask in use for the local-side numbered interface.
Ascend-Metric (225)	RIP metric for the specified route (a number from 1 to 15, default 7). If preference values are equal, the higher the metric, the less likely that the router will use the route.
Ascend-Route-Preference (126)	Preference value for the route. Valid values are from `0` to `255`. A value of `255` prevents the use of the route.
Framed-Route (22)	Static route definition, which can be used to make a user profile a private route. For details, see "Configuring static IP routes" on page 1-23.
Ascend-Assign-IP-Pool (218)	Number of the address pool from which to acquire an address. For details, see "Configuring and using address pools" on page 1-60.
Ascend-Assign-IP-Global-Pool (146)	Name of a global address pool. For details, see "Global RADIUS pools (RADIPAD)" on page 1-61.
Ascend-IP-Direct (209)	IP address of a host to which all IP packets received across the link will be directed. For details, see "Typical IP-direct configuration" on page 1-20.
Framed-Routing (10)	Enable/disable RIP updates on the interface. RIP is disabled by default. Valid values are None(0), Broadcast(1), Listen(2), Broadcast-Listen(3), Broadcast-v2(4), Listen-v2(5), and Broadcast-Listen-v2(6).
Ascend-Source-IP- Check (96)	Enable/disable antispoofing for the session. The default is Source-IP-Check-No (0). With the Source-IP-Check-Yes (1) setting, the system discards packets that do not originate on the subnet to which the remote device is attached. The system determines the subnet during IPCP negotiation. If Framed-IP-Netmask specifies a subnet, packets that originate on that subnet are accepted. If Framed-IP-Netmask specifies a 32-bit mask, only packets from a single host are accepted. Packets sent from an address that does not match are discarded.
Ascend-Multicast-Client (155)	Multicast forwarding option. (See Chapter 2, "IGMP Multicast Forwarding.")
Ascend-Multicast-Rate-Limit (152)	Multicast forwarding option. (See Chapter 2, "IGMP Multicast Forwarding.")

RADIUS attribute	Value
Ascend-Multicast-Leave-Delay (111)	Multicast forwarding option. (See Chapter 2, "IGMP Multicast Forwarding.")
Ascend-Client-Primary-DNS (135)	Client DNS option. (See "Using client DNS" on page 1-56.)
Ascend-Client-Secondary-DNS (136)	Client DNS option. (See "Using client DNS" on page 1-56.)
Ascend-Client-Assign-DNS (137)	Client DNS option. (See "Using client DNS" on page 1-56.)
Ascend-Client-Gateway (132)	Default route for traffic from this connection. For details, see "Typical configuration of client default gateways" on page 1-21.
Ascend-IP-TOS (87)	Type of service (TOS) of the data stream. The value of this attribute sets the four bits following the three most significant bits of the TOS byte. The four bits are used to choose a link according to the type of service. Specify one of the following values: ■ Ascend-IP-TOS IP-TOS-Normal (0)—Normal service ■ Ascend-IP-TOS IP-TOS-Disabled (1)—Disable TOS ■ Ascend-IP-TOS IP-TOS-Cost (2)—Minimize monetary cost ■ Ascend-IP-TOS IP-TOS-Reliability (4)—Maximize reliability ■ Ascend-IP-TOS IP-TOS-Throughput (8)—Maximize throughput ■ Ascend-IP-TOS IP-TOS-Latency (16)—Minimize delay ■ Ascend-IP-TOS IP-TOS-Dscp (128)—Differentiated services code point. See the Ascend-IP-DSCP (3) attribute for related information.
Ascend-IP-TOS-Precedence (88)	Priority level of the data stream. The three most significant bits of the TOS byte are priority bits used to set precedence for priority queuing. If TOS is enabled, set the bits to one of the following values (most significant bit first): ■ IP-TOS-Precedence-Pri-Normal (0)—Normal priority ■ IP-TOS-Precedence-Pri-One (32)—Priority level 1 ■ IP-TOS-Precedence-Pri-Two (64)—Priority level 2 ■ IP-TOS-Precedence-Pri-Three (96)—Priority level 3 ■ IP-TOS-Precedence-Pri-Four (128)—Priority level 4 ■ IP-TOS-Precedence-Pri-Five (160)—Priority level 5 ■ IP-TOS-Precedence-Pri-Six (192)—Priority level 6 ■ IP-TOS-Precedence-Pri-Seven (224)—Priority level 7 (the highest priority)

RADIUS attribute	Value
Ascend-IP-TOS-Apply-To (89)	Direction in which TOS is enabled. With the IP-TOS-Apply-To-Incoming (1024) setting, which is the default, bits are set in packets received on the interface. With the IP-TOS-Apply-To-Outgoing (2048) setting, bits are set in outbound packets only. With the IP-TOS-Apply-To-Both (3072) setting, both incoming and outgoing packets are tagged.
Ascend-IP-DSCP (3)	DSCP value if Ascend-IP-TOS IP-TOS-Dscp (128) is specified as the value of the Ascend-IP-TOS attribute. Values can range from 00 to FF hexadecimal specifying different classes of service, as defined for IPv4 headers in RFC 2474.

Typical WAN connection to another IP router

Figure 1-2 shows a Stinger unit connecting to a router customer premises equipment (CPE).

Figure 1-2. Router-to-router IP connection

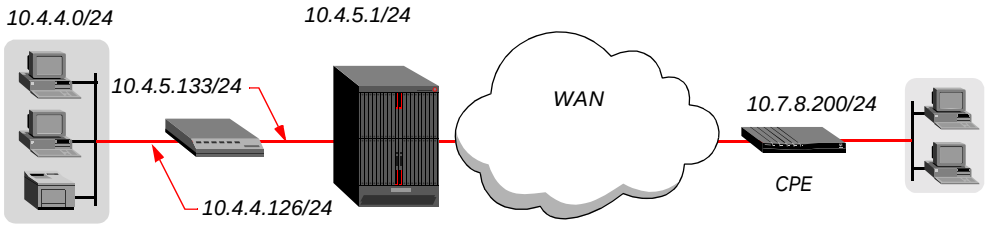

The default settings for the `ip-options` subprofile enable IP routing and Van Jacobson header compression and turn RIP off. Those settings are appropriate for the following example, which shows configuration of a `connection` profile for the CPE in Figure 1-2:

```
admin> read connection router-1
CONNECTION/router-1 read
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp send-auth-mode = pap-ppp-auth
admin> set ppp recv-password = localpw
admin> set ip-options remote = 10.7.8.200/24
admin> write
CONNECTION/router-1 written
```

Typical numbered-interface configuration

For a numbered-interface connection, assign each side of the connection with a unique address that applies only to that connection. Some applications, such as SNMP have this requirement.

The local-address value assigned to a numbered interface must be unique to the connection and to the network. You can assign a fake IP address or an IP address from one of the local subnets. A Stinger unit accepts IP packets destined for the specified address and treats them as destined for the system itself. (The packets can arrive on any interface, and the destination interface need not be in the active state.)

Caution Do not assign a local address that belongs to one of the Stinger unit's real, physical LAN interfaces. Doing so causes routing problems.

Figure 1-3 shows a numbered-interface connection. The Stinger unit's real, physical Ethernet interface has the IP address 10.5.6.7/24. The other two addresses represent the local and remote addresses of the numbered-interface connection.

Figure 1-3. A numbered-interface connection

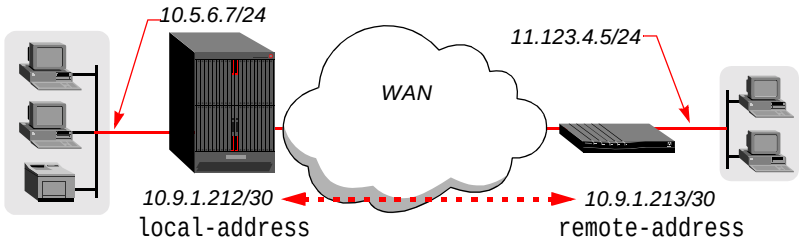

The following set of commands specifies a connection profile for the numbered interface:

```
admin> new connection numbered
CONNECTION/numbered read
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp recv-password = localpw
admin> set ip-options remote-address = 10.9.1.213/30
admin> set ip-options local-address = 10.9.1.212/30
admin> write
CONNECTION/numbered written
```

Following is a comparable RADIUS profile:

```
numbered Password = "localpw"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Ascend-Route-IP = Route-IP-Yes,
   Framed-IP-Address = 10.9.1.213,
   Framed-IP-Netmask = 255.255.255.252,
   Ascend-PPP-Addr = 10.9.1.212,
   Ascend-IF-Netmask = 255.255.255.252
```

In this example, the interface is assigned a 30-bit subnet, so four bit combinations are available for host assignments. Of the four possible host addresses, the one that is evenly divisible by 4 is the network or base address (the address that specifies zeros in the host bits). This address is added to the routing table. The other host addresses are assigned a /32 subnet mask and added as host routes. You can suppress advertisement of the host routes associated with a numbered interface by using the

`suppress-host-routes` parameter, as described in "Suppressing host-route advertisements" on page 1-44.

Typical IP-direct configuration

Packets received on an IP-direct connection bypass the routing tables and are redirected instead to a specified next-hop destination IP address. Outbound packets are routed as usual. Currently, the feature is implemented only for data calls. Figure 1-4 shows an example of the IP-direct traffic flow.

Figure 1-4. IP-Direct connections

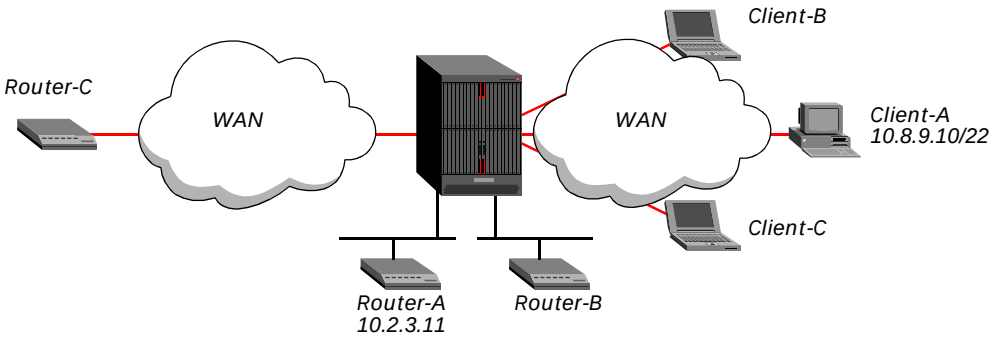

In Figure 1-4, the following conditions apply:

- Client-A's profile redirects inbound packets to router-A on a LAN interface.
- Client-B's profile redirects inbound packets to router-B on a LAN interface.
- Client-C's profile redirects inbound packets to router-C through a switched connection.

Outbound packets destined for any of the three clients are routed normally, which means that these client connections can *receive* packets from any source, not just from the IP address to which their packets are sent.

The following set of commands configures an IP-Direct `connection` profile for client-A:

```
admin> read connection client-A
CONNECTION/client-A read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp recv-password = localpw

admin> set ip-options remote = 10.8.9.10/22

admin> set ip-options ip-direct = 10.2.3.11

admin> write
CONNECTION/client-A written
```

Following is a comparable RADIUS profile:

```
client-A Password = "localpw"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Framed-IP-Address = 10.8.9.10,
```

```
   Framed-IP-Netmask = 255.255.252.0,
   Ascend-IP-Direct = 10.2.3.11
```

IP-direct connections require the following special handling:

- If the profile enables the receipt or receipt-transmission of RIP updates, all RIP packets from an incoming connection are kept locally and forwarded to the IP-Direct address, so that the system can correctly forward packets *destined* for the client.
- ARP requests received from the incoming connection are ignored.
- The caller cannot `telnet` to the Stinger unit, because the connection is passed through to the IP-direct host.

Example of making the route to a connection private

A private route is placed in the routing table but is marked with a flag that prevents routing protocols from advertising it. The following commands specify a private route in a `connection` profile:

```
admin> read connection david
CONNECTION/david read

admin> set ip-options remote = 10.8.9.10/24

admin> set ip-options private = yes

admin> set ip-options routing-metric = 3

admin> write
CONNECTION/david written
```

Following is a comparable RADIUS profile:

```
david Password = "localpw"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
   Framed-IP-Address = 10.8.9.10,
   Framed-IP-Netmask = 255.255.255.0,
   Framed-Route = "10.8.9.10/24 0.0.0.0 3 y"
```

Typical configuration of client default gateways

A client default gateway is a route that replaces the systemwide default route for a particular connection. For packets arriving on the connection, if the system consults the routing table and finds no match for the packets' destination (if it finds only the system default route or, if there is no system default route, no match at all) it forwards the packets to the client default gateway address.

The specified address must be a legitimate next hop. That is, the system must be able to reach the router directly in one hop. If this is not the case, the system drops packets and does not route them to the client default gateway.

Packets from other users or from the Ethernet network are handled normally. The system's routing table is not modified by use of this feature. The following commands specify a connection-specific default gateway:

```
admin> read connection test
CONNECTION/test read

admin> set ip-options client-default-gateway = 17.1.1.1
```

```
admin> write
CONNECTION/test written
```

Following is a comparable setting in a RADIUS profile:

```
test Password = "localpw"
   Service-Type = Framed-User,
   Ascend-Client-Gateway = 17.1.1.1
```

Example of per-session source address checking

You can configure WAN IP interfaces so that the system checks the source IP address in all received packets and drops the packets if the address does not match the address negotiated for the far-end subnet. This type of configuration enables the system to detect packets with a spoofed source IP addresses and discard them.

When the system initially detects a spoofing attempt (a mismatched source address), it logs a message that includes the port number on which the attempt occurred. For example:

```
[1/4/1/1] Spoofing Attempt:from port 1[MBID 1; 1119855018][ed-mc1-p75]
```

The following commands configure a `connection` profile for antispoofing:

```
admin> read connection ed-mc1-p75
CONNECTION/ed-mc1-p75 read

admin> set ip-options source-ip-check = yes

admin> write
CONNECTION/ed-mc1-p75 written
```

Following is a comparable setting in a RADIUS profile:

```
ed-mc1-p75 Password = "localpw"
   Service-Type = Framed-User,
   Ascend-Source-IP-Check = Source-IP-Check-Yes
```

Typical QoS and TOS policy configuration

You can configure the router to set quality of service (QoS) priority bits and type of service (TOS) classes of service on behalf of customer applications. The Stinger unit does not implement priority queuing, but it does set information that can be used by other routers to prioritize and select links for particular data streams.

To enable service-based TOS or to set QoS precedence for the traffic on a particular WAN connection, configure the TOS options in a `connection` or RADIUS profile. The settings cause the system to set bits in the TOS byte of IP packet headers that are received (the default), transmitted, or both, on the WAN interface. Another router can then interpret the bits accordingly.

You can also specify proxy-QoS and TOS policy in a TOS filter, which can then be applied to any number of `connection` or RADIUS profiles. For a `connection` or RADIUS profile that has both its own local policy and an applied TOS filter, the policy defined in the TOS filter takes precedence. For example, applying a TOS filter to a TOS-enabled connection allows you to define one priority setting for incoming packets on a connection and another for incoming packets addressed to a particular destination (the destination in a TOS filter). For details, see Chapter 7, “Packet Filters.”

The following set of commands enables TOS for incoming packets on a WAN interface. It sets the priority of the packets at 6, which means that another router that implements priority queuing will not drop the packets until it has dropped all packets of a lower priority. The commands also set TOS to prefer maximum throughput, which means that the priority-queuing router will choose a high bandwidth connection if one is available, even if it has higher cost or higher latency or is less reliable than another available link.

```
admin> read connection jfan-pc
CONNECTION/jfan-pc read
admin> set ip-options remote-address = 10.168.6.120/24
admin> set ip-options tos active = yes
admin> set ip-options tos precedence = 110
admin> set ip-options tos type = throughput
admin> write
CONNECTION/jfan-pc written
```

Following is a comparable RADIUS profile:

```
jfan-pc Password = "localpw"
   Service-Type = Framed-User,
   Framed-IP-Address = 10.168.6.120,
   Framed-IP-Netmask = 255.255.255.0,
   Ascend-IP-TOS = IP-TOS-Throughput,
   Ascend-IP-TOS-Precedence = IP-TOS-Precedence-Pri-Six,
   Ascend-IP-TOS-Apply-To = IP-TOS-Apply-To-Incoming
```

Configuring static IP routes

Any profile that specifies how to reach an IP device or subnet (such as an `ip-interface`, `connection`, or RADIUS user profile specifies a static IP route to that destination. However, sometimes administrators configure static routes in a more flexible way, to extend or fine-tune the routing table.

The default route is a special-case static route that acts as a catch-all for packets for which the Stinger unit cannot find a route. If you define a default route (with the zero destination address), the system routes all packets with unknown destinations to the specified gateway. If no default route is defined, the system drops those packets.

If the unit's LAN IP addresses include subnet specifications, you must create a static route to another LAN router to enable the system to reach local networks beyond its own subnets. You might also configure a static route to a LAN router to offload local routing overhead from the Stinger unit.

Another reason to configure static routes is to specify multipath routes, which define multiple paths to the same destination. Multipath routes, with equal metric values, distribute traffic to a single destination across multiple interfaces.

Overview of static-route settings

You can define static routes in `ip-route` profiles or in RADIUS.

Settings in ip-route profiles

Following are the parameters in a local ip-route profile (shown with default settings):

```
in IP-ROUTE/"" (new)]
name* = ""
dest-address = 0.0.0.0/0
gateway-address = 0.0.0.0
metric = 8
private-route = no
active-route = yes
ase7-adv = N/A
vrouter = ""
inter-vrouter = ""
```

Parameter	Setting
name	Name of the profile (up to 31 characters).
dest-address	Destination IP address, which can include a subnet specification. The default value is 0.0.0.0, which represents the default route. The system forwards packets whose destinations do not match an entry in the routing table to the default route.
gateway-address	IP address of a next-hop router used to reach the specified destination. A next-hop router is directly connected to the Stinger unit on the Ethernet, or is one hop away on a WAN link.
metric	RIP metric for the specified route (a number from 0 to 15, default 8). RIP is a distance-vector protocol that uses hop count as its metric. Among routes with the same destination address, the higher the metric, the less likely that the system will choose the route.
private-route	Enable/disable advertisement of the route when the router sends RIP updates. With the yes setting, the route is excluded from update packets.
active-route	Enable/disable entering the route in the routing table. (Setting the parameter to no is a useful way to make a route temporarily inactive, so you can reinstate the route later.)
vrouter	Name of a virtual router that owns the static route. For details, see Chapter 4, “Virtual Routing.”
inter-vrouter	Virtual router option. For details, see Chapter 4, “Virtual Routing.”

Settings in a RADIUS route profile

A route profile is a pseudo-user profile in which the first line has the following format:

route-*name-N* Password = "ascend", Service-Type = Outbound-User

The *name* argument is the system name (specified by the name parameter in the system profile), and *N* is a number in a sequential series, starting with 1. Make sure there are no missing numbers in the series specified by *N*. If there is a gap in the sequence of numbers, the system stops retrieving the profiles when it encounters the gap.

Each pseudo-user profile specifies one or more routes with the Framed-Route (22) attribute. The RADIUS protocol limits the number of Framed-Route definitions in a single route profile. The limit varies with the exact contents of the routes. However, 25 Framed-Route definitions per profile is the recommended maximum.

The value of the Framed-Route attribute uses the following syntax:

dest-addr gateway-addr metric [*private*] [*profile*] [*preference*] [*VRouter*]

Syntax element	Description
dest-addr	Destination IP address, which can include a subnet specification. The default value is 0.0.0.0, which represents a default route.
gateway-addr	IP address of the next-hop router to reach the specified destination.
metric	RIP metric for the specified route (a number from 1 to 15, default 8). If preference values are equal, the higher the metric, the less likely that the system will use the route.
private	Enable/disable advertisement of the route when the router sends RIP updates. The yes setting makes the route private, excluding it from update packets.
profile	Name of the user profile for the route. The default value is null.
VRouter	Virtual router option. For details, see “Defining virtual router static routes” on page 4-11.

Route settings in a RADIUS user profile

You can also include the Framed-Route (22) attribute in a RADIUS user profile to define a static route. For details about Framed-Route usage, see “Settings in a RADIUS route profile” on page 1-24.

In a user profile, you can specify the zero address as the gateway address. In this context, the 0.0.0.0 address is a wildcard entry the system replaces with the caller’s IP address. When RADIUS authenticates the caller and sends the Stinger unit an Access Accept message with a value of 0.0.0.0 for the router address, the system updates its routing tables with the Framed-Route value, but substitutes the caller’s IP address for the router address. This setting is useful when the system assigns an IP address from an address pool and RADIUS does not have access to the IP address of the caller.

If a Framed-Route definition in a user profile duplicates a route defined in the routing table or ip-route profile, the user profile definition takes precedence while the connection is active. For example, suppose a static route to network 10.10.10.10 is defined in a local ip-route profile, with a metric of 10. A RADIUS user profile defines a static route to 10.10.10.10 with a metric of 7. When the RADIUS user’s route is not in use, the routing table indicates that the route has a metric of 10. When the route is in use, the routing table indicates that the route has a metric of 7, with

an `r` in the Flags column to indicate that the route came from RADIUS. Furthermore, the route with a metric of 10 remains in the routing table, with an asterisk (*) in the flags column, indicating that it is a hidden route.

Connection-specific private static routes (RADIUS only)

The following attribute-value pair configures IP options in a RADIUS profile:

RADIUS attribute	Value
Ascend-Private-Route (104)	A private framed route known only to the profile in which it is specified. The value is a destination address and next-hop router address (in that order). For details, see “Examples of private static routes” on page 1-34.

Typical default route configurations

A route with the zero destination address is a default route. If the system does not find a route for a packet’s destination, rather than dropping the packet, it forwards it to a default route. If there is no default route in the routing table, the router drops any packet for which it cannot find a route.

The system creates an `ip-route` profile named `default`, but the profile is not valid until you specify a gateway address, so the route is not active until you assign an address and activate the route. For example:

```
admin> read ip-route default
IP-ROUTE/default read
admin> set gateway-address = 10.10.10.10
admin> set active-route = yes
admin> write
IP-ROUTE/default written
```

You can create a default route by modifying the `default` profile, or by creating one or more `ip-route` profiles that specify a zero destination and a valid gateway address.

Example of a LAN-based default route

Figure 1-5 shows a router that resides on the same subnet as one of the Stinger unit’s local IP interfaces.

Figure 1-5. Default route to a local IP router

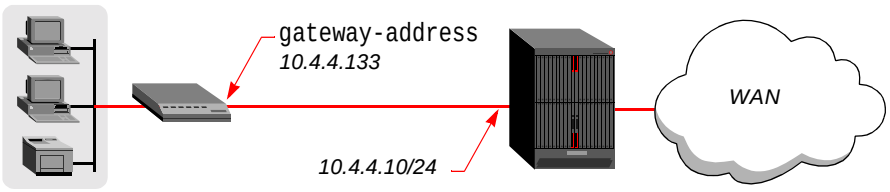

Because the Stinger unit is located on a subnet, it needs to be informed about other backbone routers that can route beyond the subnet. In this example, the system offloads part of its routing overhead by using a default route to the LAN router. The following commands define a default route to the local router:

```
admin> new ip-route lanroute-1
IP-ROUTE/lanroute-1 read

admin> set gateway-address = 10.4.4.133

admin> write
IP-ROUTE/lanroute-1 written
```

Following is a comparable RADIUS default route:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
   Framed-Route = "0.0.0.0 10.4.4.133"
```

Example of a default route across a WAN link

Figure 1-6 shows a router that resides across a frame relay DLCI interface. If the WAN link to this default route goes down for any reason, the Stinger unit removes the route from its routing table.

Figure 1-6. Default route across a frame relay DLCI interface

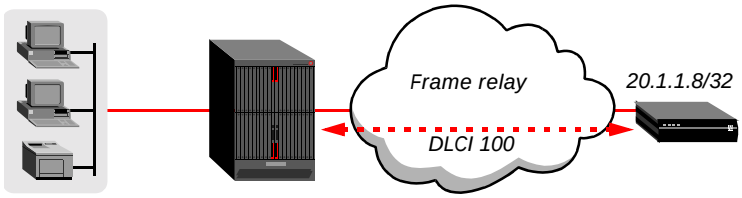

In this example, the following frame relay settings define the data link:

```
[in FRAME-RELAY/fr1]
fr-name* = fr1
active = yes
nailed-up-group = 1
link-mgmt = ansi-t1.617d
link-type = dte
```

The following connection profile defines the DLCI interface:

```
[in CONNECTION/pvc1]
station* = pvc1
active = yes
encapsulation-protocol = frame-relay
ip-options = { yes yes 20.1.1.8/32 0.0.0.0/0 1 60 120 no no 0 0.0.0.0+
telco-options = { ans-and-orig no ft1 1 no no 56k-clear 0 "" "" no no+
fr-options = { fr1 16 "" no "" 16 }
```

The following commands define a default route to the remote device:

```
admin> new ip-route dlci
IP-ROUTE/dlci read

admin> set gateway-address = 20.1.1.8

admin> set private-route = yes

admin> write
IP-ROUTE/dlci written
```

Following is a comparable RADIUS default route:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
   Framed-Route = "0.0.0.0 20.1.1.8 y"
```

Example of configuring a route to a remote subnet

When RIP is turned off on an IP interface, the router cannot reach other routers on that interface unless it has a static route. For example, if a `connection` profile specifies the destination address of a host on a remote subnet, but the packets must be routed through an intermediary device to reach that host (and RIP is not enabled), you must configure a static route specifying the intermediary device as the gateway. Figure 1-7 shows an example.

Figure 1-7. Static route to a remote subnet

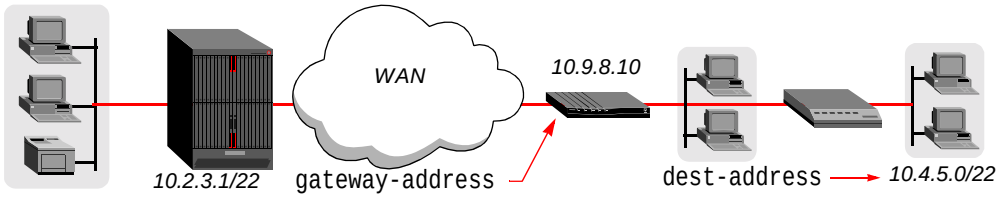

The following commands configure a static route to all hosts on the remote subnet:

```
admin> new ip-route subnet
IP-ROUTE/subnet read
admin> set dest = 10.4.5.0/22
admin> set gateway = 10.9.8.10
admin> write
IP-ROUTE/subnet written
```

Following is a RADIUS profile that shows both the default route and a route to the remote subnet:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
   Framed-Route = "10.4.5.0/22 10.9.8.10"
```

Typical multipath route configuration

Multipath static routes distribute traffic to one destination across the aggregated bandwidth of multiple interfaces. A multipath route requires that the multiple static routes have the same destination address and subnet mask, but different gateway addresses. (Otherwise, the route with the lowest values for these settings is used exclusively.)

Note Even the default routes can be multipath. If more than one route has a destination of 0.0.0.0, the Stinger unit creates multipath default routes.

Following is an example in which an administrator configures a multipath route to the network 10.76.109.0/24:

```
admin> new ip-route bdvnet-1
IP-ROUTE/bdvnet-1 read
admin> set dest = 10.76.109.0/24
admin> set gateway = 11.65.212.1
```

```
admin> set metric = 2

admin> write
IP-ROUTE/bdvnet-1 written

admin> new ip-route bdvnet-2
IP-ROUTE/bdvnet-2 read

admin> set dest = 10.76.109.0/24

admin> set gateway = 11.65.210.1

admin> set metric = 2

admin> write
IP-ROUTE/bdvnet-2 written
```

The multipath routes appear in the routing table with the M (multipath) flag. For example:

```
admin> netstat -rn
Destination      Gateway         IF   Flg   Pref  Met Use   Age
...
10.76.109.0/24 11.65.212.1   ie1-12-2 SGM   100    2   20  7772
10.76.109.0/24 11.65.210.1   ie1-12-3 SGM   100    2   24  7772
```

Configuring private routing tables

You can define private routes through the use of the Ascend-Private-Route (104) attribute in a RADIUS user profile or in private-route pseudo-user profiles, which can then be referred to by multiple RADIUS profiles, `connection` profiles, or both. These externally defined private routing tables are cached locally for a configurable interval. The `prtcache` command displays statistics about each cached RADIUS private-route profile, and enables you to flush profiles from the cache.

You can also define private routing tables locally, in the `private-route-table` profile. These profiles can then be referenced by multiple RADIUS profiles, `connection` profiles, or both.

Overview of local private-route settings

To configure private routing tables, you set the following parameters (shown with default settings):

```
[in PRIVATE-ROUTE-TABLE/""]
name* = ""

[in PRIVATE-ROUTE-TABLE/"":route-description-list[1]]
enabled = no
dest-address = 0.0.0.0/0
netmask = 0.0.0.0
gateway-address = 0.0.0.0
metric = 0

[in CONNECTION/"":ip-options]
private-route-table = ""
private-route-profile-required = no

[in ANSWER-DEFAULTS:ip-answer]
private-route-profile-required = no
```

```
[in IP-GLOBAL]
default-prt-cache-time = 1440
```

Parameter	Setting
PRIVATE-ROUTE-TABLE:name	Name of the profile, up to 23 characters. This name is used to associate a RADIUS or connection profile with the defined private routes.
PRIVATE-ROUTE-TABLE/ route:route-description-list[1]:enabled	Enable/disable the specific route for use in the private routing table. A table can contain up to 24 routes.
PRIVATE-ROUTE-TABLE/ route:route-description-list[1]:dest-address	Destination IP address, which can include a subnet specification. This setting works the same as its counterpart in an ip-route profile (see "Configuring static IP routes" on page 1-23).
PRIVATE-ROUTE-TABLE/ route:route-description-list[1]:netmask	Netmask of the destination IP address, set automatically when you specify a prefix length as part of the IP address.
PRIVATE-ROUTE-TABLE/ route:route-description-list[1]:gateway-address	IP address of a router used to reach the specified destination. This setting works the same as its counterpart in an ip-route profile (see "Configuring static IP routes" on page 1-23).
PRIVATE-ROUTE-TABLE/ route:route-description-list[1]:metric	RIP metric for the private route (a number from 0 to 15, with a default of 0). This setting works the same as its counterpart in an ip-route profile (see "Configuring static IP routes" on page 1-23).
CONNECTION/anyconnection: ip-options:private-route-table	Name of a private-route-table profile associated with the connection. The name can be that of a local profile or of a private-route pseudo-user profile in RADIUS. However, if a local connection profile does not use authentication, it cannot point to a RADIUS private-route profile.
CONNECTION/anyconnection: private-route-profile-required	Whether access to the private routing table is required for the session. Specify yes or no: ■ no (the default)—The system establishes the session even if the associated private routing table is not found. ■ yes—The system disconnects the call if the specified private routing table is not found. This parameter does not apply if the profile does not refer to a private routing table by name. In the answer-defaults profile, this parameter is used for RADIUS user profiles that refer to a private routing table and do not specify a value for Ascend-Private-Route-Required (55).

Parameter	Setting
`ANSWER-DEFAULTS: ip-answer: default-prt-cache-time`	Number of minutes to cache RADIUS private-route profiles that do not include a value for Ascend-Cache-Time (57). The default is `1440` (24 hours). Once the cache timer expires, cached profiles are deleted from system memory. The next time a private route is needed, the system retrieves the profile from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of route lookups, at the cost of some system memory. If this parameter is set to `0` (zero), the default timer is disabled so that only RADIUS profiles specifying a cache time are cached.

Overview of RADIUS private-route settings

RADIUS user profiles can refer to private-route profiles by specifying the following vendor-specific attributes (VSAs):

RADIUS attribute	Value
Ascend-Private-Route-Table-ID (54)	Name of a RADIUS private-route profile associated with the connection.
Ascend-Private-Route-Required (55)	Whether access to the private routing table is required for the session. This attribute does not apply if the profile does not refer to a private routing table by name. If no value is specified for this attribute, the setting for the `private-route-profile-required` parameter in the `answer-defaults` profile is used. Specify one of the following values: ■ Required-No (0) (the default)—The system establishes the session even if the associated private routing table is not found. ■ Required-Yes (1)—The system disconnects the call if the private routing table is not found.

In RADIUS, private route tables are defined in a pseudo-user profile. A private-route profile is a pseudo-user profile in which the first two lines have the following format:

`profile-name` `Password = "ascend" Service-Type = Outbound`

The *`profile-name`* value is any name you assign to the profile. Private-route profile definitions can include the following VSAs:

RADIUS attribute	Value
Ascend-Private-Route (104)	Destination address and next-hop router address for a private route. Each private-route profile specifies one or more private routes with this attribute, which is more fully described in the *TAOS RADIUS Guide and Reference*.

RADIUS attribute	Value
Ascend-Cache-Refresh (56)	Whether the timer for cached routes in this profile is reset each time a new session that refers to the pseudo-user profile becomes active. Refresh-No (0) does not reset the timer. Refresh-Yes (1) resets the cache timer when a session referring to the profile becomes active.
Ascend-Cache-Time (57)	Number of minutes to cache the profile. Once the cache timer expires for a RADIUS profile, the profile is deleted from system memory. The next time it is needed, the system retrieves it from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of route lookups, at the cost of some system memory. The minimum possible cache time is 0 minutes, which causes the system to retrieve the profile for every route lookup in the table. This value is usually not desirable. If no value is specified for this attribute, the setting for the `default-prt-cache-time` parameter in the `ip-global` profile is used.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the Stinger unit must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius
```

```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *Stinger Reference*.

Typical configuration of a private routing table

You can configure private routing tables locally or in RADIUS. For example, the following commands define a private routing table named `check`:

```
admin> new private-route-table check
PRIVATE-ROUTE-TABLE/check read
```

```
admin> list route-description-list 1
[in PRIVATE-ROUTE-TABLE/check:route-description-list[1] (new)]
enabled = no
dest-address = 0.0.0.0/0
netmask = 0.0.0.0
gateway-address = 0.0.0.0
metric = 0
```

```
admin> set enabled = yes
```

```
admin> set dest-address = 1.1.1.1/24
```

```
admin> set gateway-address = 2.2.2.2
```

```
admin> set metric = 2
```

```
admin> list
[in PRIVATE-ROUTE-TABLE/check:route-description-list[1]]
enabled = yes
dest-address = 1.1.1.1/24
netmask = 255.255.255.0
gateway-address = 2.2.2.2
metric = 2
admin> list .. 2
[in PRIVATE-ROUTE-TABLE/check:route-description-list[1]]
enabled = no
dest-address = 0.0.0.0/0
netmask = 0.0.0.0
gateway-address = 0.0.0.0
metric = 0
admin> set enabled = yes
admin> set dest-address = 3.3.3.3/28
admin> set gateway-address = 2.2.2.2
admin> set metric = 3
admin> write
PRIVATE-ROUTE-TABLE/check written
```

Following is a comparable RADIUS private-route profile:

```
check Password = "ascend", Service-Type = Outbound
    Ascend-Cache-Time = 3,
    Ascend-Cache-Refresh = Refresh-Yes,
    Ascend-Private-Route = "1.1.1.1/24 2.2.2.2 2",
    Ascend-Private-Route = "3.3.3.3/28 2.2.2.2 3"
```

The following commands configure the default cache time for RADIUS private-route profiles:

```
admin> read ip-global
IP-GLOBAL read
admin> set default-prt-cache-time = 180
admin> write
IP-GLOBAL written
```

Following is a sample RADIUS private-route profile that uses of the default instead of specifying a value for Ascend-Cache-Time (57):

```
my-routes Password = "ascend"
    Service-Type = Outbound,
    Ascend-Private-Route = "1.1.1.1/24 2.2.2.2",
    Ascend-Private-Route = "3.3.3.3/28 2.2.2.2"
```

Examples of using private routing tables

The following commands modify a `connection` profile so that the session has access to the routes in the private routing table, and the system disconnects the call if the private routing table is not found:

```
admin> read connection p50-v2
CONNECTION/p50-v2 read
```

```
admin> set ip-options private-route-table = check

admin> set ip-options private-route-profile-required = yes

admin> write
CONNECTION/p50-v2 written
```

The following RADIUS profile refers to the same private routing table and has the same requirements. This profile also specifies how the routes are cached for this connection.

```
p50-v2 Password = "my-password"
    Service-Type = Framed,
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.1.1.1,
    Framed-IP-Netmask = 255.0.0.0,
    Ascend-Private-Route-Table-ID = "check",
    Ascend-Private-Route-Required = Required-Yes
```

The following commands configure the system to reject incoming calls when the RADIUS user profile specifies a private routing table that is not found:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set ip-answer private-route-profile-required = yes

admin> write
ANSWER-DEFAULTS written
```

Following is a sample RADIUS profile that uses of the default instead of specifying a value for Ascend-Private-Route-Required (55):

```
p50-v2 Password = "my-password"
    Service-Type = Framed,
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.1.1.1,
    Framed-IP-Netmask = 255.0.0.0,
    Ascend-Private-Route-Table-ID = "check"
```

Examples of private static routes

A RADIUS user profile can specify a list of private routes associated with the connection. (There is no comparable functionality in local `connection` profiles.)

Private routes defined by the Ascend-Private-Route attribute in a user profile affect only packets received from the connection. The routes are not added to the global routing table. If a destination is not found in the list of private routes and there is no default private route, the global routing table is consulted for a decision on routing the packets. Otherwise, only the private routing table is consulted.

Following is an example of a user profile that creates three private routes associated with the user:

```
pipe50 Password = "ascend" User-Service = Framed
   Framed-Protocol = PPP,
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.0.0.0,
   Ascend-Private-Route = "170.1.0.0/16 10.10.10.1",
   Ascend-Private-Route = "200.1.1.1/32 10.10.10.2",
```

```
   Ascend-Private-Route = "20.1.0.0/16 10.10.10.3",
   Ascend-Private-Route = "0.0.0.0/0 10.10.10.4"
```

With this profile, the private routing table for the connection contains the following routes, the last one of which is the default route:

```
Dest/Mask        Gateway
170.1.0.0/16     10.10.10.1
200.1.1.1/32     10.10.10.2
20.1.0.0/16      10.10.10.3
0.0.0.0/0        10.10.10.4
```

Note The user profile can also specify a value for Ascend-Client-Gateway attribute, but the value will *not* modify a private default route that has been specified by the Ascend-Private-Route attribute.

When the next-hop router address specified by an Ascend-Private-Route attribute is the zero address (0.0.0.0), a lookup is performed for that route in the global routing table, providing an exit mechanism to the global table for specific private routes. For example, suppose the private routes are defined as in the following RADIUS user profile:

```
pipe50 Password = "ascend" User-Service = Framed
   Framed-Protocol = PPP,
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.0.0.0,
   Ascend-Private-Route = "170.1.0.0/16 10.10.10.1 1",
   Ascend-Private-Route = "200.1.1.1/32 10.10.10.2",
   Ascend-Private-Route = "20.1.0.0/16  0.0.0.0 1",
   Ascend-Private-Route = "0.0.0.0/0  10.10.10.4 1"
```

The private routing table for this connection contains the following routes:

```
Dest/Mask        Gateway
170.1.0.0/16     10.10.10.1
200.1.1.1/32     10.10.10.2
20.1.0.0/16      0.0.0.0
0.0.0.0/0        10.10.10.4
```

With this private routing table, a route lookup for the 20.1.0.0/16 network proceeds to the global routing table.

Setting TCP/IP routing policies

The TAOS router has many configuration settings that affect its operations. The settings that determine its routing policies include security, RIP options, IP route cache options, and other options. These settings are available only in the `ip-global` profile. They have no counterpart in RADIUS.

Note You can also configure the system to set QoS priority bits and TOS classes of service on behalf of customer applications. These settings can then be used by other routers to prioritize and select links for particular data streams. These policies are set on WAN interfaces. For details, see “Typical QoS and TOS policy configuration” on page 1-22.

Setting a system source IP address

The system IP address is the source address used for all packets generated by the system. For example, this address is used for RADIUS requests, ATMP tunnel requests, or a `telnet`, `traceroute`, or `ping` command originating from the unit. It must be the real address of one of the unit's LAN IP interfaces, or the interface-independent address described in "Typical definition of the soft interface" on page 1-9.

Following is the parameter for specifying a system address:

```
[in IP-GLOBAL]
system-ip-addr = 0.0.0.0
```

With the default zero address, the Stinger unit uses the IP address assigned to the control-module Ethernet interface as the source address for packets it generates. One reason for setting a system address other than the default is that doing so simplifies access control. For example, most RADIUS servers keep a database of known remote access server (RAS) clients and their authentication keys. If you do not specify a system address, the RADIUS database must include a complete list of all the system's interface addresses. If you specify a system address, it is used for all RADIUS request packets.

Another reason for setting a system address is to ensure that packets sent from an ATMP home agent to foreign agents have a single, standard source address. A system address is recommended for ATMP home agents that have multiple interfaces into the IP cloud that separates them from foreign agents, to prevent communication problems if a route changes within the IP cloud. For details, see "System IP address recommendation" on page 6-2.

Following is an example of setting the `system-ip-addr` parameter to an address assigned to a port on a T1000 module in slot 3:

```
admin> dir ip-interface
    6  09/14/2001  10:13:24  { { any-shelf any-slot 0 } 0 }
    8  09/14/2001  10:13:24  { { shelf-1 first-control-module 1 } 0 }
   19  09/14/2001  10:14:02  { { shelf-1 second-control-module 1 } 0 }
    8  09/14/2001  11:36:32  { { shelf-1 slot-3 2 } 0 }
   64  09/14/2001  11:53:12  { { shelf-1 slot-3 1 } 0 }

admin> get ip-interface { { 1 3 1 } 0 } ip-address
ip-address = 10.2.3.4

admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 10.2.3.4

admin> write
IP-GLOBAL written
```

Setting router security policies

The following parameters (shown with default settings) affect router security:

```
[in IP-GLOBAL]
must-accept-address-assign = no
shared-prof = no
```

```
telnet-password = ""
user-profile = ""
```

Parameter	Setting
must-accept-address-assign	Enable/disable rejection of an assigned IP address by an incoming caller during PPP negotiation.
shared-prof	Enable/disable multiple callers sharing a single connection profile.
telnet-password	Password required for telnet access to the Stinger unit.
user-profile	Name of a default user profile for telnet sessions.

Requiring acceptance of dynamic address assignment

During PPP negotiation, a calling station can reject an IP address offered by the router and present the caller's own IP address for consideration. For security purposes, many sites set must-accept-address-assign to yes to ensure that the Stinger unit terminates such a call, as shown in the following example:

```
admin> read ip-global
IP-GLOBAL read

admin> set must-accept-address-assign = yes

admin> write
IP-GLOBAL written
```

For address assignment to occur, the Stinger unit must have an address available for assignment, the answer-defaults profile must enable dynamic assignment, the caller's profile must specify dynamic assignment, and the caller's PPP software must be configured to acquire its IP address dynamically. For details, see "Examples of assigning an address from a pool" on page 1-65.

Shared profiles

In low-security situations, you can allow multiple callers to share a single connection profile. Sharing profiles is recommended only for *low-security* networks and requires dynamic address assignment to the individual callers. When the shared profile uses dynamic address assignment, each call is a separate connection that shares the same name and password, but a separate IP address is assigned dynamically to each caller. Because the password is shared by multiple callers, this application is unsuitable for high-security networks.

Note If you do enable shared profiles, the profile must not result in a shared IP address (two callers at different locations cannot share the same address). The profile must either not assign an address or must assign an IP address dynamically. For details about dynamic IP address assignment, see "Examples of assigning an address from a pool" on page 1-65.

You can enable shared profiles on one of two possible levels: globally in the ip-global profile, or per-connection in a connection profile. Once you enable shared profiles globally, you cannot disable it for an individual connection. The following commands enable shared profiles globally:

```
admin> read ip-global
IP-GLOBAL read
```

```
admin> set shared-prof = yes

admin> write
IP-GLOBAL written
```

With this setting, the shared-prof parameter in an individual connection profile has no effect.

However, if you shared profiles are disabled globally, you can allow sharing of a specific connection by enabling the shared-prof parameter in the connection profile. This functionality is also available in RADIUS profiles via the Ascend-Shared-Profile-Enable attribute. Use the following parameters (shown with its default setting) to enable or disable shared profiles on a per user basis and specify the maximum number of users that can be connected at the same time while sharing the profile:

```
[in CONNECTION/""]
shared-prof = no
max-shared-users = 0
```

Parameter	Setting
shared-prof	Enable/disable multiple callers to share the connection profile, provided that IP address conflicts do not result and the ip-global profile sets shared-prof to no. If the ip-global profile sets shared-prof to yes, this setting has no effect.
max-shared-users	Maximum number of users that can be simultaneously connected using a shared profile. The default zero value indicates no limit on the number of users sharing a profile at the same time.

For example, with the following settings, only the connection profile named shared-1 can be shared by multiple callers:

```
admin> get ip-global shared-prof
[in IP-GLOBAL:shared-prof]
shared-prof = no

admin> read connection shared-1
CONNECTION/shared-1 read

admin> set shared-prof = yes

admin> write
CONNECTION/shared-1 written
```

Specifying a default user profile for telnet access

RADIUS use the following attribute-value pair to specify a default user profile for RADIUS-authenticated telnet access to the Stinger unit:

RADIUS attribute	Value
Ascend-Telnet-Profile (91)	Name of a local user profile to be used for authenticating telnet logins.

When a user attempts to telnet into the system's interface, the system first looks for a user profile matching the login name and password given by the user. If that fails, the system uses the server specified in the external-auth profile to locate a RADIUS user profile. If the RADIUS server returns a profile that includes the Ascend-Telnet-Profile attribute, the system uses the specified user profile to authenticate and set permissions for the session. Only RADIUS profiles that specify a value for this attribute can be used to authenticate a telnet login to the command interface. Following is a sample RADIUS profile that enables telnet access to the Stinger unit with administrator permissions:

```
admin Password = "secret-pw"
    Service-Type = Framed-User,
    Ascend-Telnet-Profile = admin
```

Restricting telnet access to the system

A user can initiate a telnet session to the command line from a local workstation or from a WAN connection. In both cases, the system authenticates the session by means of a user profile, which defines a permission level for the user logging in. (For details about user profiles, see the *Stinger Reference*.)

In addition to the password required by a user profile, you can specify that telnet requires its own password authentication, which occurs before any user profile authentication.

The commands in the following example set the telnet-password parameter and specify the Default user profile for telnet logins. The default profile enables minimal permissions and requires no password.

```
admin> read ip-global
IP-GLOBAL read

admin> set telnet-password = !234#@

admin> set user-profile = default

admin> write
IP-GLOBAL written
```

When users telnet to the system, they are allowed three tries, each with a 60-second time limit, to enter the correct telnet password. If all three attempts fail, the connection times out. If they specify the correct telnet password, the system prompts again for a username and password to authenticate a user profile. In the following example, a user starts a telnet session to a Stinger unit named taos01, for which a telnet password has been specified.

```
% telnet taos01
<taos01> Enter Password:

Trying 10.1.2.3 ...
Connected to taos01.abc.com.
Escape character is '^]'.
User:
```

After entering the correct telnet password, the user is prompted for a username and password to authenticate a user profile.

Setting systemwide routing policies

The following parameters, (shown with default settings) specify system-wide routing policies:

```
[in IP-GLOBAL]
ignore-icmp-redirects = no
icmp-reply-directed-bcast = no
drop-source-routed-ip-packets = no
static-pref = 100
```

Parameter	Setting
ignore-icmp-redirects	Enable/disable processing of ICMP Redirect packets.
icmp-reply-directed-bcast	Enable/disable responding as a host to directed-broadcast ICMP Echo Requests.
drop-source-routed-ip-packets	Enable/disable forwarding of IP packets that have the source route option set.
static-pref	Default preference given to static IP routes.

Ignoring ICMP packets

ICMP Redirect packets can be counterfeited and used to change the way a device routes packets. For security purposes, many sites choose to ignore ICMP Redirects.

ICMP Echo Requests to the broadcast address have been used in denial-of-service attacks. To prevent the TAOS router from being used in a denial-of-service attack when an attacker compromises another router on the same Ethernet network as the Stinger unit, you can prevent the Stinger unit from responding to directed-broadcast ICMP Echo Requests sent to the IP broadcast address.

The following commands configure the unit to ignore both types of ICMP packets. (By default, it does not respond to ICMP Echo Requests to the broadcast address.)

```
admin> read ip-global
IP-GLOBAL read
admin> set ignore-icmp-redirects = yes
admin> write
IP-GLOBAL written
```

Dropping source-routed packets

The default setting for the drop-source-routed-ip-packets parameter is no, which causes the router to forward all source-routed packets as described in RFC 1812. When the parameter is set to yes, the router drops all packets that have either a Loose or a Strict source route among their IP options. The following set of commands instructs the router to drop source-routed packets:

```
admin> read ip-global
IP-GLOBAL read
admin> set drop-source-routed-ip-packets = yes
admin> write
IP-GLOBAL written
```

Setting routing protocol options

The following parameters (shown with default settings) define how the system handles routing protocol updates:

```
[in IP-GLOBAL]
rip-policy = Poison-Rvrs
summarize-rip-routes = no
rip-trigger = yes
rip-pref = 100
dialout-poison = no
rip-queue-depth = 0
ignore-def-route = yes
suppress-host-routes = no
ospf-pref = 10
ospf-ase-pref = 150
rip-tag = c8:00:00:00
rip-ase-type = 1
```

```
[in IP-GLOBAL:ospf-global]
as-boundary-router = yes
```

Parameter	Setting
rip-policy	Policy for sending update packets that include routes received on the same interface.
summarize-rip-routes	Enable/disable summarization of subnet information in RIP-v1 updates. This setting has no effect on RIP-2 updates.
rip-trigger	Enable/disable RIP triggering. With a yes setting (the default), RIP updates include only changed routes.
rip-pref	Default preference for routes that the Stinger unit learns from RIP updates. When choosing the routes to put in the routing table, the unit first compares their preference values, preferring the lowest number. If the preference values are equal, the router compares the metric values, using the route with the lowest metric. Specify a number from 0 through 255. A value of 255 prevents the use of the route. Following are the default preferences for different types of routes: 0 (zero)—Connected routes 10—Open Shortest Path First (OSPF) routes 30—Routes learned from Internet Control Message Protocol (ICMP) redirects 100—Routes learned from RIP 100—Static routes 100—Ascend Tunnel Management Protocol (ATMP) routes

Parameter	Setting
dialout-poison	Enable/disable advertisement of dial-out routes when no trunks are available. Stinger units do not dial out, so leave this parameter at its default setting.
ignore-def-route	Enable/disable exclusion of advertised default routes from the routing table.
rip-queue-depth	Maximum number of RIP packets to be held for processing. Valid values are 0 to 1024. The default (0) means that the router will not drop any RIP packets, no matter how far behind it gets.
suppress-host-routes	Enable/disable suppression of host routes for interfaces with a subnet mask of less than 32 bits.
ospf-pref	OSPF option (see "Setting routing protocol options" on page 1-41).
ospf-ase-pref	OSPF option (see "Setting routing protocol options" on page 1-41).
rip-tag	OSPF option (see "Setting routing protocol options" on page 1-41).
rip-ase-type	OSPF option (see "Setting routing protocol options" on page 1-41).
ospf-global:as-boundary-router	OSPF option (see "Setting routing protocol options" on page 1-41).

RIP policy for propagating updates back to the originating subnet

You can specify a split-horizon or poison-reverse policy for outgoing update packets that include routes received on the same interface on which the update is sent. Split-horizon means that the router does not propagate routes back to the subnet from which they were received. Poison-reverse means that it propagates routes back to the subnet from which they were received, but with a metric of 16 (infinite metric).

The following set of commands specifies the split-horizon policy:

```
admin> read ip-global
IP-GLOBAL read

admin> set rip-policy = split

admin> write
IP-GLOBAL written
```

RIP triggering

RIP triggering enables the router to tag routes that have been updated in the routing table and send updates that include only the changed routes. The result is reduced processing overhead for both the TAOS router and its neighbors.

With the default value (yes), the router tags changes to its routing table and includes only the tagged routes in its next update. Changes occur when a call arrives or disconnects, RIP learns a route from another router, or the administrator modifies a route-related profile. The router broadcasts updates 5 to 8 seconds after the first

change in the routing table is detected. The delay helps to prevent constant updates during peak traffic conditions.

If `rip-trigger` is set to `no`, the router sends full table updates every 20 to 40 seconds. To prevent RIP routers on a network from synchronizing and sending large updates in unison, the full table update is no longer broadcast at fixed 30-second intervals.

Limiting the size of UDP packet queues

When the router is very busy and receives a flood of UDP packets from SNMP requests or RIP updates, a backlog of packets waiting for processing can create enough delay in routing to cause sporadic problems with time-sensitive packets, such as LCP negotiation or frame relay management packets.

To prevent such problems, UDP processing runs at a lower priority than processing of routed packets. On a system busily routing packets, UDP processing might be delayed, and a backlog of UDP packets builds up. The `rip-queue-depth` parameter in the `ip-global` profile and the `queue-depth` parameter in the `snmp` profile specify the maximum size of this backlog.

When you set one of these parameters to specify a queue depth, the system is more likely to drop UDP packets when it is busy routing packets. However, time-sensitive routed packets are less likely to be delayed and system memory is used more efficiently.

In following sample commands sets both queue depths to 50. Fifty of each type of packet is held for processing, and if additional packets of either type are received when the queue is full, they are dropped.

```
admin> read ip-global
IP-GLOBAL read

admin> set rip-queue-depth = 50

admin> write
IP-GLOBAL written

admin> read snmp
SNMP read

admin> set queue-depth = 50

admin> write
SNMP written
```

The `netstat` command output shows the queue depth of various UDP ports, and the total packets received and total packets dropped on each port. The total packets received count includes dropped packets. In the following example, the SNMP queue depth was set to 32:

```
admin> netstat udp
udp:
Socket  Local Port  InQLen  InQMax    InQDrops    Total Rx
  0        1023       0       1           0           0
  1       route       0      50           0         509
  2        echo       0      32           0           0
  3         ntp       0      32           0           0
  4        1022       0     128           0           0
  5        SNMP      32      32        5837       20849
```

Ignoring default routes when updating the routing table

Lucent Technologies recommends enabling the ignore-def-route parameter to prevent routing updates from modifying the default route in the routing table. The following set of commands protects the default route from RIP updates:

```
admin> read ip-global
IP-GLOBAL read

admin> set ignore-def-route = yes

admin> write
IP-GLOBAL written
```

Suppressing host-route advertisements

If you set the suppress-host-routes parameter to yes, routes are suppressed according to the following rules:

- If a connection profile includes a subnet mask of less than 32 bits in the remote-address setting, host routes for the interface are suppressed while the session is being negotiated, and after the session is established, only network routes are advertised for the interface.
- If a connection profile includes a subnet mask of /32 in the remote-address setting, host routes for the interface are not suppressed. (Pool addresses also have a 32-bit mask, so they are not suppressed.)

The following set of commands configures the router to suppress host routes for connections that specify a subnet mask of less than 32 bits:

```
admin> read ip-global
IP-GLOBAL read

admin> set suppress-host-routes = yes

admin> write
IP-GLOBAL written
```

Enabling protocol options

The following parameters (shown with default settings) configure TCP/IP protocol options:

```
[in IP-GLOBAL]
bootp-enabled = no
rarp-enabled = no
udp-cksum = yes
tcp-timeout = 0
finger = no

[in IP-GLOBAL:bootp-relay]
active = no

[in IP-GLOBAL:bootp-relay:bootp-servers]
bootp-servers[1] = 0.0.0.0
bootp-servers[2] = 0.0.0.0

[in IP-GLOBAL:sntp-info]
enabled = no
gmt-offset = utc+0000
host = [0.0.0.0 0.0.0.0 0.0.0.0]
```

```
[in IP-GLOBAL:sntp-info:host]
host[1] = 0.0.0.0
host[2] = 0.0.0.0
host[3] = 0.0.0.0
```

Parameter	Setting
bootp-enabled	Enable/disable querying a BOOTP server.
rarp-enabled	Enable/disable obtaining the system's IP addresses from a RARP server.
udp-cksum	Enable/disable UDP checksums.
tcp-timeout	Interval for TCP retry attempts. Valid values are from 0 to 200 seconds.
finger	Enable/disable response to remote Finger queries. When Finger is set to No (the default), the Stinger unit rejects queries from Finger clients and sends a message that the Finger online user list is denied.
bootp-relay:active	Enable/disable BOOTP Relay.
bootp-relay:bootp-servers[1]	IP address of up to two BOOTP servers. Only one address is required.
bootp-relay:bootp-servers[2]	
sntp-info:enabled	Enable/disable the Simple Network Time Protocol (SNTP).
sntp-info:gmt-offset	Current time zone as an offset from the Universal Time Configuration (UTC). UTC is in the same time zone as Greenwich Mean Time (GMT).
sntp-info:host[1] sntp-info:host[2] sntp-info:host[3]	IP addresses for up to three SNTP servers. Only one address is required.

Enabling boot protocol and reverse ARP

The Boot Protocol (BOOTP) is a UDP/IP-based protocol that enables a host to obtain its configuration dynamically from a BOOTP server. Reverse ARP (RARP) enables a host to obtain its address from a RARP server. The following commands enable both BOOTP and RARP:

```
admin> read ip-global
IP-GLOBAL read

admin> set bootp-enabled = yes

admin> set rarp-enabled = yes

admin> write
IP-GLOBAL written
```

Enabling UDP checksums

If data integrity is of the highest concern for your network, and redundant checks are important, you can turn on UDP checksums to generate a checksum whenever a UDP

packet is transmitted. UDP packets are transmitted for queries and responses related to ATMP, SYSLOG, DNS, ECHOSERV, RADIUS, TACACS, RIP, SNTP, and TFTP.

The following commands enable UDP checksums for transmitted packets:

```
admin> read ip-global
IP-GLOBAL read

admin> set udp-cksum = yes

admin> write
IP-GLOBAL written
```

Setting a TCP timeout

The tcp-timeout parameter adjusts the TCP retry timer. At the default value (0), the system attempts a fixed number of retries at escalating intervals adding up to about 170 seconds total. (Other limits in the system terminate TCP retries after about 170 seconds, even if the parameter is set to a higher number.) If you set TCP-Timeout to a nonzero value, the value specifies the number of seconds TCP retries persist. After the specified number of seconds, the retries stop and the connection is considered lost.

The tcp-timeout setting applies to all TCP connections initiated from the Stinger unit, including telnet, rlogin, and the TCP portion of DNS queries. The parameter applies to both established TCP connections and initial attempts to connect. A situation in which you might adjust the TCP retry timer would be, for example, when a user employs client software to enter a hostname in a terminal-server session, and DNS returns a list of IP addresses for the host. If the first address proves unreachable and the timeout on each attempt is long, the client software often times out before finding a good address.

The following commands set the timeout to 50 seconds:

```
admin> read ip-global
IP-GLOBAL read

admin> set tcp-timeout = 50

admin> write
IP-GLOBAL written
```

The optimal setting for the tcp-timeout parameter depends on the characteristics of the TCP destination (server) hosts, and therefore must be based on experience. For example, if the destinations are all on a LAN under the same administrative control as the Stinger unit and are lightly loaded, a short timeout (such as a few seconds) might be reasonable, because a host that does not respond within that interval is probably down. Conversely, if the environment includes servers with longer network latency times (for example, those connected across the WAN), or load is high in the network or the router, or the characteristics of the remote hosts are not well-known, a longer timeout is appropriate. Values of 30 to 60 seconds are common in UNIX TCP implementations.

Enabling response to finger queries

If finger (described in RFC 1288) is enabled in the ip-global profile, the system can return user information to a remote finger query. The following commands enable the Stinger unit to accept finger queries and return the requested active session details to a remote client:

```
admin> read ip-global
IP-GLOBAL read

admin> set finger = yes

admin> write
IP-GLOBAL written
```

When the `finger` parameter is set to `yes`, a client (such as a UNIX client) can request session information for the system named TAOS1 by entering the following command:

```
# finger @taos1
```

The above command displays the information in narrow (80-character-wide) format. The client can request the information in wide format by using the command with the `-l` option. For example, the following command displays a wide (140-character-wide) format of session information for the system named TAOS1:

```
# finger -l @taos1
```

The client can also request the details of all sessions or of a single session. For example, the following command would request information about a single user named Gavin:

```
# finger gavin@taos1
```

The `finger` forwarding service is not supported. It uses the following hostname format:

@host1@host2

A remote client that uses the forwarding request format receives the following message:

```
Finger forwarding service denied.
```

Enabling BOOTP-Relay

If a host requesting an address does not reside on the same IP network as a BOOTP server, an intervening system is required to transfer messages between the client and server. The intervening host is a BOOTP Relay Agent.

The following commands enable the BOOTP Relay feature and specify the address of a BOOTP server:

```
admin> read ip-global
IP-GLOBAL read

admin> list bootp-relay
[in IP-GLOBAL:bootp-relay]
active = no
bootp-servers = [ 0.0.0.0 0.0.0.0 ]

admin> set active = yes

admin> set bootp-servers 1 = 10.178.10.125

admin> write
IP-GLOBAL written
```

If more than one server is specified, the Stinger unit uses the first server until it becomes unavailable. Once the unit starts using the second server, the unit continues

using that server until it becomes unavailable, at which time the unit switches back to using the first server again.

Using SNTP to set and maintain the system time

The Stinger unit can use Simple Network Time Protocol (SNTP), which is described in RFC 1305, to set and maintain its system time by communicating with an SNTP server.

You specify the system's time zone as an offset from the Universal Time Configuration (UTC). UTC is in the same time zone as Greenwich Mean Time (GMT). The offset specifies hours and minutes from UTC, using a 24-hour clock. Because some time zones, such as Newfoundland, do not have an even hour boundary, the offset includes four digits and requires half-hour increments.

For example, in Newfoundland the time is 1.5 hours earlier than UTC, so the offset is UTC-0130. For San Francisco, which is 8 hours earlier than UTC, the offset is UTC -0800. For Frankfurt, which is 1 hour later than UTC, the offset is UTC +0100.

The commands in the following example specify the time zone for San Francisco and the address of one SNTP server:

```
admin> read ip-global
IP-GLOBAL read
admin> list sntp-info
enabled = no
gmt-offset = utc+0000
host = [0.0.0.0 0.0.0.0 0.0.0.0]
admin> set enabled = yes
admin> set gmt = utc-0800
admin> set host 1 = 10.2.3.4
admin> write
IP-GLOBAL written
```

The system always communicates with the first address unless it is inaccessible. In that case, the unit attempts to communicate with the second address, trying the third address only if the other two are inaccessible.

Configuring port redirection

Port redirection enables you to configure a `connection` or RADIUS profile to redirect certain packet types to a specified server. For example, you could redirect Hypertext Transfer Protocol (HTTP) traffic to a Web cache server on a local network. However, port redirection is not limited to HTTP traffic. You can use the feature to redirect any TCP or UDP packet on the basis of its protocol and port information.

Overview of Connection profile settings

To configure port redirection in a `connection` profile, set the following parameters (shown with default settings):

```
[in CONNECTION/"":port-redirect-options]
protocol = none
port-number = 0
redirect-address = 0.0.0.0
```

Parameter	Setting
`protocol`	Protocol type. Valid settings are `none` (the default, which disables port redirection), `udp`, and `tcp`. The specified setting, together with the `port-number` setting, defines a type of packet. For example, `tcp` with `21` represents FTP traffic, and `tcp` with `23` represents Telnet traffic. For HTTP traffic, set the parameter to `tcp`.
`port-number`	Port number to be compared with the destination port of a packet. TCP and UDP port numbers are typically assigned to services. For example, HTTP traffic uses TCP port 80. For a list of assigned port numbers, see RFC 1700.
`redirect-address`	IP address to which matching packets are redirected.

Overview of RADIUS settings

RADIUS uses the following attribute-value pairs for port redirection:

RADIUS attribute	Value
Ascend-Port-Redir-Protocol (82)	Protocol type. Valid values are `Ascend-Proto-TCP (6)` and `Ascend-Proto-UDP (17)`. The specified value, together with the Ascend-Port-Redir-Portnum value defines a type of packet. For example, `Ascend-Proto-TCP` with 21 represents FTP traffic, and `Ascend-Proto-TCP` with 23 represents Telnet traffic. For HTTP traffic, specify `Ascend-Proto-TCP (6)`.
Ascend-Port-Redir-Portnum (83)	Port number to be compared with the destination port of a packet. TCP and UDP port numbers are typically assigned to services. For example, HTTP traffic uses TCP port 80. For a list of assigned port numbers, see RFC 1700, *Assigned Numbers*.
Ascend-Port-Redir-Server (84)	IP address to which matching packets are redirected.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the system must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius
```

```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *Stinger Reference*.

Example of configuring port redirection

In this example, a Stinger unit redirects a PPP client's browser requests to a Web cache server at 1.1.1.1. The Web cache server can respond directly if a cached entry is

found, or forward the browser request to its original destination if no cache entry is found. The sample setup is shown in Figure 1-8.

Figure 1-8. Port redirection to an HTTP server

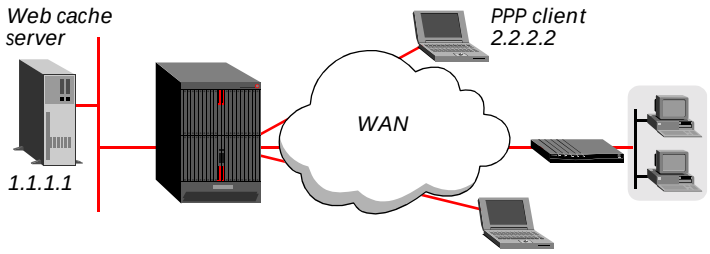

The following commands configure a local profile for the PPP client, redirecting its HTTP traffic to the server at 1.1.1.1:

```
admin> new connection atcp50
CONNECTION/atcp50 read
admin> set active = yes
admin> set ip-options remote-address = 2.2.2.2/32
admin> set ppp-options recv-password = test
admin> set port-redirect-options protocol = tcp
admin> set port-redirect-options port-number = 80
admin> set port-redirect-options redirect-address = 1.1.1.1
admin> write
CONNECTION/atcp50 written
```

Following is a comparable RADIUS profile:

```
atcp50 Password = "test"
   Service-Type = Framed,
   Framed-Protocol = MPP,
   Framed-IP-Address = 2.2.2.2,
   Framed-IP-Netmask = 255.255.255.255,
   Ascend-Port-Redir-Protocol = Ascend-Proto-TCP,
   Ascend-Port-Redir-Portnum = 80,
   Ascend-Port-Redir-Server = 1.1.1.1
```

Configuring DNS

Domain Name System (DNS) is a TCP/IP service for centralized management of address resolution. Service providers can maintain multiple DNS servers, each one dedicated to a particular client or location. In that case, for security reasons you might need to ensure that connections are always directed to the correct DNS service. With per-connection DNS access, a service provider can direct specific users to the DNS servers appropriate to their services or locations.

In the Stinger unit, DNS configuration includes settings for enabling local DNS lookups and supporting a DNS list, settings for a local DNS table maintained in RAM, and client DNS for directing connections to a particular DNS service.

Configuring DNS lookups and a DNS list

You enable DNS lookups by specifying a domain name and the IP addresses of local servers. Some DNS servers also support a list feature that enables them to return multiple addresses for a hostname in response to a DNS query.

Overview of DNS settings

Following are the parameters (shown with default settings) for configuring DNS to allow lookups and support a DNS list:

```
[in IP-GLOBAL]
domain-name = ""
dns-primary-server = 0.0.0.0
dns-secondary-server = 0.0.0.0
netbios-primary-ns = 0.0.0.0
netbios-secondary-ns = 0.0.0.0
dns-list-attempt = no
dns-list-size = 6
sec-domain-name = ""
```

Parameter	Setting
`domain-name`	Primary domain name to use for DNS lookups. The system appends this domain name to hostnames when performing lookups.
`dns-primary-server`	Address of the primary local DNS server to use for lookups.
`dns-secondary-server`	Address of the secondary local DNS server to use for lookups. Used only if the primary server is not found.
`netbios-primary-ns` `netbios-secondary-ns`	Addresses of a primary and secondary NetBIOS server.
`dns-list-attempt`	Enable/disable a DNS list.
`dns-list-size`	Maximum number of hosts in a DNS list, up to 35.
`sec-domain-name`	Secondary domain name to use for DNS lookups if the hostname is not found in the primary domain.

Specifying domain names for lookups

When the system receives a hostname to look up, it tries various combinations, including appending the domain name specified in the `ip-global` profile. The following commands specify a primary and secondary domain name for DNS lookups:

```
admin> read ip-global
IP-GLOBAL read

admin> set domain-name = abc.com

admin> set sec-domain-name = eng.abc.com

admin> write
IP-GLOBAL written
```

If a lookup fails with the first domain name, the router tries again with the secondary domain name.

Specifying local DNS server addresses

To enable the system to look up addresses via DNS, specify DNS server addresses as shown in the following example:

```
admin> read ip-global
IP-GLOBAL read

admin> set dns-pri = 10.2.3.56

admin> set dns-sec = 10.2.3.107

admin> write
IP-GLOBAL written
```

If the primary server is unavailable, the system attempts a lookup on the secondary server. To execute a lookup manually, use the `nslookup` command. For example:

```
admin> nslookup techpubs
Resolving host techpubs.
IP address for host techpubs is 10.6.212.19.
```

Local DNS servers provide information about the local network, and are sometimes isolated from incoming callers for security purposes. For details, see “Using client DNS” on page 1-56.

Supporting a DNS list

Some DNS servers support a list feature that enables them to return multiple addresses for a hostname in response to a DNS query. However, the responses do not include information about availability of the hosts in the list. Users typically attempt to access the first address in the list. If that host is unavailable, the user must try the next host, and so forth.

When the DNS list is used for an immediate connection by a client (for example, an immediate `telnet` connection to a local host), and the first attempt fails, the physical connection is torn down. To avoid tearing down and then reestablishing the connection before attempting to access the next host in the list, enable the DNS list feature. The following example shows how to enable a DNS list with a maximum of 14 hosts in the list:

```
admin> read ip-global
IP-GLOBAL read

admin> set dns-list-attempt = yes

admin> set dns-list-size = 14

admin> write
IP-GLOBAL written
```

For related information, see “Using the auto-update feature” on page 1-55.

Setting up a local DNS table

The Stinger unit can maintain in RAM a DNS table of up to eight hostnames and their IP addresses. It consults the table in RAM for address resolution only if requests to the

DNS server fail. The local table acts as a safeguard to ensure that the system can resolve the local set of DNS names even if all DNS servers become unreachable or fail.

The local DNS table is propagated to RAM from a configured `dns-local-table` subprofile in the `ip-global` profile. At startup, the system copies values in the profile to the table in RAM. If you subsequently modify the `dns-local-table` subprofile, the changes are propagated to the table in RAM when the profile is written.

The DNS table in RAM has space for up to 35 IP addresses per `host-name` entry (35 is the maximum setting for `dns-list-size`). The `dns-local-table` subprofile allows a single IP address per hostname. (For related information, see "Using the auto-update feature" on page 1-55.)

Overview of DNS table settings

To set up the local DNS table, configure the following parameters (shown with their default values) in the `ip-global` profile:

```
[in IP-GLOBAL:dns-local-table]
enabled = no
auto-update = no

[in IP-GLOBAL:dns-local-table:table-config]
table-config [1] = { "" 0.0.0.0 }
table-config [2] = { "" 0.0.0.0 }
table-config [3] = { "" 0.0.0.0 }
table-config [4] = { "" 0.0.0.0 }
table-config [5] = { "" 0.0.0.0 }
table-config [6] = { "" 0.0.0.0 }
table-config [7] = { "" 0.0.0.0 }
table-config [8] = { "" 0.0.0.0 }

[in IP-GLOBAL:dns-local-table:table-config[1]]
host-name = ""
ip-address = 0.0.0.0
```

DNS-Local Table parameter	Setting
`enabled`	Whether the local DNS table in RAM will be available if DNS queries fail. With a setting of `no` (the default), if a DNS query times out, the request fails. With a setting of `yes`, the system attempts to resolve the query by consulting the DNS table in RAM. If the hostname in the DNS query has an entry in the table in RAM, the system returns the associated IP address(es) to the requester.
`auto-update`	Whether regular successful DNS queries update the local DNS table. For details about `auto-update`, see "Using the auto-update feature" on page 1-55.
`table-config[1–8]`	An array of up to eight hostnames and IP addresses for inclusion in the local DNS table.

DNS-Local Table parameter	Setting
table-config:host-name	A hostname, which must be unique within the table and meet the requirements described in "Hostname matching" on page 1-54.
table-config:ip-address	A valid IP address for the host-name setting, or the zero address. If auto-update is enabled and ip-address specifies the default zero address, successful DNS queries will gradually build the local table.

Hostname matching

A hostname in the local DNS table must start with an alphabetic character and must have fewer than 256 characters. Trailing periods are ignored in the comparison.

The name can be a hostname or a fully qualified domain name. If the name does not include a domain name, and you have specified one or more domain-name settings, the system appends the specified domain name when looking up the hostname. For example, if you have entered the settings shown in "Specifying domain names for lookups" on page 1-51, a DNS query for hostname wheelers results in a search for the following fully qualified domain names:

wheelers.eng.abc.com

wheelers.abc.com

Defining the local table

Following is an example of configuring a local table that specifies three hosts:

```
admin> read ip-global
IP-GLOBAL read

admin> list dns-local
enabled = no
auto-update = no
table-config = [ { "" 0.0.0.0 } {"" 0.0.0.0 } {"" 0.0.0.0 } {"" 0.0.0.0+

admin> set enabled = yes

admin> list table 1
hostname = ""
ip-address = 0.0.0.0

admin> set hostname = host1.abc.com

admin> set ip-address = 10.1.2.3

admin> list ..
table-config[1] = { host1.abc.com 10.1.2.3 }
table-config[2] = { "" 0.0.0.0 }
table-config[3] = { "" 0.0.0.0 }
table-config[4] = { "" 0.0.0.0 }
table-config[5] = { "" 0.0.0.0 }
table-config[6] = { "" 0.0.0.0 }
table-config[7] = { "" 0.0.0.0 }
table-config[8] = { "" 0.0.0.0 }

admin> set 2 hostname = host2.xyz.
```

```
admin> set 2 ip-address = 11.1.2.3
admin> set 3 hostname = localhost
admin> set 3 ip-address = 10.0.0.1
admin> write
IP-GLOBAL written
```

If you specify an IP address without also specifying a hostname, a message such as the following appears when you write the profile:

```
error: dns-local-table: host-name missing (#3 1.2.3.4)
```

If you enter an invalid hostname, a message such as the following appears when you write the profile:

```
error: dns-local-table: host-name must start with alpha char (#5 11foo)
```

Using the auto-update feature

If the `auto-update` parameter is set to `no` (the default), successful DNS queries do not affect the contents of the local table. With a setting of `yes`, when a regular DNS query succeeds, the system performs a lookup on that hostname in the local table. If there is an entry for the hostname, the IP address or addresses associated with that hostname in the local table are replaced by the query response. The number of addresses added to the table depends on the `dns-list-attempt` and `dns-list-size` settings. If `dns-list-attempt` is set to `no`, a successful DNS query returns a single address for a given hostname. In the DNS table in RAM, that address is stored and the remaining 34 addresses are cleared (set to zero).

If `dns-list-attempt` is set to `yes`, a successful DNS query returns the number of addresses it finds for the host, up to the value of `dns-list-size`. In the DNS table in RAM, those addresses are stored, overwriting the configured address or the addresses retrieved from earlier DNS queries. If the table in RAM contains more addresses than `dns-list-size` specifies, the excess addresses are cleared at each update to prevent the accumulation of stale addresses.

Note If you modify the `dns-local-table` subprofile, assigning a single address to a host, the newly configured address is propagated to the table in RAM. The first address of the `host-name` entry is overwritten with the configured address, and all remaining addresses are cleared. If the `auto-update` parameter is set to `yes`, the next successful DNS query overwrites the configured address and restores the multiple addresses (up to the value of `dns-list-size`).

The following commands configure eight hostnames with null addresses and then sets `auto-update` to `yes`. The changes to `dns-local-table` are propagated to RAM, and successful DNS queries to the specified hostnames are builds the local table with up to 14 addresses for each of the hosts.

```
admin> read ip-global
IP-GLOBAL read
admin> set dns-list-attempt = yes
admin> set dns-list-size = 14
admin> list dns-local
enabled = no
auto-update = no
table-config = [ { "" 0.0.0.0 } {"" 0.0.0.0 } {"" 0.0.0.0 } {"" 0.0.0.0+
```

```
admin> set enabled = yes
admin> set auto-update = yes
admin> list table
table-config[1] = { "" 0.0.0.0 }
table-config[2] = { "" 0.0.0.0 }
table-config[3] = { "" 0.0.0.0 }
table-config[4] = { "" 0.0.0.0 }
table-config[5] = { "" 0.0.0.0 }
table-config[6] = { "" 0.0.0.0 }
table-config[7] = { "" 0.0.0.0 }
table-config[8] = { "" 0.0.0.0 }
admin> set 1 host = mercury
admin> set 2 host = venus
admin> set 3 host = earth
admin> set 4 host = mars
admin> set 5 host = jupiter
admin> set 6 host = saturn
admin> set 7 host = uranus
admin> set 8 host = neptune
admin> write
IP-GLOBAL written
```

Using client DNS

Client DNS specifies particular servers for remote clients. ISPs use client DNS to direct callers to servers belonging to particular locations or customers, and to prevent those callers from accessing other clients' host information.

Client DNS can be specified systemwide to allow all remote clients to access one or two DNS servers. Or it can be configured on a connection basis, to allow each appropriately configured connection to access one or two specific servers. At the system level, client DNS also provides an exit mechanism to the local servers if the client servers are inaccessible.

The addresses configured for client DNS servers are presented to WAN connections during IPCP negotiation.

Overview of client DNS settings

You can configure client DNS at the system level in the `ip-global` profile. At the connection level, you can specify client DNS servers in `connection` or RADIUS profiles.

Settings in the `ip-global` *profile*

The following parameters (shown with default values) specify client DNS at the system level:

```
[in IP-GLOBAL]
client-primary-dns-server = 0.0.0.0
```

```
client-secondary-dns-server = 0.0.0.0
allow-as-client-dns-info = True
```

Parameter	Setting
client-dns-primary-server	Address of a client DNS server for remote clients.
client-dns-secondary-server	Address of a secondary DNS server for remote clients.
allow-as-client-dns-info	Enable/disable an exit mechanism to local servers if the client DNS servers are not found. To isolate local network information, set this parameter to false.

Settings in connection *profiles*

The following parameters (shown with default settings) specify client DNS at the connection level:

```
[in CONNECTION/"":ip-options]
client-dns-primary-addr = 0.0.0.0
client-dns-secondary-addr = 0.0.0.0
client-dns-addr-assign = yes
```

Parameter	Setting
client-dns-primary-addr	Address of a client DNS server for the connection.
client-dns-secondary-addr	Address of a secondary client DNS server for the connection.
client-dns-addr-assign	Enable/disable client DNS for the connection. With the yes setting (the default), the system presents client DNS server addresses while negotiating the connection. The addresses it presents can be specified in the connection profile or ip-global profile.

Settings in a RADIUS profile

The following attribute-value pairs configure client DNS in RADIUS profiles:

RADIUS attribute	Value
Ascend-Client-Primary-DNS (135)	Address of a client DNS server for the connection.
Ascend-Client-Secondary-DNS (136)	Address of a secondary client DNS server for the connection.
Ascend-Client-Assign-DNS (137)	Enable/disable client DNS for the connection. With the DNS-Assign-Yes (1) value, the system presents client DNS server addresses while negotiating the connection. The addresses it presents can be specified in the RADIUS profile or ip-global profile.

Example of configuring client DNS servers at the system level

The following commands configure client DNS servers at the system level:

```
admin> read ip-global
IP-GLOBAL read
admin> set client-dns-pri = 10.22.17.56
admin> set client-dns-sec = 10.22.17.107
admin> set allow-as-client-dns-info = false
admin> write
IP-GLOBAL written
```

The secondary server is accessed only if the primary one is inaccessible. If neither of these client DNS servers is accessible and the caller's profile does not specify client DNS servers, the system does *not* allow the client to access local DNS servers.

Examples of configuring client DNS at the connection level

The following commands identify two DNS servers for this connection. The secondary server is accessed only if the primary one is inaccessible.

```
admin> read connection cherry
CONNECTION/cherry read
admin> set ip-options client-dns-primary-addr = 10.2.3.4
admin> set ip-options client-dns-secondary-addr = 10.2.3.56
admin> set ip-options client-dns-addr-assign = yes
admin> write
CONNECTION/cherry written
```

Following are comparable settings in a RADIUS profile:

```
cherry Password = "localpw"
   Service-Type = Framed-User,
   Ascend-Client-Primary-DNS = 10.2.3.4,
   Ascend-Client-Secondary-DNS = 10.2.3.56,
   Ascend-Client-Assign-DNS = DNS-Assign-Yes
```

Configuring Microsoft WINS assignment

In the current software version, you can specify a primary and secondary Windows Internet Name Service (WINS) server on a per-connection basis, either in local `connection` profiles or in RADIUS.

In previous releases, the unit allowed systemwide configuration of a primary and secondary NetBIOS WINS server, to support WINS name resolution for machines connected to a NetBIOS network.

Note The PC initiating the session request must have the Dynamic Host Configuration Protocol (DHCP) for WINS enabled in its Network settings for this feature to work.

Overview of WINS settings

Following are the local parameters (shown with default settings) for configuring client WINS servers:

```
[in CONNECTION/"":ip-options]
client-wins-primary-addr = 0.0.0.0
client-wins-secondary-addr = 0.0.0.0
client-wins-addr-assign = yes
```

Parameter	Setting
client-wins-primary-addr	Address of a client WINS server for the connection.
client-wins-secondary-addr	Address of a secondary client WINS server for the connection.
client-wins-addr-assign	Enable/disable client WINS for the connection. With the yes setting (the default), the system presents client WINS server addresses while negotiating the connection.

For more details about these parameters, see the *Stinger Reference*. For information about specifying NetBIOS servers in the ip-global profile, see "Configuring DNS lookups and a DNS list" on page 1-51.

The following attribute-value pairs configure client WINS servers in RADIUS profiles:

RADIUS attribute	Value
Ascend-Client-Primary-WINS (78)	Address of a client WINS server for the connection.
Ascend-Client-Secondary-WINS (79)	Address of a secondary client WINS server for the connection.
Ascend-Client-Assign-WINS (80)	Enable/disable the use of client WINS servers for the connection. With the WINS-Assign-Yes (1) value, the system presents client WINS server addresses while negotiating the connection.

For more details about these attributes, see the *TAOS RADIUS Guide and Reference*. To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the system must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius
```

```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *Stinger Reference*.

Examples of configuring client WINS servers

The following commands identify two WINS servers for a configured connection. The secondary server is accessed only if the primary one is inaccessible.

```
admin> read connection pc-1
CONNECTION/pc-1 read
```

```
admin> set ip-options client-wins-primary-addr = 10.2.3.4
```

```
admin> set ip-options client-wins-secondary-addr = 10.2.3.56
admin> set ip-options client-wins-addr-assign = yes
admin> write
CONNECTION/pc-1 written
```

Following are comparable settings in a RADIUS profile:

```
pc-1 Password = "localpw", Service-Type = Framed
    Framed-Protocol = PPP,
    Framed-IP-Address = 1.1.1.1,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-Client-Primary-WINS = 10.2.3.4,
    Ascend-Client-Secondary-WINS = 10.2.3.56,
    Ascend-Client-Assign-WINS = WINS-Assign-Yes
```

Configuring and using address pools

An address pool is a range of contiguous addresses on a local IP network or subnet. Pool addresses are available for assignment to incoming callers that request an address. When the call terminates, the address is returned to the pool, making it available again for assignment.

If you designate a subnet for IP address pools, you must make sure that other IP hosts on the local network know the route to that subnet. You must also make sure that the pools do not overlap (do not contain duplicate addresses).

For related information, see "Defining address pools for a virtual router" on page 4-9.

Overview of settings for defining pools

You can define up to 128 address pools locally in the `ip-global` profile. Those pools can be used to assign addresses to callers authenticated locally (in `connection` profiles) or by RADIUS. If you are using RADIUS authentication, you can choose to define address pools in RADIUS instead of, or in addition to, those defined locally. If you have the RADIPAD program installed, you can use it to manage address pools centrally on a single RADIUS server.

Settings in the `ip-global` profile

The following parameters (shown with default values) configure address pools locally:

```
[in IP-GLOBAL]
pool-summary = no
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0+
pool-name = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" ""+
```

Parameter	Setting
`pool-summary`	Set/clear the Pool Summary flag. For details, see "Example of configuring summarized address pools" on page 1-64.
`pool-base-address`	Base address of a pool of contiguous addresses on a local network or subnet.

Parameter	Setting
assign-count	Number of addresses in the pool.
pool-name	A pool name, required only when TACACS+ authentication is in use. If TACACS+ authentication is not in use, the name is treated as a comment.

Settings in RADIUS pseudo-user profiles

You can define address pools in a RADIUS pools pseudo-user profile. The first line of pools pseudo-user profile uses the following format:

pools-*name* Password = "ascend", Service-Type = Outbound-User

The *name* argument is the system name (specified by the name parameter in the system profile). Subsequent lines in the profile define IP address pools by using the Ascend-IP-Pool-Definition (217) attribute. The value of the Ascend-IP-Pool-Definition attribute uses the following syntax:

"*pool-num base-addr assign-count*"

Syntax element	Description
pool-num	Pool number. If you use the same number to designate two pools, one locally and one in RADIUS, the RADIUS definition takes precedence. So if you have defined some pools in the ip-global profile and do not wish to override them, start numbering the pools at the next number. For example, if you defined 10 pools in the ip-global profile, start with number 11 in RADIUS. Otherwise, start with 1.
base-addr	The base address in a pool of contiguous addresses on the local network or subnet.
assign-count	Number of addresses included in the pool.

Global RADIUS pools (RADIPAD)

RADIUS IP Address Daemon (RADIPAD) is a program that works with RADIUS authentication to manage IP address pools centrally, so that connections can all acquire an address from a global pool, regardless of which system answers the call.

RADIPAD runs on one RADIUS server on the network. A Stinger unit sends an authentication request to RADIUS, and if the user profile contains an attribute to allocate an IP address from the global pool, RADIUS sends a request to RADIPAD to acquire the address.

The Stinger unit does not communicate directly with RADIPAD, so it does not require additional configuration to use RADIPAD. To configure RADIPAD, you define the global pools of addresses, specify which RADIUS server is running RADIPAD, and (optionally) specify which systems can obtain addresses from those pools. You can then create RADIUS user profiles that acquire an IP address from the global pool.

At startup, syslog notes RADIUS requests to release RADIUS-allocated IP addresses. Some versions of the RADIUS server might time out the request, resulting in log messages indicating the release of global-pool addresses.

Defining global pools

Global address pools are defined in a `global-pools` pseudo-user profile on the server running RADIPAD. The first line of a `global-pools` pseudo-user profile uses the following format:

`global-pools-`*`name`* `Password = "ascend", Service-Type = Outbound-User`

The *name* argument is a designation for any class of users. You can create multiple global pool profiles for multiple user classes. For example, you could create profiles named `global-pool-ppp`, `global-pool-slip`, and so forth. Subsequent lines in the profile define IP address pools by using the Ascend-IP-Pool-Definition (217) attribute. This is the same attribute described in "Settings in RADIUS pseudo-user profiles" on page 1-61, and it follows the same rules for global pools. In addition, when the Stinger unit assigns an address from a pool managed by the RADIPAD daemon, RADIPAD tries to allocate an address from the pools in order, by pool number, and chooses an address from the first pool with an available IP address.

Specifying the RADIPAD host

Each RADIUS server must specify the host running RADIPAD and (optionally) the systems that can access the global pools. These settings are defined in a `radipa-hosts` pseudo-user profile, which uses the following format in the first line of the profile:

`radipa-hosts Password = "ascend", Service-Type = Outbound-User`

Subsequent lines in the profile use the following attribute-value pairs to define which hosts can assign addresses from the pools managed by RADIPAD:

RADIUS attribute	Value
Ascend-Assign-IP-Client (144)	Address of a system that is allowed to access the global address pools managed by RADIPAD. You can specify multiple instances of this attribute. If no client addresses are specified, all units listed in the RADIUS clients file can access RADIPAD pools.
Ascend-Assign-IP-Server (145)	Address of the server running RADIPAD. Only one instance of this attribute can appear in the profile, and it must specify the correct IP address.

For example:

```
radipa-hosts Password ="ascend", Service-Type = Outbound-User
   Ascend-Assign-IP-Server = 10.31.4.34,
   Ascend-Assign-IP-Client = 10.31.4.10,
   Ascend-Assign-IP-Client = 10.31.4.11
```

You can specify only one RADIPAD server, but you can include multiple clients. The sample profile indicates that two systems (10.31.4.10 and 10.31.4.11) can access RADIPAD pools as clients.

Preventing the use of class boundary addresses

If you define address pools that contain more than 254 addresses, be aware that the system allocates the class boundary addresses (*n.n.n*.0 and *n.n.n*.255) as valid caller

addresses. According to CIDR, this is permitted because the pool is not a /24 network. However, some client systems, such as Windows, do not tolerate the class boundary addresses well. For example, because Windows assumes a /24 network, it broadcasts NetBIOS over IP name service to the .255 address, which could overwhelm a connection assigned the .255 host address.

To prevent client software from using a host address for broadcasts, you must explicitly apply a filter that prevents the system from using the class boundary addresses. For example, if you are using RADIUS authentication, you can apply a data filter, in the `answer-defaults` profile, that drops packets from any source to pool address *n.n.n*.0 or *n.n.n*.255.

Examples of configuring address pools

For a pool that is not summarized, each assigned address is advertised as its own host route. Such a pool can start at any base address. Addresses do not accept a subnet mask component, because they are always advertised as host routes.

The following commands define three address pools, each containing 50 addresses. Pool 1 contains 10.2.3.4 through 10.2.3.54. Pool 2 contains 11.5.7.51 through 11.5.7.101. Pool 3 contains 12.7.112.15 through 12.7.112.65.

```
admin> read ip-global
IP-GLOBAL read
admin> set pool-base-address 1 = 10.2.3.4
admin> set pool-base-address 2 = 11.5.7.51
admin> set pool-base-address 3 = 12.7.112.15
admin> set assign-count 1 = 50
admin> set assign-count 2 = 50
admin> set assign-count 3 = 50
admin> write
IP-GLOBAL written
```

Following is a comparable RADIUS `pools` profile (for use by a single RADIUS server):

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 10.2.3.4 50",
   Ascend-IP-Pool-Definition = "2 11.5.7.51 50",
   Ascend-IP-Pool-Definition = "3 12.7.112.15 50"
```

Following is a comparable global pools definition (for use with RADIPAD):

```
global-pool-ppp Password ="ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 10.2.3.4 50",
   Ascend-IP-Pool-Definition = "2 11.5.7.51 50",
   Ascend-IP-Pool-Definition = "3 12.7.112.15 50"
```

Although some client software assumes a default subnet mask of 255.255.255.0 for PPP interfaces, you can define pools on subnets wider than /24. For example, the following commands define an address pool on a /23 subnet:

```
admin> read ip-global
IP-GLOBAL read
admin> set pool-base-address 1 = 10.55.178.1
admin> set assign-count 1 = 510
```

```
admin> write
IP-GLOBAL written
```

This pool definition translates to 10.55.178.0/23 (a subnet mask of 255.255.254.0). Following are comparable RADIUS definitions:

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 10.55.178.1 510"
global-pool-ppp Password ="ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 10.55.178.1 510"
```

Example of configuring summarized address pools

The pool-summary feature reduces routing overhead associated with address pools. Instead of advertising each address assigned from a pool as a host route, the system suppresses the host route advertisements and instead advertises a static route to the pool itself.

To use summarized pools locally or in RADIUS, you must set the pool-summary flag to yes in the ip-global profile, and you must define all pools to be network-aligned.

Setting the pool-summary flag

The following commands enable the pool-summary flag:

```
admin> read ip-global
IP-GLOBAL read
admin> set pool-summary = yes
admin> write
IP-GLOBAL written
```

Defining network-aligned pools

Following are the rules for network-aligned address pools:

- The specified number of addresses in the pool must be two less than the total number of addresses in the pool. (Add 2 to the assign-count value for the total number of addresses in the subnet, and calculate the mask for the subnet on the basis of this total.)

  assign-count + 2 = number of subnet hosts

- The specified base address of the pool must be the first host address. (Subtract 1 from the pool-base-address value for the base address for the subnet.)

  pool-base-address - 1 = network-aligned subnet address

The following commands enable the pool-summary flag and define a network-aligned pool:

```
admin> read ip-global
IP-GLOBAL read
admin> set pool-summary = yes
admin> set assign-count 1 = 62
admin> set pool-base-address 1 = 10.12.253.1
admin> write
IP-GLOBAL written
```

In the preceding sample configurations, the `assign-count` parameter is set to 62. When you add 2 to this value, you get 64. The subnet mask for 64 addresses is 255.255.255.192 (256 – 64 = 192). The prefix length for a 255.255.255.192 mask is /26.

The `pool-base-address` parameter is set to 10.12.253.1. When you subtract 1 from this value, you get 10.12.253.0, which is a valid network-aligned base address for the 255.255.255.192 subnet mask. (Note that 10.12.253.64, 10.12.253.128, and 10.12.253.192 are also valid zero addresses for the same mask.) The resulting address pool subnet is 10.12.253.0/26.

Following is a comparable RADIUS `pools` profile (for use by a single RADIUS server).

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 10.12.253.1 62"
```

Following is a comparable global pools definition (for use with RADIPAD):

```
global-pool-ppp Password ="ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 10.12.253.1 62"
```

The system still creates (but does not advertise) a host route for each assigned address in the pool. Host routes take precedence over subnet routes, so packets whose destination matches an assigned IP address from the pool are routed properly. However, because the system advertises the entire pool as a route, and only privately knows which IP addresses in the pool are active, a remote network might improperly send the Stinger unit a packet for an inactive IP address. If that occurs, the packets are routed to the Reject (`rj0`) interface (127.0.0.2). Packets routed to the Reject interface are bounced back to the sender with an `ICMP unreachable` message.

Examples of assigning an address from a pool

When an incoming call requests an IP address, the Stinger unit assigns one from a pool. A host requests an address if its client software has settings such as those shown in Figure 1-9:

Figure 1-9. Client software settings requesting dynamic address assignment

```
Username=victor
Accept Assigned IP=Yes
IP address=Dynamic (or Assigned or N/A)
Netmask=255.255.255.255 (or None or N/A)
Default Gateway=None or N/A
Name Server=10.2.3.55
Domain suffix=abc.com
Baud rate=38400
Hardware handshaking ON
VAN Jacobson compression ON
```

Figure 1-10 shows a remote host requesting and being assigned an IP address.

Figure 1-10. Remote CPE requiring assigned IP address

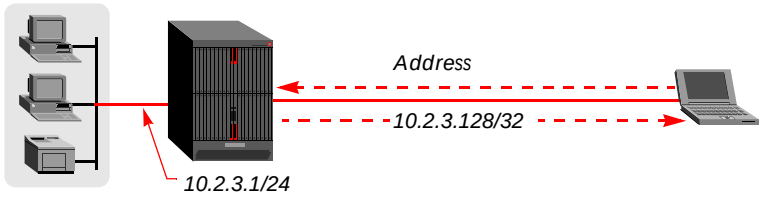

The following commands enable dynamic address assignment systemwide:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set ip-answer assign = yes

admin> write
ANSWER-DEFAULTS written
```

For information about ensuring that connections must accept the address offered, see “Requiring acceptance of dynamic address assignment” on page 1-37.

The following commands configure a profile to acquire an address from the first pool that has available addresses:

```
admin> new connection victor
CONNECTION/victor read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp recv-password = localpw

admin> set ip-options address-pool = 0

admin> write
CONNECTION/victor written
```

Following is a comparable RADIUS profile:

```
victor Password = "localpw"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Ascend-Assign-IP-Pool = 0
```

Following is a comparable RADIUS profile that acquires an address from any global pool managed by the RADIPAD daemon:

```
victor Password = "localpw"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Ascend-Assign-IP-Pool = 65535,
   Ascend-Assign-ip-global-Pool = "global-pool-ppp"
```

IP pool chaining

Because the addresses within a pool must be contiguous, many sites have defined a large number of pools, with each pool containing only a small range of addresses. For example, the following RADIUS profile defines six pools, each containing 10 addresses:

```
pools-JFAN-TNT Password = "ascend"
    Service-Type = Outbound,
    Ascend-IP-Pool-Definition = "1 11.168.6.10 10",
    Ascend-IP-Pool-Definition = "2 12.168.6.10 10",
    Ascend-IP-Pool-Definition = "3 13.168.6.10 10",
    Ascend-IP-Pool-Definition = "7 17.168.6.10 10",
    Ascend-IP-Pool-Definition = "8 18.168.6.10 10",
    Ascend-IP-Pool-Definition = "9 19.168.6.10 10"
```

In earlier versions of the software, you could allow a caller to acquire an address from any pool (by assigning the pool number 0 in the caller's profile) or from a single specified pool, such as pool 1. IP pool chaining enables you to allow a caller to acquire an address from any pool within a chain.

When IP pool chaining is enabled, contiguous pools are treated as one *pool space* with shared addresses. When the system assigns an address to an end user, it begins searching for an available address in the first pool of the chain and stops when it either finds an available address or encounters a null pool definition. So, the pools within a chain must be defined in a contiguous sequence. For example, the following profile contains two IP pool chains (pools 1, 2, 3 and pools 7, 8, 9), with each pool chain containing 30 addresses:

```
pools-JFAN-TNT Password = "ascend", Service-Type = Outbound
    Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes,
    Ascend-IP-Pool-Definition = "1 11.168.6.10 10",
    Ascend-IP-Pool-Definition = "2 12.168.6.10 10",
    Ascend-IP-Pool-Definition = "3 13.168.6.10 10",
    Ascend-IP-Pool-Definition = "7 17.168.6.10 10",
    Ascend-IP-Pool-Definition = "8 18.168.6.10 10",
    Ascend-IP-Pool-Definition = "9 19.168.6.10 10"
```

Note To support IP pool chaining in RADIUS profiles, the RADIUS server must support vendor-specific attributes (VSA) and the system must be configured in VSA compatibility mode. For details, see "Pool chaining in RADIUS" on page 1-69.

IP pool chaining is supported both for RADIUS-defined address pools and for pools defined locally in the `ip-global` profile. For example, the following settings in the `ip-global` profile enable pool chaining and define a pool chain (pools 1 and 2) that contains 252 addresses:

```
[in IP-GLOBAL]
pool-chaining = yes
pool-base-address = [ 172.20.31.1 172.20.33.1 0.0.0.0 153.37.21.1 0.0+
assign-count = [ 126 126 0 30 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0+
```

Pool chaining in local profiles

Whether pool chains are defined locally or in RADIUS, the pool addresses are available for dynamic assignment regardless of where the caller's profile is authenticated.

Overview of local profile settings

Following are the parameters, shown with default settings, relevant to IP pool chaining:

```
[in IP-GLOBAL]
pool-chaining = no
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0+
[in CONNECTION/"":ip-options]
address-pool = 0
```

Parameter	Setting
pool-chaining	Enable/disable IP pool chaining. With the yes setting, the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a caller.
pool-base-address	An array of up to 128 IP addresses to be used as the first address in a pool. These values are used with the assign-count values to define address pools locally. A pool chain contains all of the pools defined in sequence within the array, such as 1, 2, 3. To end a pool chain, leave a null value in the array.
assign-count	An array of up to 128 numbers that specify the number of addresses in a pool that starts with the corresponding pool-base-address.
address-pool	Number of an address pool from which to acquire an address. When pool chaining is enabled, a pool number within a chain includes addresses defined in all other pools within the chain. For example, if pools 1, 2, and 3 are in a pool chain, setting this parameter to 1 has the same effect as setting it to 2 or 3.

Example of local pool chain definition

The following commands define five address pools, which form two pool chains. Notice that the pool numbers (their indexes in the pool-base-address and assign-count arrays) are contiguous within a chain.

```
admin> read ip-global
IP-GLOBAL read
admin> set pool-chaining = yes
admin> set pool-base-address 1 = 10.1.1.1
admin> set pool-base-address 2 = 11.1.1.1
admin> set pool-base-address 3 = 12.1.1.1
admin> set assign-count 1 = 50
admin> set assign-count 2 = 50
admin> set assign-count 3 = 50
admin> set pool-base-address 7 = 13.1.1.1
admin> set pool-base-address 8 = 14.1.1.1
admin> set assign-count 7 = 50
admin> set assign-count 8 = 50
```

```
admin> write
IP-GLOBAL written
```

The following commands enable dynamic address assignment systemwide:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set ip-answer assign = yes

admin> write
ANSWER-DEFAULTS written
```

The following commands configure profiles to acquire an address from the first pool chain. When the end users initiate a session request, they can acquire an address from 10.1.1.1 to 10.1.1.51, from 11.1.1.1 to 11.1.1.51, or from 12.1.1.1 to 12.1.1.51. If no addresses are available within those ranges, the connection is refused.

```
admin> new connection jfan
CONNECTION/jfan read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp-options recv-password = localpw

admin> set ip-options address-pool = 2

admin> write
CONNECTION/jfan written

admin> new connection ravi
CONNECTION/ravi read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp-options recv-password = localpw

admin> set ip-options address-pool = 1

admin> write
CONNECTION/ravi written
```

Following are comparable RADIUS profiles:

```
jfan Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 2

ravi Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 1
```

Pool chaining in RADIUS

Whether pool chains are defined locally or in a RADIUS pool's pseudo-user profile, the pool addresses are available for dynamic assignment regardless of where the caller's profile is authenticated.

Overview of RADIUS profile settings

RADIUS servers use the following attribute-value pairs to define and apply pool chains:

RADIUS attribute	Value
Ascend-IP-Pool-Chaining (85)	Enable/disable IP pool chaining in a pseudo-user profile that defines address pools. If this attribute is set to IP-Pool-Chaining-Yes (1), the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a caller. With a value of IP-Pool-Chaining-No (0), the system treats each address pool as a separate space. **Note** When this attribute is specified in a RADIUS profile, its value overrides the Pool-Chaining setting in the `ip-global` profile.
Ascend-IP-Pool-Definition (217)	Address pool definition in a pseudo-user profile. The value has the following syntax: *`pool-number base-addr assign-count`* The *`pool-number`* value is an integer that identifies the pool. A pool chain contains all of the pools defined in sequence, such as 1, 2, 3. To end a pool chain, leave a gap in the sequence of *`pool-number`* values. The *`base-addr`* value is an IP address to be used as the first address in a pool, and the *`assign-count`* value specifies the number of addresses in a pool.
Ascend-Assign-IP-Pool (218)	Number of the address pool from which the RADIUS user profile should acquire an address. When pool chaining is enabled, a pool number within a chain includes addresses defined in all other pools within the chain. For example, if pools 1, 2, and 3 are in a pool chain, setting this value to 1 has the same effect as setting it to 2 or 3.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the system must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius
```

```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *Stinger Reference*.

Example of pool chaining in RADIUS

The following pseudo-user profile defines five address pools, which form two pool chains. Notice that the pool numbers are contiguous within a chain.

```
pools-JFAN-TNT Password = "ascend"
    Service-Type = Outbound,
    Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes,
    Ascend-IP-Pool-Definition = "1 10.1.1.1 50",
    Ascend-IP-Pool-Definition = "2 11.1.1.1 50",
    Ascend-IP-Pool-Definition = "3 12.1.1.1 50",
    Ascend-IP-Pool-Definition = "7 13.1.1.1 50",
    Ascend-IP-Pool-Definition = "8 14.1.1.1 50"
```

The following commands configure local connection profiles to acquire an address from the first pool chain. When the end users initiate a session request, they can acquire an address from 13.1.1.1 to 13.1.1.51, or from 14.1.1.1 to 14.1.1.51. If no addresses are available within those ranges, the connection is refused.

```
admin> new connection hanif
CONNECTION/hanif read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp-options recv-password = localpw

admin> set ip-options address-pool = 7

admin> write
CONNECTION/hanif written

admin> new connection hasnain
CONNECTION/hasnain read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp-options recv-password = localpw

admin> set ip-options address-pool = 8

admin> write
CONNECTION/hasnain written
```

Following are comparable RADIUS user profiles:

```
hanif Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 7

hasnain Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 8
```

IGMP Multicast Forwarding

2

Introduction to multicast forwarding

The IP multicast backbone (MBONE) provides one-to-many and many-to-many communication, rather than the point-to-point communication used by many other types of network applications. Video and audio transmissions use the MBONE as a much cheaper and faster way to communicate the same information to multiple hosts.

MBONE routers maintain multicast groups, in which hosts must register to receive a multicast transmission. Multicast group functions are handled with the Internet Group Management Protocol (IGMP). The Stinger unit forwards IGMP version-1 or version-2 packets, including IGMP MTRACE (multicast trace).

Figure 2-1 shows a Stinger unit forwarding multicast traffic from an MBONE router across the WAN to two WAN multicast client interfaces and a LAN multicast client interface.

Figure 2-1. Stinger unit forwarding multicast traffic to LAN and WAN clients

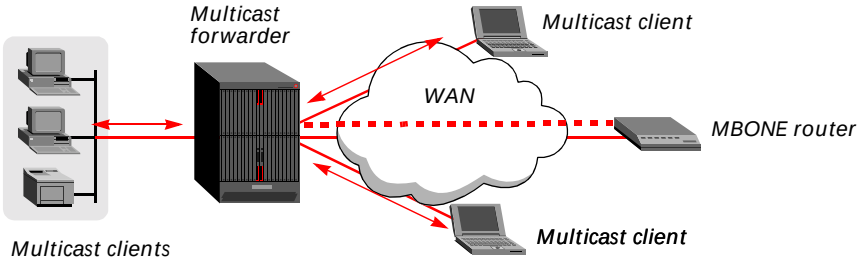

The interface to the MBONE router is the MBONE interface. The Stinger unit can have one and only one MBONE interface, which can be either a LAN or WAN IP interface.

To MBONE routers, the Stinger unit appears to be a multicast client, because it responds as a client to IGMP packets. To multicast clients, the Stinger unit appears to be an MBONE router, because it forwards IGMP queries to those clients, receives their responses, and forwards multicast traffic.

Enabling multicast forwarding in the system

The following parameters (shown with default settings) initiate multicast forwarding at the system level:

```
[in IP-GLOBAL]
multicast-forwarding = no
mbone-profile = ""
mbone-lan-interface = { { any-shelf any-slot 0 } 0 }
multicast-hbeat-addr = 0.0.0.0
multicast-hbeat-port = 0
multicast-hbeat-slot-time = 0
multicast-hbeat-number-slot = 0
multicast-hbeat-alarm-threshold = 0
multicast-hbeat-src-addr = 0.0.0.0
multicast-hbeat-src-addr-mask = 0.0.0.0
multicast-member-timeout = 360
```

Note Heartbeat monitoring is optional. It is not required for multicast forwarding.

Parameter	Setting
`multicast-forwarding`	Enable/disable multicast forwarding. When you change the value to `yes` and write the profile, the multicast subsystem reads the values in the `ip-global` profile and initiates the forwarding function.
`mbone-profile`	Name of a local `connection` profile for an MBONE router on a WAN interface. This parameter and the `mbone-lan-interface` parameter are mutually exclusive. For details, see “Configuring the MBONE interface” on page 2-5.
`mbone-lan-interface`	Interface address (shelf, slot, and port) to MBONE router on a LAN interface. This parameter and the `mbone-profile` parameter are mutually exclusive. For details, see “Configuring the MBONE interface” on page 2-5.
`multicast-hbeat-addr`	Multicast address to be monitored for determining a minimal level of traffic (heartbeat).
`multicast-hbeat-port`	UDP port number to be monitored. The system counts only packets received on this port.
`multicast-hbeat-slot-time`	Polling interval (in seconds) during which the system polls for multicast traffic.

Parameter	Setting
multicast-hbeat-number-slot	Number of times to poll for the specified interval before comparing the number of heartbeat packets received to the alarm threshold.
multicast-hbeat-alarm-threshold	Number of packets that represents normal multicast traffic. If the number of monitored packets falls below this number, the SNMP alarm trap is sent.
multicast-hbeat-src-addr	Source IP address to be ignored. Packets received from that address are ignored for heartbeat-monitoring purposes.
multicast-hbeat-src-addr-mask	Subnet mask to apply to a multicast-hbeat-src-addr value before comparing it to the source address in a packet.
multicast-member-timeout	Timeout (in seconds) for client responses to multicast polling messages. If it does not receive responses on a client interface in the specified number of seconds, the system stops forwarding multicast traffic on the interface.

Specifying a timeout for group memberships

The multicast-member-timeout parameter specifies the timeout (in seconds) for client responses to multicast polling messages. If no client responds to the polling messages within the amount of time you specify for multicast-member-timeout, the system stops forwarding multicast traffic on the interface. The following commands set the timeout value to 60 seconds:

```
admin> read ip-global
IP-GLOBAL read

admin> set multicast-member-timeout = 60

admin> write
IP-GLOBAL written
```

Monitoring the multicast traffic heartbeat

Heartbeat monitoring is optional. It enables you to monitor possible multicast connectivity problems by continuously polling for a certain level of multicast traffic and generating the following SNMP alarm trap in the event of a traffic breakdown:

```
Trap type: TRAP_ENTERPRISE
Code:       TRAP_MULTICAST_TREE_BROKEN (19)
Arguments:
1) Multicast group address being monitored (4 bytes),
2) Source address of last heartbeat packet received (4 bytes)
3) Slot time interval configured in seconds (4 bytes),
4) Number of slots configured (4 bytes).
5) Total number of heartbeat packets received before the unit started
sending SNMP Alarms (4 bytes).
```

Enabling heartbeat monitoring

To enable multicast heartbeat monitoring, you specify a polling frequency and the threshold below which the alarm is generated.

With the following sample configuration, the system polls 10 times at 10-second intervals and then compares the total traffic count to the threshold value. If fewer than 30 packets have been received, the unit generates the SNMP alarm.

```
admin> read ip-global
IP-GLOBAL/ read
admin> set multicast-hbeat-slot-time = 10
admin> set multicast-hbeat-number-slot = 10
admin> set multicast-hbeat-alarm-threshold = 30
admin> write
IP-GLOBAL/ written
```

Specifying which packets to monitor

To fine-tune heartbeat monitoring, you can specify which packets the system counts as multicast traffic. You can do this in one or more of the following ways:

- Specify a particular multicast address to be used for monitoring.
- Specify a UDP port number (all packets received on that port will be used for monitoring).
- Specify a source address (all packets from that host will be ignored for monitoring purposes).
- Specify a subnet mask to be applied to the source address (all packets from the subnet or network will be ignored for monitoring purposes).

The following example shows how to monitor only packets forwarded to and received from the 224.1.1.1 multicast address:

```
admin> read ip-global
IP-GLOBAL/ read
admin> set multicast-hbeat-addr = 224.1.1.1
admin> write
IP-GLOBAL/ written
```

The next sample configuration limits monitoring to packets forwarded to and received from the multicast address 224.1.1.1 on UDP port 16387:

```
admin> read ip-global
IP-GLOBAL/ read
admin> set multicast-hbeat-addr = 224.1.1.1
admin> set multicast-hbeat-port = 16387
admin> write
IP-GLOBAL/ written
```

The following example shows how to specify that multicast packets from the 10.1.0.0 subnet will be ignored for heartbeat-monitoring purposes:

```
admin> read ip-global
IP-GLOBAL/ read
```

```
admin> set multicast-hbeat-src-addr = 10.1.2.3
admin> set multicast-hbeat-src-addr-mask = 255.255.0.0
admin> write
IP-GLOBAL/ written
```

Configuring the MBONE interface

The MBONE interface is the single LAN or WAN IP interface on which an MBONE router resides. The MBONE interface cannot support multicast clients.

To enable a Stinger unit to forward traffic to and from an MBONE router, you must configure both the `ip-global` settings and the appropriate settings in an `ip-interface` or `connection` profile.

Overview of MBONE interface settings

The following parameter (shown with its default setting) is used on the MBONE interface:

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 } ]
multicast-allowed = no
[in CONNECTION/"":ip-options]
multicast-allowed = no
```

Parameter	Setting
`multicast-allowed`	Enable/disable handling of IGMP requests and responses on the interface. The system does *not* forward multicast traffic on the basis of this setting.

Typical local MBONE router configuration

Figure 2-2 shows an MBONE router on one of the system's LAN IP interfaces.

Figure 2-2. MBONE router on a LAN interface

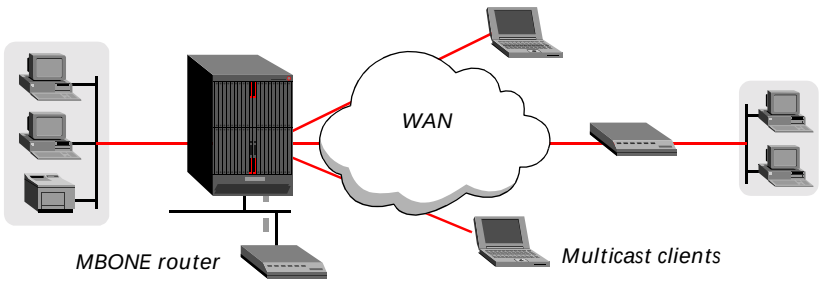

The following commands configure the first Ethernet interface of a T1000 module as the MBONE interface:

```
admin> read ip-global
IP-GLOBAL read
admin> set multicast-forwarding = yes
admin> set mbone-lan-interface = { { 1 3 1 } 0}
```

```
admin> write
IP-GLOBAL written
admin> read ip-interface { { 1 3 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } read
admin> set multicast-allowed = yes
admin> write
IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } written
```

Typical configuration of an MBONE router on a WAN interface

Figure 2-3 shows an MBONE router on a WAN interface.

Figure 2-3. MBONE router on a WAN interface

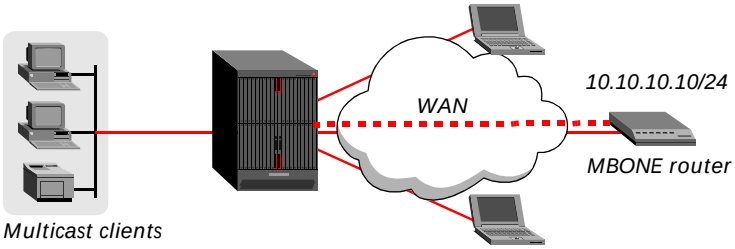

The following commands configure a WAN IP interface to the MBONE router:

```
admin> read ip-global
IP-GLOBAL read
admin> set multicast-forwarding = yes
admin> set mbone-profile = multicast-router
admin> write
IP-GLOBAL written
admin> read connection multicast-router
CONNECTION/multicast-router read
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 10.10.10.10/24
admin> set ip-options multicast-allowed = yes
admin> set atm-options atm1483type = aal5-llc
admin> set atm-options vci = 101
admin> set atm-options nailed-group = 802
admin> set atm-qos-options usr-up-stream-contract = cbr
admin> set atm-qos-options usr-dn-stream-contract = cbr
admin> write
CONNECTION/multicast-router written
```

Managing multicast group memberships

To receive a multicast transmission, a client interface must join a specific multicast group (a Class D IP address from 224.0.0.0 to 239.255.255.255). When data is sent to an address in that range, it is multicast to all hosts that have joined that group.

The `mcast-service` profile provides a way to manage which multicast groups can be accessed by a client interface. Each profile specifies a number of multicast groups (up to 256), and a filter type. The filter type determines how the list of multicast groups is used: to allow access only to those groups, or allow access to all groups *except* those listed.

You can configure multiple `mcast-service` profile, one for each level of multicast services you provide. You can define the profiles locally or via RADIUS. You then apply the appropriate profile to a client interface by specifying the profile name in the client's `connection` or RADIUS profile (for details, see "Configuring multicast client interfaces" on page 2-10).

Overview of `mcast-service` settings

Following are the `mcast-service` settings, shown with default values:

```
[in MCAST-SERVICE/""]
service-name* = ""
active = no
snmp-trap-enable = no
filter-type = none
filter-list = [ { no 0.0.0.0 } { no 0.0.0.0 } { no 0.0.0.0 } { no 0.0.0.0 }+

[in MCAST-SERVICE/"":filter-list[1]]
active = no
mcast-ip-address = 0.0.0.0
```

Parameter	RADIUS attribute	Setting
`service-name`	Ascend-Multicast-Service-Name (276)	Name assigned to the service profile, up to 31 characters.
`active`	Ascend-Multicast-Service-Active (277)	Enables or disables the profile. *Note:* If the profile is disabled, none of the clients on the interfaces whose profiles point to this service profile are allowed to join any multicast groups. If you want the opposite effect, to allow clients on those interfaces access to all multicast groups, use the `filter-type` setting instead.
`snmp-trap-enable`	Ascend-Multicast-Service-Snmp-Trap (278)	Enables or disables sending a trap for a multicast link up or link down event for all client interfaces associated with this profile. To send these traps, you must set this parameter to `yes`. In addition, you must also enable the traps at the system level by setting the `ascend-multicast-link-trap-enabled` parameter in the `trap` profile. The objects reported in these traps are contained in the `mcastserv.mib`.

Parameter	RADIUS attribute	Setting
`filter-type`	Ascend-Multicast-Service-Filter-Type (279)	Specifies whether access to the multicast groups defined in the filter list will be filtered inclusively or exclusively. With inclusive filtering, client interfaces have access only to those groups specified in the filter list. With exclusive filtering, clients have access to all multicast groups *except* those in the list. If you set this to `none`, access to all multicast groups is allowed.
`filter-list[`*n*`]: active`	Ascend-Multicast-Filter-Active (280)	The filter list contains 256 indexed subprofiles, each of which specifies a multicast group address filter. The `active` parameter enables or disables the filter. When the filter is enabled, access to the address specified in the `mcast-ip-address` parameter is controlled as specified in the `filter-type` setting. If it is disabled, the filter has no effect.
`filter-list[`*n*`]: mcast-ip-address`	Ascend-Multicast-Filter-Address (281)	Class D IP address from 224.0.0.0 to 239.255.255.255. When data is sent to this address, it is multicast to all hosts that have joined that group.

Sample multicast service configurations

In this example, a multicast video server supports two multicast group addresses. Multicast clients can subscribe to either the “bronze” or the “gold” multicast service. Bronze service permits access to the group at 239.255.129.119. Gold service permits access to both 239.255.129.119 and a premium group at 239.255.129.120.

The following commands configure the `mcast-service` profiles:

```
admin> new mcast-service bronze-service
MCAST-SERVICE/bronze-service read
admin> set active = yes
admin> set snmp-trap-enable = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
admin> set filter-list 1 mcast-ip-address = 239.255.129.119
admin> write -f
MCAST-SERVICE/bronze-service written
admin> new mcast-service gold-service
MCAST-SERVICE/gold-service read
admin> set active = yes
admin> set snmp-trap-enable = yes
admin> set filter-type = inclusive
admin> set filter-list 1 active = yes
admin> set filter-list 1 mcast-ip-address = 239.225.129.119
admin> set filter-list 2 active = yes
admin> set filter-list 2 mcast-ip-address = 239.255.129.120
admin> write -f
```

```
MCAST-SERVICE/gold-service written
admin> dir mcast-service
   802  11/05/2002 20:12:09  bronze-service
   809  11/05/2002 20:13:19  gold-service
```

Following are comparable RADIUS profiles:

```
mcastService-ipstinger-1 password = "pwd"
    Ascend-Multicast-Service-Name = "bronze-service"
    Ascend-Multicast-Service-Active = "Multicast-Service-Yes"
    Ascend-Multicast-Service-Snmp-Trap = "Multicast-Snmp-Trap-Yes"
    Ascend-Multicast-Service-Filter-Type = "Multicast-Filter-Inclusive"
    Ascend-Multicast-Filter-Address = "239.255.129.119"
mcastService-ipstinger-2 password = "pwd"
    Ascend-Multicast-Service-Name = "gold-service"
    Ascend-Multicast-Service-Active = "Multicast-Service-Yes"
    Ascend-Multicast-Service-Snmp-Trap = "Multicast-Snmp-Trap-Yes"
    Ascend-Multicast-Service-Filter-Type = "Multicast-Filter-Inclusive"
    Ascend-Multicast-Filter-Address = "239.255.129.119"
    Ascend-Multicast-Service-Snmp-Trap = "Multicast-Snmp-Trap-Yes"
    Ascend-Multicast-Service-Filter-Type = "Multicast-Filter-Inclusive"
    Ascend-Multicast-Filter-Address = "239.255.129.120"
```

Note You must set the Ascend-Multicast-Filter-Active attribute after each Ascend-Multicast-Filter-Address setting, because it increments the index for the next Ascend-Multicast-Filter-Address setting.

The following command displays information about the multicast service profiles:

```
admin> igmp profiles
IGMP Service Profiles

          Service Name               : gold-service
          SNMP Trap                  : Enabled
          Call logging               : Disabled
          Filter Type                : MCAST_FILTER_INCLUSIVE
          Filter List                :
                 224.255.129.120
                 224.225.129.119

          Service Name               : bronze-service
          SNMP Trap                  : Enabled
          Call logging               : Disabled
          Filter Type                : MCAST_FILTER_INCLUSIVE
          Filter List                :
                 224.255.129.119
```

Setting the multicast rate limit

The `multicast-rate-limit` parameter specifies the rate at which the Stinger unit accepts multicast packets from clients on the interface.

Note By default, `multicast-rate-limit` is set to 100. This setting disables multicast forwarding on the interface. (The forwarder handles IGMP packets, but does not accept packets from clients or forward multicast packets from the MBONE router.) To

enable multicast forwarding on the interface, you must set the `multicast-rate-limit` parameter to a number *less than* 100.

For example, if you set `multicast-rate-limit` to 5, the system accepts one packet every 5 seconds from multicast clients on the interface. Any subsequent packets received within that 5-second window are discarded.

In addition to multicast rate limiting, the Stinger unit also supports prioritized packet dropping for high-bandwidth data, voice, and audio multicast applications. If the system is the receiving device under extremely high loads, it drops packets according to a priority ranking, which is determined by the following UDP port ranges:

- Traffic on ports 0–16384 (unclassified traffic) has the lowest priority (50).
- Traffic on ports 16385–32768 (audio traffic) has the highest priority (70).
- Traffic on ports 32769–49152 (whiteboard traffic) has medium priority (60).
- Traffic on ports 49153–65536 (video traffic) has low priority (55).

Specifying a delay for clearing IGMP groups

The `multicast-group-leave-delay` parameter specifies the number of seconds the Stinger unit waits before forwarding to the MBONE router an IGMP version-2 Leave Group message it receives across a multicast client interface. Typically, these messages indicate that the IGMP group session can be cleared. However, a multicast interface in the Stinger unit can support many clients, some of which might establish multiple multicast sessions for identical groups, in which case a Leave Group message from a single client must be treated in a special way.

If `multicast-group-leave-delay` is set to zero (the default), the system forwards the Leave Group messages immediately.

If you set `multicast-group-leave-delay` to a nonzero value, the system does not immediately forward a Leave Group message it receives from a client on the interface. Instead, it sends back a query to make sure there are no clients on the interface with active multicast sessions for that group. If the Stinger unit receives a response before the specified `multicast-group-leave-delay` interval expires, it does not forward the Leave Group message. If the unit does not receive a response, it forwards the message and clears the IGMP group session from its tables after the specified interval.

If users might establish multiple multicast sessions for identical groups, set this parameter to a value from 10 to 20.

Configuring multicast client interfaces

A multicast client interface is an IP-enabled ATM PVC across a DSL interface to a CPE router such as a CellPipe®.

Overview of multicast client `ip-options` settings

In addition to the ATM and IP options needed to configure a terminating PVC, the following parameters, shown with default values, are used to enable IP multicast on a client interface:

```
[in CONNECTION/"":ip-options]
multicast-allowed = no
```

```
multicast-rate-limit = 100
multicast-group-leave-delay = 0
multicast-group-leave-delay-msec = 0
multicast-service-profile = ""
multicast-max-groups = 0
```

Parameter	RADIUS attribute	Setting
`multicast-allowed`	Ascend-Multicast-Client (155)	Enable/disable handling of IGMP requests and responses on the interface. The system does *not* forward multicast traffic on the basis of this setting.
`multicast-rate-limit`	Ascend-Multicast-Rate-Limit (152)	The rate at which the Stinger unit accepts multicast packets from clients on the interface. For example, if you set the rate to 5, the system accepts one packet every 5 seconds from multicast clients on the interface. Any subsequent packets received within that 5-second window are discarded. The default value of 100 disables multicast forwarding on the interface—the Stinger acts as a forwarder and handles IGMP packets, but does not accept packets from clients or forward multicast packets from the MBONE router. To enable multicast forwarding on the interface, you must set the rate to a number *less than* 100.
`multicast-group-leave-delay`	Ascend-Multicast-Grp-Leave-Delay (111)	Number of seconds to delay before forwarding a Leave Group message. The sum of (`multicast-group-leave-delay` × 1000) plus `multicast-group-leave-delay-msec` is the number of milliseconds the system waits before forwarding to the MBONE router an IGMP version-2 Leave Group message it receives across a multicast client interface. With the zero default values, the system forwards the Leave Group messages immediately. For client interfaces that support multiple multicast sessions to the same group, set these parameters to a value from 10 to 20. When these parameters have nonzero values, the system sends back a query to make sure there are no active multicast sessions on the interface for that group, and if it receives a response before the specified, specified delay expires, it does not forward the Leave Group message.
`multicast-group-leave-delay-msec`	Ascend-Multicast-GLeave-Delay-Msec	Number of milliseconds to add to the value of (`multicast-group-leave-delay` × 1000) to determine the total delay before forwarding a Leave Group message.

Parameter	RADIUS attribute	Setting
multicast-service-profile	Ascend-Multicast-Service-Profile-Name (274)	Name of a configured multicast services profile that filters multicast group access for this client interface. The filtering affects all new client links on the interface, and affects old client links after the expiration of query-response-interval. (See "Setting IGMP-v2 timers (local profiles only)" on page 2-12.) For information about configuring a multicast service profile, see "Managing multicast group memberships" on page 2-7.
multicast-max-groups	Ascend-Multicast-Max-Groups (275)	Maximum number of accessible multicast group (from 0 to 512) for this client interface. You can set this to a lower number to limit multicast traffic to the interface. This value of this parameter limits the activation of new client links, but does not affect the existing client links.

Setting IGMP-v2 timers (local profiles only)

The following parameters, shown with default values, are used to configure the timers defined in RFC 2236, *Internet Group Management Protocol Version 2*, on multicast client interfaces:

```
[in CONNECTION/"":ip-options:igmp-options]
robust-count = 2
query-interval = 125
query-response-interval = 100
last-member-query-interval = 10
last-member-query-count = 2
```

Parameter	Setting
robust-count	A threshold of packet losses (from 2 to 10) up to which the multicast subsystem will remain robust. If the interface is expected to have a high rate of packet loss, increase this value. IGMP is robust to (robust-count minus 1) packet losses. It cannot be set to zero and should not be set to 1. The default is 2.
query-interval	Number of seconds (from 0 to 1024) between general queries. You can increase this value from its default of 125 seconds to reduce the number of IGMP queries sent on the interface. The sum of query-interval plus (query-response-interval divided by 10) must be less than the value of multicast-member-timeout.

Parameter	Setting
`query-response-interval`	Maximum response time in tenths of a second (from 0 to 1024) inserted into general queries. You can increase this value from its default of 10 seconds to make IGMP traffic less bursty, because host responses will be spread out over a larger interval. The number of seconds response time (this value divided by 10) must be less than the `query-interval` value.
`last-member-query-interval`	Maximum response time in tenths of a second (from 0 to 1024) inserted into group-specific queries sent in response to Leave Group messages. You can reduce this value from its default of 1 second to reduce the time it takes to detect that the last member of a group has left. The response time (this value divided by 10) must be less than the `query-interval` value.
`last-member-query-count`	Number of group-specific queries sent before the multicast router assumes there are no local members.

OSPF Routing

3

Introduction to OSPF

Open Shortest Path First (OSPF) is a next-generation Internet routing protocol. OSPF is defined by RFC 2328, *OSPF Version 2*. The *Open* in its name refers to the fact that OSPF was developed in the public domain as an open specification. *Shortest Path First* refers to an algorithm developed by Dijkstra in 1978 for building a self-rooted shortest-path tree from which routing tables can be derived. (For a description of the algorithm, see "Link-state routing algorithm" on page 3-8.)

RIP limitations solved by OSPF

The rapid growth of the Internet has pushed Routing Information Protocol (RIP) beyond its capabilities, particularly in the areas of distance-vector metrics, the 15-hop limitation, and slow convergence due to excessive routing traffic.

Distance-vector metrics

RIP is a distance-vector protocol, which uses a hop count to select the shortest route to a destination network. RIP always uses the lowest hop count, regardless of the speed or reliability of a link.

In contrast, OSPF is a link-state protocol, which can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network.

15-hop limitation

With RIP, a destination that requires more than 15 consecutive hops is considered unreachable, and this limitation inhibits the maximum size of a network. OSPF has no hop limitation. You can add as many routers to a network as you want.

Excessive routing traffic and slow convergence

RIP creates a routing table and then propagates it throughout the internetwork of routers, hop by hop. The time required for all routers to receive information about a topology change is called *convergence*. Slow convergence can result in routing loops and errors.

A RIP router broadcasts its routing table every 30 seconds. On a 15-hop network, convergence can be as high as 7.5 minutes. In addition, a large table can require multiple broadcasts for each update, which consumes a lot of bandwidth. OSPF uses a topological database to represent the network and propagates only changes to the database. (For more information about propagation, see “Exchange of routing information” on page 3-4.)

TAOS implementation of OSPF

The primary goal of the OSPF implementation is to allow a Stinger unit to communicate with other routers within a single autonomous system (AS).

Limited border router capability

A Stinger unit acts as an OSPF internal router with limited border router capability.

The Stinger unit does not currently function as an interior gateway protocol (IGP) gateway, although it performs autonomous system border router (ASBR) calculations for external routes (such as WAN links that do not support OSPF). The Stinger unit imports external routes into its OSPF database and flags them as autonomous system external (ASE). It redistributes these routes via OSPF ASE advertisements, and propagates its OSPF routes to remote WAN routers running RIP.

Authentication

The Stinger unit supports null authentication, simple password authentication, and message-digest MD5) cryptographic authentication. For details, see “Security” on page 3-3.

One active IP interface per port

The TAOS OSPF implementation conforms to RFC 2328, it does not support virtual IP interfaces. If more than one IP address is assigned to the same physical port, only one of the logical interfaces can have OSPF enabled. For example, in the following listing the first port on the Ethernet card in slot 15 (shelf 1, slot 15, port 1) has three virtual interfaces:

```
admin> dir ip-interface
18  02/28/2002 10:47:30  { { any-shelf any-slot 0 } 0 }
28  02/28/2002 10:49:14  { { shelf-1 first-control-module 1 } 0 }
18  02/28/2002 10:47:30  { { shelf-1 second-control-module 1 } 0 }
```

OSPF can be enabled on any one of the port’s IP interfaces, but not on more than one interface for the same port.

OSPF diagnostic commands

The `ospf` diagnostic-level commands enable an administrator to display information related to OSPF routing, including the Link-State Advertisements (LSAs), border router information, and the OSPF areas, interfaces, statistics, and routing table. See the *Stinger Administration Guide* for additional information.

OSPF traps

The Stinger unit supports OSPF traps as defined in RFC 1850, *OSPF Version 2 Management Information Base*. For an OSPF trap to be generated when the trap condition occurs, you must enable OSPF traps, either in the `trap` profile or by setting the corresponding bit in the `ospfSetTrap` MIB object, which is defined in RFC 1850. In addition, the individual trap that represents the trap condition must be enabled. See the *Stinger Administration Guide* for additional information.

OSPF features

This section provides a brief overview of OSPF routing to help you configure the Stinger unit properly. For details about how OSPF works, see RFC 2328, *OSPF Version 2*.

An autonomous system is a group of OSPF routers exchanging information, typically under the control of one company. An autonomous system can include a large number of networks, all of which are assigned the same autonomous system number. All information exchanged within the autonomous system is *interior*.

Exterior gateway protocols (EGPs) are used to exchange routing information between autonomous systems. The autonomous system number can be used by border routers to filter out certain EGP routing information. OSPF can make use of EGP data generated by other border routers and added into the OSPF system as ASE information, as well as static routes configured locally or in RADIUS.

Security

All OSPF protocol exchanges are authenticated by null authentication. Only trusted routers can participate in the autonomous system's routing. A variety of authentication schemes can be used. In fact, different authentication types can be configured for each area. For a discussion of areas, see "Hierarchical routing (areas)" on page 3-6.

Authentication provides added security for the routers that are on the network. Routers that do not have the password are unable to gain access to the routing information, because authentication failure prevents a router from forming adjacencies. (For a discussion of adjacencies, see "Exchange of routing information" on page 3-4.) If both sides of a connection do not support the same authentication method, packet error messages can result.

In addition to null and simple password authentication, Stinger units support the MD5 cryptographic authentication method for OSPF as described RFC 2328. For details about MD5 encryption, see RFC 2328.

Support for variable-length subnet masks

OSPF routers handle variable-length subnet masks (VLSMs). Each route distributed by OSPF has a destination address and subnet mask, and two different subnets of the

same IP network can use different size subnet masks. A packet is routed to the best (longest or most specific) match. Host routes are considered to be subnets whose masks are all ones (0xFFFFFFFF).

Note OSPF is very useful for networks that make use of VLSMs. However, to prevent excessive link-state calculations by all OSPF routers on the network, make every effort to assign subnets that are as continuous as possible.

Exchange of routing information

OSPF stores its information about the network in a topological database and propagates only changes to the database. Selected neighboring routers form relationships, referred to as *adjacencies*, for the purpose of exchanging routing information. Not every pair of neighboring routers become adjacent. Routers connected by point-to-point networks and virtual links always become adjacent. On multiaccess networks, all routers become adjacent to routers identified as the designated router (DR) and the backup designated router (BDR).

As the adjacency is established, the neighbors exchange databases and build a consistent, synchronized database between them. When an OSPF router detects a change on one of its interfaces, it modifies its topological database and multicasts the change to its adjacent neighbors, which in turn propagate the change to their adjacent neighbors, until all routers within an area have synchronized topological databases. This process provides quick convergence among routers.

A link state advertisement (LSA) is a packet that describes various aspects of an OSPF route. Each LSA is flooded throughout the routing domain. The collected link state advertisements of all routers and networks forms the OSPF topological database. The LSA types include:

- Type 1 (RTR) router—Describes the collected states of the router's interfaces.
- Type 2 (NET) network—Describes the set of routers attached to the network.
- Types 3 and 4 (STUB) summary—Describes point-to-point routes to networks or area border routers (ABRs).
- Type 5 (ASE) AS-external—Describes routes to destinations external to the autonomous system. An AS-external-LSA can also describe a default route for the autonomous system. For example, other routers send LSAs to only the designated router by using the All-Designated-Routers multicast address of 224.0.0.6.
- Type 7 (ASE) NSSA—For not-so-stubby-areas (NSSAs), all routes imported to OSPF have the P-bit set (P stand for propagate). When the P-bit is enabled, ABRs translate type 7 LSAs to type 5 LSAs, which can then be transmitted to the backbone. These external routes are considered type 7 LSAs

Designated and backup designated routers

In OSPF terminology, a broadcast network is any network that has more than two OSPF routers attached and supports the capability to address a single physical message to all the attached routers. Figure 3-1 shows such a network.

Figure 3-1. OSPF broadcast network

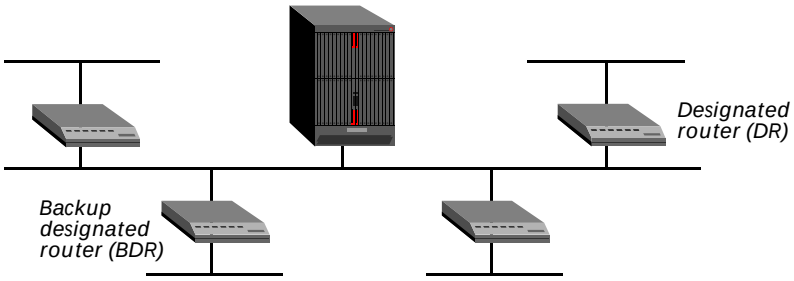

To reduce the number of adjacencies each router must form, OSPF calls one of the routers the designated router. As routers begin to form adjacencies, they elect a designated router and then all other routers on the network establish adjacencies, primarily with the designated router. This process simplifies the routing table update procedure and reduces the number of link-state records in the database. The designated router plays other important roles as well to reduce the overhead of OSPF link-state procedures.

To prevent the designated router from becoming a serious liability to the network if it fails, OSPF routers also elect a backup designated router at the same time. Other routers maintain adjacencies with both the designated router and its backup, but the backup router leaves as many of the processing tasks as possible to the designated router. If the designated router fails, the backup immediately becomes the designated router and a new backup is elected.

You choose the designated router on the basis of the processing power, speed, and memory of the system, then assign priorities to other routers on the network in case the backup designated router is also down at the same time.

Note The Stinger unit can function as a designated router or backup designated router. However, many sites choose to assign a LAN-based router for these roles to dedicate the Stinger unit to WAN processing.

Configurable cost metrics

You assign a cost to the output side of each router interface. The lower the cost, the more likely the interface is to be used to forward data traffic. Costs can also be associated with the externally derived routing data.

The OSPF cost can also be used for preferred-path selection. If two paths to a destination have equal costs, you can assign a higher cost to one of the paths so it can be a backup to be used only when the primary path is not available.

Figure 3-2 shows how costs are used to direct traffic over high-speed links. For example, if Router-2 in Figure 3-2 receives packets destined for Host B, it routes them through Router-1 across two T1 links (`Cost=20`) rather than across one 56Kbps B channel to Router-3 (`Cost=240`).

Figure 3-2. OSPF costs for different types of links

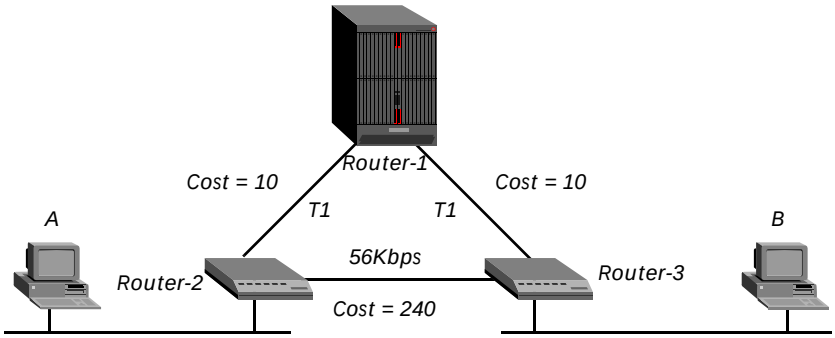

The Stinger unit has a default cost of 1 for a connected route (Ethernet) and 10 for a WAN link. If two paths have the same destination, the Stinger unit uses the path with the lower cost unless route preferences change the equation. For information about route preferences, see “Setting routing protocol options” on page 1-41.

When assigning costs, remember to account for the bandwidth of a connection. For example, for a single B-channel connection, the cost is be 24 times greater than for a T1 link.

Note Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.

Hierarchical routing (areas)

If a network becomes too large, the size of the database, time required for route computation, and related network traffic become excessive. You can partition an autonomous system into areas to provide hierarchical routing, with a backbone area connecting the other areas. The backbone area is special and always has the area number 0.0.0.0. The backbone consists of networks not contained in any area, their attached routers, and routers that belong to multiple areas.

The backbone must be contiguous. You can use virtual links to connect two backbone routers that have an interface to a common nonbackbone area. OSPF treats two routers joined by a virtual link as if they were connected by an unnumbered point-to-point network. The backbone distributes routing information between areas and has all the properties of an area. The topology of the backbone is invisible to each of the areas, while the backbone itself has no information about area topology.

Each area acts as its own network: All area-specific routing information stays within the area, and all routers within an area must have a synchronized topological database. To tie the areas together, some routers belong to the backbone area and also to one of the other areas. These routers are area border routers (ABRs). In Figure 3-3, all the routers are ABRs.

Figure 3-3. Dividing an OSPF autonomous system into areas

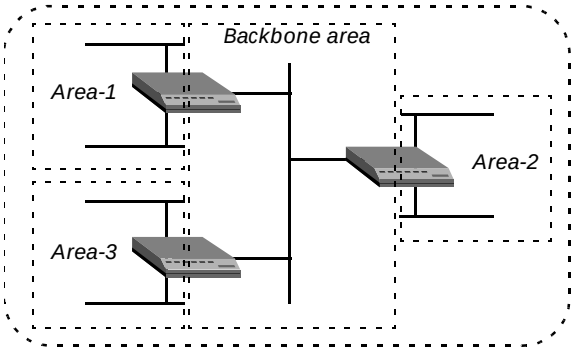

With the ABRs and area boundaries set up correctly, link-state databases are unique to an area. You can configure the Stinger unit to route in normal areas, stub areas, and NSSAs. These different kinds of areas handle the autonomous system external (ASE) routes originated by ASBRs in the following ways.

Normal areas

An OSPF normal area allows type 5 LSAs to be flooded throughout the area.

Stub areas

Areas that are connected only to the backbone area by one ABR have one exit point and need not maintain information about external routes. To reduce the cost of routing, OSPF supports stub areas in which a default route summarizes all external routes. A stub area allows no type 5 LSAs to be propagated into or throughout the area, and instead depends on default routing to external destinations.

NSSAs

NSSAs are like stub areas in that they do not receive or originate type 5 LSAs. However, NSSAs rely solely on default routing for external routes. They employ type 7 LSAs for carrying ASE route information within the area. Type 7 LSAs use a P-bit to flag the NSSA border router to translate the type 7 LSA into a type 5 LSA, which can then be propagated into other areas.

When the Stinger unit is routing OSPF in an NSSA, it imports ASE routes defined in local or RADIUS profiles as type 7 LSAs. These imported ASE LSAs always have the P-bit enabled, which flags border routers to translate them into type 5 LSAs.

You can list the router IDs of NSSA border routers that are translating type 7 LSAs to type 5 LSAs, by entering the `ospf translators` command. For example:

```
admin> ospf translators
Area ID     Router ID
0.0.0.1     10.105.0.13
0.0.0.2     12.1.1.1
```

Note For details about the NSSA specification, see RFC 1587, *The OSPF NSSA Option*.

Link-state routing algorithm

The link-state routing algorithms require that all routers within a domain maintain synchronized (identical) topological databases, and that the databases describe the complete topology of the domain. An OSPF router's domain can be an autonomous system or an area within an autonomous system.

OSPF routers create and update a link-state database from information exchanged with other routers. Link-state databases are synchronized between pairs of adjacent routers (as described in "Exchange of routing information" on page 3-4). In addition, each OSPF router uses its link-state database to calculate a self-rooted tree of shortest paths to all destinations. The routing table is built from these calculated shortest-path trees. For example, consider the network topology in Figure 3-4.

Figure 3-4. Sample OSPF topology

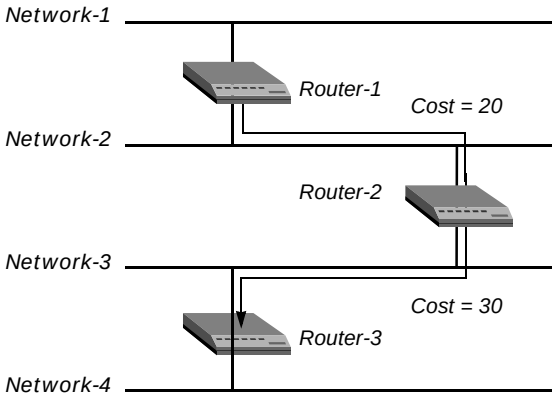

Table 3-1 shows the relevant information in the routers' link-state databases.

Table 3-1. Link-state databases for OSPF topology in Figure 3-4

Router-1	Router-2	Router-3
Network-1/Cost 0	Network-2/Cost 0	Network-3/Cost 0
Network-2/Cost 0	Network-3/Cost 0	Network-4/Cost 0
Router-2/Cost 20	Router-1/Cost 20	Router-2/Cost 30
	Router-3/Cost 30	

From the link-state database, each router builds a self-rooted shortest-path tree, and then calculates a routing table stating the shortest path to each destination in the autonomous system. (See Table 3-2, Table 3-3, and Table 3-4.) The table also includes externally derived routing information.

All the routers calculate a routing table of shortest paths, based on the link-state database. Externally derived routing data is advertised throughout the autonomous system but is kept separate from the link-state data. Each external route can also be tagged by the advertising router, enabling the passing of additional information between routers on the boundary of the autonomous system.

Table 3-2. Shortest-path tree and resulting routing table for Router-1

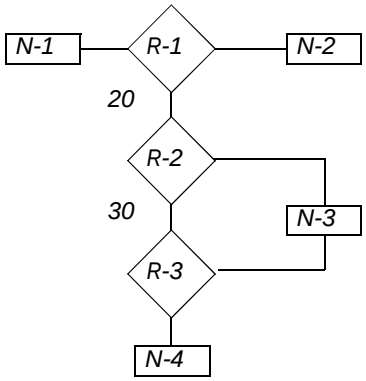

Destination	*Next hop*	*Metric*
Network-1	Direct	0
Network-2	Direct	0
Network-3	Router-2	20
Network-4	Router-2	50

Table 3-3. Shortest-path tree and resulting routing table for Router-2

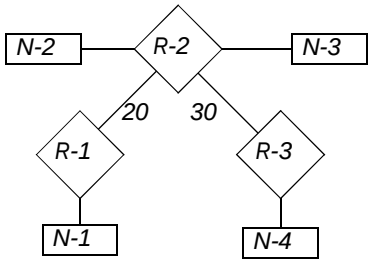

Destination	*Next hop*	*Metric*
Network-1	Router-1	20
Network-2	Direct	0
Network-3	Direct	0
Network-4	Router-2	30

Table 3-4. Shortest-path tree and resulting routing table for Router-3

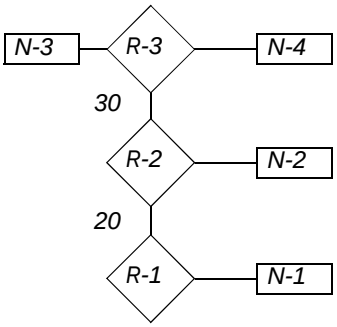

Destination	*Next hop*	*Metric*
Network-1	Router-2	50
Network-2	Router-2	30
Network-3	Direct	0
Network-4	Direct	0

Adding a Stinger unit to an OSPF network

Before it can run OSPF, a Stinger unit must be configured for IP routing, as described in “IP Routing” on page 1-1.

System reset requirement

After enabling OSPF routing, you must reset the system. The only time the system brings up OSPF routing on an interface is following a reset. As the system comes up

with OSPF enabled on one or more interfaces, it begins to form adjacencies and build its routing table.

If you change the value of the `pool-ospf-adv-type` parameter in the `ip-global` profile, you must reset the system for the change to take effect.

If you change the OSPF `area-type` from `normal` to `nssa` or vice versa, you must reset the system for the change to take effect.

Overview of LAN and WAN OSPF settings

The OSPF subprofiles of the `ip-interface` and `connection` profiles (for configuring local and WAN interfaces, respectively) both contain the same parameters. Following are the OSPF parameters, shown with their default values:

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:ospf (new)]
active = no
area = 0.0.0.0
area-type = normal
hello-interval = 10
dead-interval = 40
priority = 5
authen-type = simple
auth-key = *******
key-id = 0
cost = 1
down-cost = 16777215
ase-type = type-1
ase-tag = c0:00:00:00
transit-delay = 1
retransmit-interval = 5
non-multicast = no
network-type = Broadcast
poll-interval = 10
md5-auth-key = *******

[in CONNECTION/"":ip-options:ospf-options (new)]
active = no
area = 0.0.0.0
area-type = normal
hello-interval = 30
dead-interval = 120
priority = 5
authen-type = simple
auth-key = *******
key-id = 0
cost = 10
down-cost = 1000
ase-type = type-1
ase-tag = c0:00:00:00
transit-delay = 1
retransmit-interval = 5
non-multicast = no
```

```
network-type = Point-to-Point
poll-interval = 10
md5-auth-key = *******
```

Parameter	Setting
active	Enable/disable OSPF on an interface.
area	Area number in dotted decimal notation. The default area number is 0.0.0.0, which represents the OSPF backbone. Note that area numbers are not IP addresses, although they use a similar format. For a discussion of areas, see "Hierarchical routing (areas)" on page 3-6.
area-type	Type of area. The default is the normal area type, in which external routes are advertised throughout the autonomous system.
hello-interval	Number of seconds between Hello packets. Specify a value from 1 through 65535. In the ip-interface profile, the default setting is 10. In the connection profile, the default setting is 30.
dead-interval	Number of elapsed seconds without receiving a Hello packet the router will wait before considering its neighbor inoperable and instituting a link-state change. Specify a value from 0 through 65535. In the ip-interface profile, the default setting is 40. In the connection profile, the default setting is 120.
priority	Priority value, used to elect a designated router and backup designated router. Specify a value from 0 through 255. A setting of 0 excludes the Stinger unit from becoming a designated router or backup designated router. A setting of 1 or greater places the Stinger unit on the list of possible designated routers. The higher the priority value of the Stinger unit relative to other OSPF routers on the network, the better the chances that it will become one of these routers. For details, see "Designated and backup designated routers" on page 3-4.
authen-type	Type of authentication to use for validating OSPF packet exchanges. Specify one of the following values: ■ none—No authentication is required. ■ simple (the default)—The router uses the password supplied in the auth-key parameter to validate OSPF packet exchanges. ■ md5—The router uses MD5 encryption and the authentication key ID supplied by the key-id parameter to validate OSPF packet exchanges. For related information, see "Security" on page 3-3.
auth-key	Secret key for authenticating traffic in the router's area. Enter a text string of up to 8 characters. When authen-type is set to md5, you must set the md5-auth-key parameter to specify a key.

Parameter	Setting
key-id	Number from 0 to 255, used to encrypt the secret key when authen-type is set to md5.
cost	Cost of routing to the interface. The lower the cost assigned to a route, the more likely it is to be used to forward traffic. For details, see "Configurable cost metrics" on page 3-5. Specify a value from 1 through 16777215. For the Ethernet (ip-interface profile), the default setting is 1. For the WAN (connection profile), the default setting is 10.
down-cost	Cost to be applied to the interface when it is unavailable. The output cost when the link is physically unavailable but virtually active. Specify a value from 1 through 16777215. For the Ethernet (ip-interface profile), the default setting is 16777215. For the WAN (connection profile), the default setting is 1000.
ase-type	Type of metric to apply to routes learned from RIP. ■ type-1 (the default)—Expresses the metric in the same units as the interface cost. ■ type-2—Metric is larger than any link-state path. This parameter applies in a connection profile only when OSPF is *not* active
ase-tag	Hexadecimal number that appears in management utilities and flags the route as external. It can also be used by border routers to filter a record. This parameter applies in a connection profile only when OSPF is *not* active.
transit-delay	Estimated number of seconds required to transmit a Link State Update packet over this interface, taking into account transmission and propagation delays. On a connected route, you can leave the default of 1.
retransmit-interval	Number of seconds between LSA retransmissions for adjacencies belonging to this interface. Its value is also used when retransmitting database description and link-state request packets. On a typical connected route, accept the default setting of 5.

Parameter	Setting
non-multicast	Enable/disable a Stinger unit to run OSPF over a frame relay link to a GRF® switch. GRF® handles frame relay as a nonbroadcast multiaccess (NBMA) network, while the Stinger unit handle frame relay as a serial (point-to-point) network. If non-multicast is set to yes, all multicast packets are remapped to a directed neighbor address, which enables adjacencies to form between neighbors. This setting is ignored on an Ethernet broadcast network. **Caution** Lucent Technologies recommends that you do not enable this parameter for unnumbered interfaces, because it causes these interfaces to drop packets. **Caution** Do not enable this parameter for NBMA networks. (See "OSPF NBMA support" on page 3-23).
network-type	Type of network attachment. Specify one of the following values: ■ broadcast—Any broadcast-capable network, such as Ethernet. ■ nonbroadcast—An OSPF NBMA network. An NBMA network has multiple points of access (more than two routers) and does not support broadcast capability. Frame relay and X.25 are typically NBMA networks. ■ Point-to-Point—An interface connected to one other node on the remote end. For the Ethernet (ip-interface profile), the default setting is broadcast. For the WAN (conenction profile), the default setting is Point-to-Point.
poll-interval	Poll interval, in seconds, for NBMA networks. Specify a value from 1 through 65535. The default value is 10. For NBMA details, see "OSPF NBMA support" on page 3-23.
md5-authen-key	Secret key to be used for the MD5 cryptographic authentication method, up to 16 characters. The default value is ascend0. When authen-type is set to md5, you must supply a value for this parameter.

Example of configuring a LAN OSPF interface

Figure 3-5 shows three OSPF routers in the backbone area of an autonomous system. Because all OSPF routers are in the same area, the units form adjacencies and synchronize their databases. This example shows how to configure the LAN interface of the unit labeled TAOS-2.

Figure 3-5. OSPF on a LAN interface

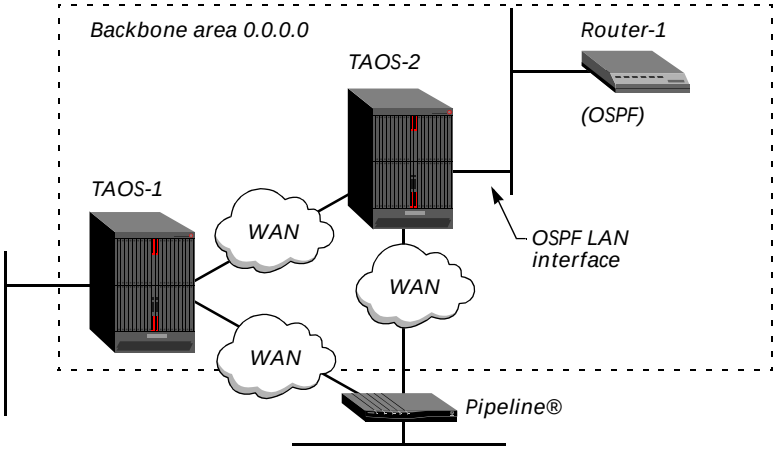

All OSPF routers in Figure 3-5 have RIP turned off. Running both RIP and OSPF is unnecessary, and turning RIP off reduces processor overhead. OSPF can learn routes from RIP interfaces, incorporate them in the routing table, assign them an external metric, and tag them as external routes.

Although RFC 2328 does not specify a limitation for the number of routers in the backbone area, keep the number of routers relatively small, because changes that occur in area zero are propagated throughout the autonomous system. Another way to configure the same units is to create a second area (such as 0.0.0.1) in one of the existing OSPF routers, and add the Stinger unit to that area. You can then assign the same area number (0.0.0.1) to all OSPF routers reached through the Stinger unit across a WAN link.

The following sample commands show how to configure TAOS-2 in Figure 3-5. The commands assign the IP address 10.168.8.17/24 to the local interface and configure the OSPF router in the backbone area:

```
admin> read ip-int {{ 1 12 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } read

admin> set ip-address = 10.168.8.17/24

admin> set rip-mode = routing-off

admin> set ignore-def-route = yes

admin> set ospf active = yes

admin> write
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } written
```

The following sample commands show how to configure the IP interface for MD5 authentication:

```
admin> read ip-interface { { 1 12 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } read

admin> set ospf authen-type = md5

admin> set ospf md5-auth-key = 12!secret*34key

admin> write
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } written
```

Example of configuring WAN OSPF interfaces

This example shows how to configure `connection` profiles in the Stinger units shown in Figure 3-6, to enable them to route OSPF across the WAN that separates them. In this example, the unit labeled TAOS-1 has the IP address 10.2.3.4/24, and the unit labeled TAOS-2 has the address 10.168.8.17/24.

Figure 3-6. OSPF on a WAN interface

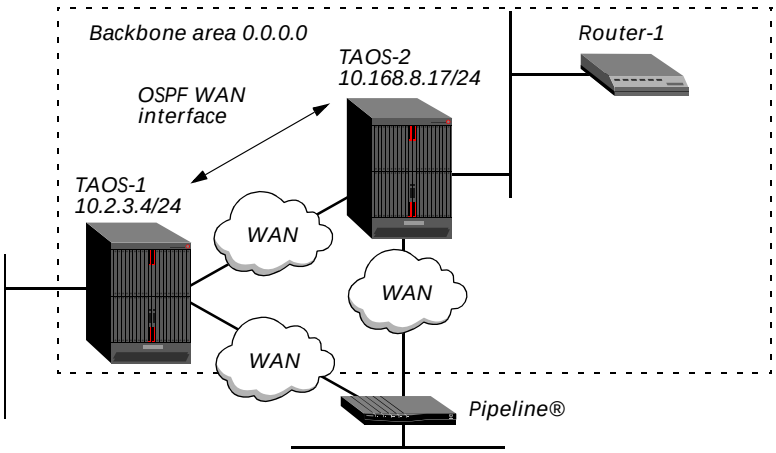

The WAN interfaces of the Stinger unit form point-to-point networks, because each link joins a single pair of routers. Point-to-point networks typically do not provide a broadcasting or multicasting service, so all advertisements are sent point to point.

The following commands configure the OSPF WAN link in TAOS-1 in Figure 3-6:

```
admin> read conn taos2ink
CONNECTION/taos2link read
admin> set ip-options remote = 10.168.8.17/24
admin> set ip-options rip = routing-off
admin> set ip-options ospf active = yes
admin> write
CONNECTION/taos2link written
```

The following commands configure the OSPF WAN link in TAOS-2 in Figure 3-6:

```
admin> read conn taos1link
CONNECTION/taos1link read
admin> set ip-options remote = 10.2.3.4/24
admin> set ip-options rip = routing-off
admin> set ip-options ospf active = yes
admin> write
CONNECTION/taos1link written
```

Example of integrating a RIP-2 interface

In Figure 3-7, each Stinger has a WAN interface to a remote Pipeline® unit. The Pipeline® is an IP router that supports RIP-2, and has the IP address 10.6.7.168/24.

The route to the Pipeline® LAN, and any routes the Stinger learns about from the remote Pipeline®, are ASE routes (external to the OSPF autonomous system).

Figure 3-7. Including ASE routes in the OSPF environment

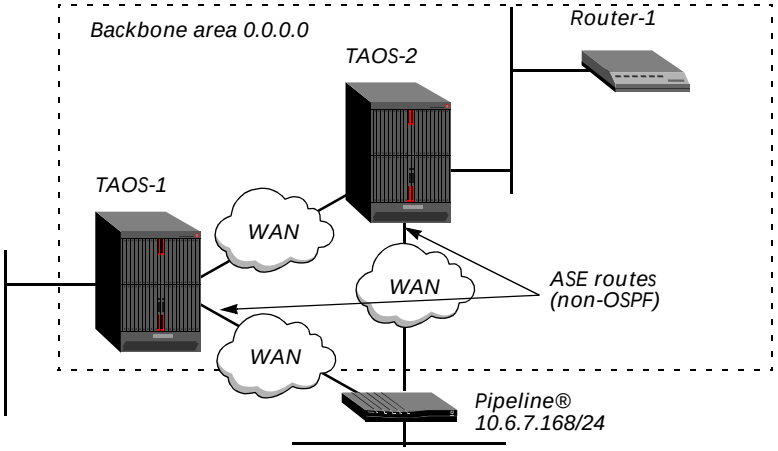

To enable OSPF to add routes learned from RIP-2 to the routing table, you can configure RIP-2 normally in the `connection` profiles. The global `rip-ase-type` parameter in the `ip-global` profile determines the ASE metric applied when the routes are imported to OSPF. For details about `rip-ase-type`, see "Configuring route options" on page 3-18.

However, in the following example, RIP is turned off on the link, so the Stinger unit does not forward or receive routing updates on the interface. The following sample commands specify a cost of `240` for the link to the Pipeline®, disable RIP, and specify ASE information for the `connection` profile's static route:

```
admin> read conn pipeline1
CONNECTION/pipeline1 read
admin> set ip-options remote = 10.6.7.168/24
admin> set ip-options rip = routing-off
admin> set ip-options ospf active = no
admin> set ip-options ospf cost = 240
admin> set ip-options ospf ase-type = type-2
admin> set ip-options ospf ase-tag = cfff8000
admin> write
CONNECTION/pipeline1 written
```

The `ase-type` and `ase-tag` settings causes the OSPF router to import the route to 10.6.7.168/24 as a type 2 LSA and tag it with the specified hexadecimal number. The cost assigned is appropriate for the bandwidth of a single B-channel connection, and the cost is 24 times greater than for a T1 link.

Configuring an OSPF area range

An OSPF area is defined by a list of address ranges. Networks and hosts belong to an area if their addresses fall into one of the area's defining address ranges. Routers can belong to multiple areas, depending on their attached networks' area membership.

A network subnet can be its own OSPF area. You can configure the area to be a single address range whose IP address is the address of the network subnet with a class A, B, or C address mask. The Stinger unit advertises a single route for the area, describing the entire subnetted network.

You define an OSPF range area by setting the parameters in the `ospf-area-range` profile, shown here with default settings.

```
[in OSPF-AREA-RANGE/"" (new)]
name* = ""
area-id = 0.0.0.0
area-network-addr = 0.0.0.0/0
area-network-mask = 0.0.0.0
advertize = no
```

Parameter	Setting
`name`	Name for this OSPF area range. Enter a text string of up to 31 characters.
`area-id`	OSPF area ID for this area range. Area numbers are not IP addresses, although they use a similar format. For a discussion of areas, see "Hierarchical routing (areas)" on page 3-6.
`area-network-addr`	Enter an address in dotted decimal notation. If the area represents a network subnet, enter the IP network number of the network subnet.
`area-network-mask`	OSPF area network mask. Specify a netmask in dotted decimal notation.
`advertize`	Enable/disable the advertisement of this area range. Specify `yes` or `no` (the default). Routing information is condensed at area boundaries. Unadvertised ranges allow certain networks to be intentionally hidden from other areas.

Configuring virtual links

You can configure virtual connections between any pair of area border routers having interfaces in a common nonbackbone area, which is referred to as the virtual link's *transit area*. You create a virtual link by configuring each router with the router ID of the virtual link's *other* end point and the transit area through which the virtual link runs. You cannot configure virtual links through stub areas.

A virtual link appears in router LSAs (for the backbone) as a separate router interface to the backbone. This interface has parameters associated with a router interface (see "Overview of LAN and WAN OSPF settings" on page 3-10).

A virtual link behaves like an unnumbered point-to-point link, but it has an associated IP interface address. The system uses this address as the IP source in OSPF protocol packets it sends along the virtual link.

The interface output cost for a virtual link is the cost of the path between the two routers, and is set dynamically during the routing table build process.

The `ospf-virtual-link` profile contains the parameters that you use for configuring a virtual link, shown here with default values.

```
[in OSPF-VIRTUAL-LINK/0.0.0.0 (new)]
neighbor-router-id* = 0.0.0.0
transit-area-id = 0.0.0.0
rexmit-delay = 5
xmit-delay = 1
hello-interval = 30
dead-interval = 120
authen-type = simple
authen-key = ascend0
key-id = 0
md5-authen-key = ascend0
```

Parameter	Setting
neighbor-router-id	Router ID of the other end point router of the virtual link. Specify an IP address or netmask, in dotted decimal notation.
transit-area-id	Address for the transit area.
rexmit-delay	Number of seconds between LSA retransmissions for adjacencies belonging to this virtual link. Its value is also used when retransmitting database description and link-state request packets. The default setting is 5.
xmit-delay	Estimated number of seconds required to transmit a Link State Update packet over this virtual interface, taking into account transmission and propagation delays. The default is 1. Specify a number from 0 through 65535.
hello-interval	See "hello-interval" on page 3-11. The default setting is 30.
dead-interval	See "dead-interval" on page 3-11. The default setting is 120.
authen-type	See "authen-type" on page 3-11. The default setting is simple.
authen-key	See "auth-key" on page 3-11. The default setting is ascend0.
key-id	See "key-id" on page 3-12. The default setting is 0.
md5-authen-key	See "md5-authen-key" on page 3-13. The default setting is ascend0.

Configuring route options

The ip-globalprofile contains several settings that apply only when OSPF routing is in use. Following are the relevant parameters, shown here with their default settings: Parameter descriptions follow the profile listings.

```
[in IP-GLOBAL]
pool-summary = no
pool-chaining = no
pool-ospf-adv-type = type-1
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 +
pool-name = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" ""+
rip-pref = 100
rip-queue-depth = 50
ospf-pref = 10
```

```
ospf-ase-pref = 150
ospf-global = { no yes 0 }
rip-tag = c8:00:00:00
rip-ase-type = 1
iproute-cache-enable = yes
iproute-cache-size = 0
ipport-cache-enable = yes
```

```
[in IP-GLOBAL:ospf-global]
enable = no
as-boundary-router = yes
ospf-max-lsa = 0
```

Parameter	Setting
pool-summary	See "Settings in the ip-global profile" on page 1-60.
pool-chaining	See "Pool chaining in local profiles" on page 1-67.
pool-ospf-adv-type	Type of ASE metric applied to summarized pools imported into OSPF as external routes. The pool-summary parameter must be set to yes and OSPF must be enabled for this setting to have any effect. Specify one of the following values: ■ type-1 (the default)—Expresses the metric in the same units as the interface cost. ■ type-2—Metric is larger than any link-state path. ■ internal—Imports pool routes as intra-area routes, which enables them to work with stub areas.
pool-base-address	See "Settings in the ip-global profile" on page 1-60.
assign-count	See "Settings in the ip-global profile" on page 1-60.
pool-name	See "Settings in the ip-global profile" on page 1-60.
rip-pref	See "Setting routing protocol options" on page 1-41.
rip-queue-depth	See "Setting routing protocol options" on page 1-41.
ospf-pref	Preference value for routes learned from OSPF. Valid values are 0 to 255. The default value is 10.
ospf-ase-pref	Preference value for routes learned from RIP, ICMP, or another non-OSPF protocol. Specify a value from 0 through 255. By default, routes learned dynamically from another routing protocol are assigned a preference value of 150.
rip-tag	Hexadecimal number associated with routes learned from RIP. OSPF border routers can use the tag to filter a record.
rip-ase-type	Type of ASE metric applied to routes learned from RIP. ■ type-1 (the default)—Expresses the metric in the same units as the interface cost. ■ type-2—Metric is larger than any link-state path.
IP-GLOBAL:ospf-global: enable	Enable/disable OSPF routing protocol for the Stinger unit. The default setting is no.

Parameter	Setting
IP-GLOBAL:ospf-global: as-boundary-router	Enable/disable autonomous system border router (ASBR) calculations related to external routes. Normally, when the Stinger unit imports external routes from RIP (for example, when it establishes a WAN link with a caller that does not support OSPF), it performs the ASBR calculations for those routes. If necessary, you can prevent the Stinger unit from performing ASBR calculations by setting as-boundary-router to no.
IP-GLOBAL:ospf-global: ospf-set-trap	Enable/disable OSPF traps. This parameter can be set using SNMP. Specify a hexadecimal value of 4 bytes.
IP-GLOBAL:ospf-global: ospf-max-lsa	Maximum number of LSAs allowed in the link-state database. Specify a number from 0 through 4294967295. The default setting is 0.

Example of importing a summarized pool as an ASE

For information about defining summarized address pools, see "Example of configuring summarized address pools" on page 1-64. The following commands configure a summarized pool and import it to OSPF with a type 1 OSPF metric:

```
admin> read ip-global
IP-GLOBAL read
admin> set pool-summary = yes
admin> set pool-base-address 1 = 10.12.253.1
admin> set assign-count 1 = 62
admin> set pool-ospf-adv-type = type-1
admin> write
IP-GLOBAL written
```

When pool-summary is set to yes and OSPF is enabled, the OSPF subsystem uses the pool-ospf-adv-type parameter to determine how to import summarized routes into OSPF. If this parameter is set to type-1, the metric for the route to a summarized pool is expressed in the same units as the link-state metric (interface cost).

If pool-ospf-adv-type is set to type-2, the unit considers the routing between autonomous systems as the major cost of routing a packet, and conversion of external costs to internal link-state metrics is unnecessary. If the parameter is set to internal, the summarized pool addresses are imported into OSPF as intra-area routes, which enables them to work properly with stub areas.

Example of setting ASE preferences

The ospf-pref and ospf-ase-pref parameters determine the preference values assigned to routes learned from other OSPF routers and those imported from other dynamic routing protocols. The default settings place a much lower preference on OSPF routes, which means that the routes learned from other protocols (ASE routes) are more likely to be used. The following commands decrease to 100 the preference assigned to ASE routes (the default is 150):

```
admin> read ip-global
IP-GLOBAL read
```

```
admin> set ospf-ase-pref = 100
admin> write
IP-GLOBAL written
```

Configuring OSPF static-route information

The following parameters in the ip-route profile (shown with sample settings), apply only when OSPF is enabled:

```
in IP-ROUTE/"" (new)]
cost = 1
third-party = no
ase-type = type-1
ase-tag = c0:00:00:00
ase7-adv = N/A
```

Parameter	Setting
cost	Cost of routing to the interface. The lower the cost assigned to a route, the more likely it is to be used to forward traffic. See "Configurable cost metrics" on page 3-5.
third-party	Enable/disable advertisement of routes to external destinations on behalf of another gateway (a third party). See "Example of specifying a third-party route" on page 3-23.
ase-type	Type of metric to apply to routes learned from RIP. ■ type-1 (the default)—Expresses the metric in the same units as the interface cost. ■ type-2—Metric is larger than any link-state path. This parameter applies in a connection profile only when OSPF is *not* active
ase-tag	Hexadecimal number that appears in management utilities and flags this route as external. It can also be used by border routers to filter this record. This parameter applies in a connection profile only when OSPF is *not* active.
ase7-adv	*Currently not used.*

Example of configuring a type 7 LSA in an NSSA

For background information about NSSAs, see "Hierarchical routing (areas)" on page 3-6. To configure the Stinger unit to route OSPF in an NSSA, *all* OSPF interfaces in the Stinger unit must specify the NSSA area-type.

To configure a type 7 LSA, you must specify a static route in an ip-route profile. Following are the related parameters (shown with sample settings):

```
[in IP-ROUTE/external]
name* = external
dest-address = 10.4.5.0/22
gateway-address = 10.4.5.7
metric = 0
cost = 1
```

```
preference = 100
third-party = no
ase-type = type-1
ase-tag = c0:00:00:00
private-route = yes
active-route = yes
ase7-adv = n/a
```

The following procedure configures the Stinger unit to route in an NSSA and import a type 7 LSA that specifies an external route across the WAN link:

1 **Assign an NSSA area type to each IP interface that is running OSPF. For example:**

```
admin> read ip-int {{ 1 12 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } read

admin> set ospf area-type = nssa

admin> write
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } written
```

2 **Configure the WAN link that represents an ASE route. For example:**

```
admin> read connection ase-like
CONNECTION/ase-link read

admin> set ip-options remote = 10.4.5.7/22

admin> set ip-options rip = routing-off

admin> set ip-options ospf active = yes

admin> write
CONNECTION/ase-link written
```

3 **Configure a static route to the remote site. For example:**

```
admin> new ip-route type7
IP-ROUTE/type7 read

admin> set dest = 10.4.5.0/22

admin> set gateway = 10.4.5.7

admin> write
IP-ROUTE/type7 written
```

Example of assigning a cost to a static route

The lower the cost assigned to a route, the more likely the router is to choose the route to forward traffic. Typically, you account for the bandwidth of a connection when assigning costs. For example, the cost for a single-channel connection is 24 times greater than for a T1 link.

The Stinger unit has a default cost of 1 for a connected route (Ethernet) and 10 for a WAN link. If two paths have the same destination, the Stinger unit uses the path with the lower cost.

Note Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.

In the following example, an administrator assigns a cost of 25 to a static route:

```
admin> new ip-route 56klink
IP-ROUTE/56klink read

admin> set dest = 10.1.2.0/24

admin> set gateway = 10.9.8.10

admin> set cost = 25

admin> write
IP-ROUTE/56klink written
```

Example of specifying a third-party route

OSPF can advertise routes to external destinations on behalf of another gateway (a third party). This function is commonly known as *advertising a forwarding address*. If third-party routing is disabled, the Stinger unit advertises itself as the forwarding address to an external destination. When third-party routing is enabled, the Stinger unit advertises the IP address of another gateway.

Depending on the topology of the network, other routers might be able to use this type of third-party LSA to route directly to the forwarding address without involving the advertising router, thus increasing the total network throughput. This feature can be used only if all OSPF routers know how to route to the forwarding address. For the route to be known, the forwarding address must be on a local network that has an OSPF router acting as the forwarding router. Or a designated router must send LSAs for that Ethernet network to any area that sees the static route's forwarding-address LSAs. Note that third-party routing cannot be used when ASE type 7 LSAs are advertised (as specified in RFC 1587).

In the following sample route, the Stinger unit advertises a third-party route (a forwarding address) for the destination 10.1.2.0. The forwarding address is 10.9.8.10.

```
admin> new ip-route fwd
IP-ROUTE/fwd read

admin> set dest = 10.1.2.0/24

admin> set gateway = 10.9.8.10

admin> set third-party = yes

admin> write
IP-ROUTE/fwd written
```

OSPF NBMA support

An OSPF nonbroadcast multiaccess (NBMA) network is any network that has multiple points of access (more than two routers) and does not support broadcast capability. Frame relay and X.25 are typically NBMA networks.

OSPF routers operate on an NBMA network much as they do on a broadcast network, by using the Hello protocol to form adjacencies and identify the designated router. However, because the routers cannot discover their neighboring routers dynamically by means of broadcasts, you must specify some additional parameters.

The Stinger unit forms adjacencies with other OSPF routers on an NBMA network. Adjacencies enable the unit to route OSPF over frame relay networks, and to

interoperate with the switches that do not support serial (point-to-point) connections over frame relay.

Note The non-multicast parameter in the ospf-options and ospf subprofiles for IP interfaces causes the translation of the multicast traffic to directed traffic. This parameter is typically used with a serial link, such as a point-to-point connection over frame relay, and is not intended for use with NBMA. Do not configure the non-multicast parameter for NBMA configurations.

Overview of OSPF NBMA settings

Following are the OSPF parameters (shown with default settings) related to NBMA. For information about the other OSPF parameters in the ospf and ospf-options subprofiles, see "Overview of LAN and WAN OSPF settings" on page 3-10.

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:ospf]
network-type = Point-to-Point
poll-interval = 0
```

```
[in CONNECTION/"":ip-options:ospf-options]
network-type = Point-to-Point
poll-interval = 0
```

```
[in OSPF-NBMA-NEIGHBOR/"" (new)]
name* = ""
host-name = ""
ip-address = 0.0.0.0
dr-capable = no
```

Parameter	Setting
network-type	Type of network to which the interface connects. For the parameter description, see page 3-13.
poll-interval	Interval, in seconds, at which to send Hello packets to a neighboring router that has become inactive. For the parameter description, see page 3-13.
name	Name of the ospf-nbma-neighbor profile. Specify a name of up to 24 characters.
host-name	Station name of a local connection profile that defines the connection to the neighboring router.
ip-address	IP address of the neighboring router.
dr-capable	Whether the neighboring router can be the designated router. Specify yes or no (the default).

Example of an OSPF NBMA configuration

In an NBMA network, a router that is eligible to become the designated router is configured with a list of all other OSPF routers connected to the network. At startup, these routers send Hello packets to each other to discover the designated router. The designated router then begins sending Hello packets to the entire list of routers on the network. When an NBMA interface becomes active on the Stinger unit, the unit sends Hello packets only to neighboring routers that are eligible to become the designated router, until it is notified about which router is the designated router.

Figure 3-8 shows an OSPF NBMA network using frame relay. For the purposes of this example, assume that the unit named FR-Router is eligible to become the designated router, and that the MAX-Router unit is not eligible.

Figure 3-8. OSPF NBMA network

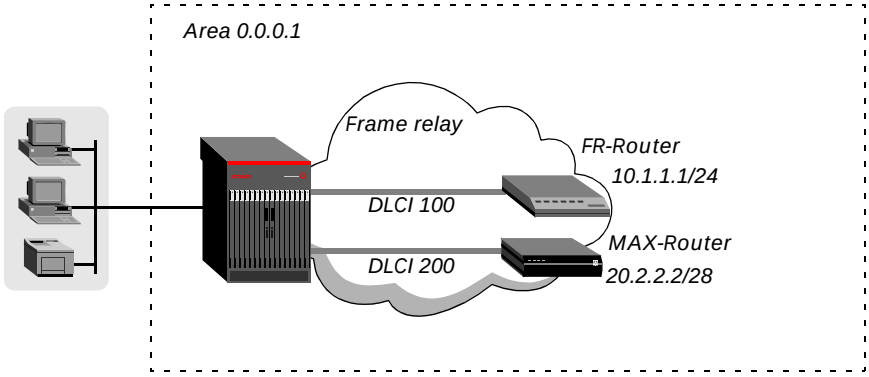

Example of configuring a designated router-capable neighboring router

The following set of commands defines a sample `frame-relay` profile for the interface to FR-Router in Figure 3-8:

```
admin> new frame-relay fr-dce
FRAME-RELAY/fr-dce read
admin> set active = yes
admin> set link-type = dce
admin> set nailed-up-group = 36
admin> set link-mgmt = ccitt
admin> write
FRAME-RELAY/fr-dce written
```

The next set of commands defines a `connection` profile for connection to FR-Router:

```
admin> new conn FR-Router
[in CONNECTION/FR-Router (new)]
admin> set active = yes
admin> set encapsulation-protocol = frame-relay
admin> set ip-options remote-address = 10.1.1.1/24
admin> set ip-options ospf active = yes
admin> set ip-options ospf area = 0.0.0.1
admin> set ip-options ospf network-type = nonbroadcast
admin> set ip-options ospf poll-interval = 60
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = fr-dce
admin> set fr-options dlci = 100
```

```
admin> write
CONNECTION/FR-Router written
```

The next set of commands enables the Stinger unit to form an adjacency with MAX-Router:

```
admin> new ospf-nbma-neighbor fr-router
[in OSPF-NBMA-NEIGHBOR/fr-router (new)]
admin> set host-name = FR-Router
admin> set ip-address = 10.1.1.1/24
admin> set dr-capable = yes
admin> write
OSPF-NBMA-NEIGHBOR/fr-router written
```

Example of configuring a non-DR-capable neighbor

The following set of commands defines a `frame-relay` profile for link operations on the interface to the unit named MAX-Router in Figure 3-8:

```
admin> new frame-relay fr-dte
FRAME-RELAY/fr-dte read
admin> set active = yes
admin> set link-type = dte
admin> set nailed-up-group = 11
admin> set link-mgmt = ccitt
admin> write
FRAME-RELAY/fr-dte written
```

The next set of commands defines a `connection` profile for connection from the Stinger unit to MAX-Router:

```
admin> new conn MAX-Router
[in CONNECTION/MAX-Router (new)]
admin> set active = yes
admin> set encapsulation-protocol = frame-relay
admin> set ip-options remote-address = 20.2.2.2/28
admin> set ip-options ospf active = yes
admin> set ip-options ospf area = 0.0.0.1
admin> set ip-options ospf network-type = nonbroadcast
admin> set ip-options ospf poll-interval = 60
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = fr-dte
admin> set fr-options dlci = 200
admin> write
CONNECTION/MAX-Router written
```

The next set of commands enables the Stinger unit to form an adjacency with FR-Router:

```
admin> new ospf-nbma-neighbor max-router
[in OSPF-NBMA-NEIGHBOR/max-router (new)]
admin> set host-name = MAX-Router
admin> set ip-address = 20.2.2.2/28
admin> write
OSPF-NBMA-NEIGHBOR/max-router written
```

Disabling OSPF

To globally disable the OSPF protocol, set the following parameter (shown with its default value):

```
[in IP-GLOBAL:ospf-global]
enable = yes
```

Parameter	Setting
enable	Enable/disable the OSPF protocol. A change to the setting takes effect immediately after you write the profile.

Although you can also deactivate OSPF manually on each OSPF interface, this parameter provides a system-wide mechanism for disabling the protocol. It can also be used to prevent OSPF from reinitializing several times if you are modifying many OSPF-related profiles. In that case, set the parameter to no, write the OSPF changes, and then set the parameter to yes again.

Virtual Routing

4

Virtual routing enables high-density circuit termination with secure logical partitioning and multiple route tables. Each virtual router (VRouter) is a grouping of interfaces in the Stinger unit. Each virtual router has its own associated routing table, ARP table, route cache, and address pools, and maintains its own routing and packet statistics.

If you do not configure any virtual routers, the IP router operates exactly as it is documented in Chapter 1, “IP Routing.” When one or more virtual routers are specified, the main router operates as the *global virtual router.* All interfaces that are not explicitly grouped with a defined virtual router are grouped with the global virtual router.

Introduction to virtual routing

Figure 4-1 shows a Stinger unit with one virtual router operating for Corporation A. Interfaces related to Corporation A are grouped and handled by one virtual router, creating a virtual private network (VPN) for Corporation A. Corporation A’s WAN interfaces can dial in to a local system, which can be on a public network, to reach Corporation A’s private LANs.

Figure 4-1. Virtual IP routing

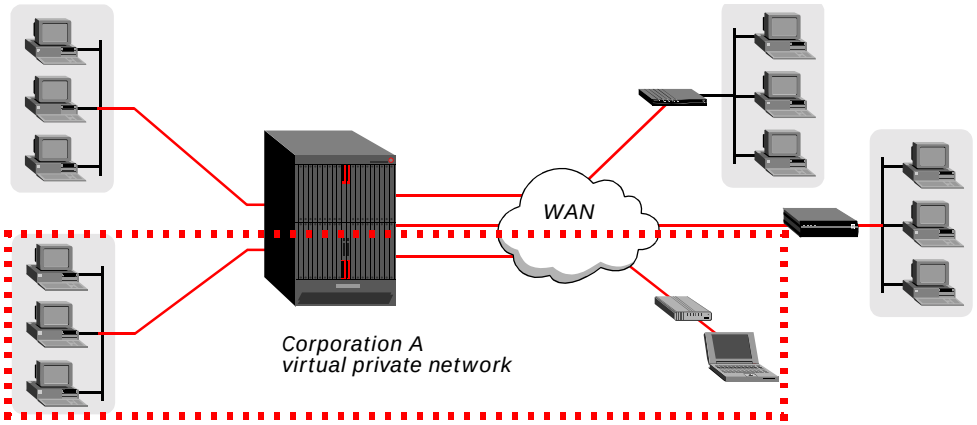

How virtual routers affect the routing table

When virtual routers are not defined, the main IP router maintains a single IP routing table that enables the router to reach any of its many interfaces. In that context, each interface known to the system requires a unique address.

With virtual routers, addresses must be unique within the virtual router's routing domain, but not necessarily within the Stinger unit. Because each virtual router maintains its own routing table, and because it knows about only those interfaces that explicitly specify the same virtual router, there is no requirement that the private networks maintain unique address spaces.

How virtual routers affect network commands

The commands partially described in Table 4-1 support virtual routing. If no virtual router name (*vroutername*) is specified on the command line, the system applies the command for the global virtual router (the main IP router). If a virtual router name is specified, the command performs its usual function but applies only to the specified virtual router:

Table 4-1. Network commands showing optional virtual router arguments

Command	Usage with optional arguments	Explanation of arguments
arptable	`arptable [`*`vroutername`*`] [[-a `*`hostname MAC_address`*`] | [-d `*`hostname`*`] | [-f]]`	`-a `*`hostname MAC_address`* Adds a hostname entry with a particular MAC address to the ARP table. `-d `*`hostname`* Deletes a hostname from the ARP table `-f` Clears an entire ARP cache For more information, see "Displaying the ARP cache" on page A-1 or the *Stinger Reference*.

Table 4-1. Network commands showing optional virtual router arguments (Continued)

Command	Usage with optional arguments	Explanation of arguments
ip-pools	ip-pools [*vroutername*]	Displays the status of all IP pools, or only those for a virtual router. For more information, see the *Stinger Reference*.
iproute	iproute add [-r *vroutername*] *destination_IPaddress/subnet_mask gateway_IPaddress* [*preference*] [metric] iproute delete [-r *vroutername*] *destination_IPaddress/subnet_mask* [gateway]	-r *vroutername* Adds or deletes a static IP route for a virtual router. For more information, see the *Stinger Reference*.
netstat	netstat [*vroutername*] [-i] [-r [*vroutername*]] [?] [-n] [-d] [-s *identifiers*] [-z]	-i Displays the interface table. -r *vroutername* Displays the IP routing table; entry for a virtual router. -? Displays a usage summary. -n Displays numeric addresses (the default). -d Displays symbolic names. -s *identifier* Displays protocol statistics. Identifiers are udp, tcp, icmp, ip, igmp, and mcast. -z Displays zombie RIP routes. For more information, see "Routes and interfaces" on page 1-1, "Displaying the IP interface table" on page A-2, "Displaying IP routes" on page A-4, or the *Stinger Reference*.
ping	ping [-q | -v] [-i *delay_sec* | -I *delay_msec*] [-s *packetsize*] [-r *vroutername*] [-x *source_IPaddress*] *hostname*	-q | -v Displays summary messages only, or almost all output. -i delay | I delay Delays the ping packet by a specified number of seconds or milliseconds. -s *packetsize* Sets the number of packet data bytes. -r *vroutername* Pings the virtual router. -x source_IPaddress For more information, see "Verifying point-to-point connectivity with ping" on page A-7 or the *Stinger Reference*.

Table 4-1. Network commands showing optional virtual router arguments (Continued)

Command	Usage with optional arguments	Explanation of arguments
telnet	telnet [-a | -b | -t] [-v *vroutername*] [-l[e] | -r[e]] *hostname* [*portnumber*]	-a | -b | -t Telnet mode: ASCII, binary, or transparent. -v vroutername Telnets to a virtual router. -l[e] | -r[e] Turns local echo on/off. For more information, see the *Stinger Reference*.
traceroute	traceroute [-n] [-v] [-m *max_ttl*] [-p *port*] [-q *nqueries*] [-w *waittime*] [-r *vroutername*] [-s src_IPaddr] *hostname* [*datasize*]	-n Displays hop addresses numerically only. -v Displays verbose output. -m *max_ttl* Sets the maximum number of hops used in trace probes. -p *port* Sets the base UDP port number in trace probes. -q *nqueries* Sets the maximum number of queries for each hop. -w *waittime* Sets the query response wait time. -r *vroutername* Traces the IP route for a virtual router. -s *src_IPaddr* Specifies a source host. For more information, see "Diagnosing problems with traceroute" on page A-8 or the *Stinger Reference*.

Current virtual router limitations

Currently, SNMP management does not display information about the Stinger unit on a per-virtual router basis. Errors and events are not logged on a per-virtual router basis. The syslog host defined in the system's log profile must be accessible to the main virtual router.

Currently, an ATMP Home Agent in router mode presents incoming packets only to the main virtual router. In addition, servers defined in the following profiles must be accessible to the main virtual router:

- debug
- trap
- external-auth
- ip-global (for SNTP and multicast)
- call-logging
- snmp

Creating a virtual router

When at least one vrouter profile is configured, the system-ip-address parameter and the global-vrouter parameter in the ip-global profile apply to the global virtual router. All interfaces that are not explicitly assigned to another virtual router are grouped with the global virtual router.

For each virtual router in the system, an instance of RIP is created to process routes. The new instance of RIP sends and receives update packets only on the interfaces associated with its particular virtual router and manipulates only that virtual router's routing table. A default instance of RIP is always created for the global virtual router.

When you create a virtual router, the new instance of RIP sends and receives packets only on the interfaces associated with that virtual router and manipulates only that virtual router's routing table. All RIP-related parameters in a vrouter profile use default settings that are recommended for most sites.

Overview of vrouter profile settings

A vrouter profile contains the following parameters (shown here with default values):

```
[in VROUTER/"" ]
name* = ""
active = yes
vrouter-ip-addr = 0.0.0.0
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 +
pool-name = [ "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "" "+
pool-summary = no
share-global-pool = yes
rip-policy = Poison-Rvrs
summarize-rip-routes = no
rip-trigger = yes
domain-name = ""
sec-domain-name = ""
dns-primary-server = 0.0.0.0
dns-secondary-server = 0.0.0.0
client-primary-dns-server = 0.0.0.0
client-secondary-dns-server = 0.0.0.0
allow-as-client-dns-info = True
```

Parameter	Setting
name	Unique name for the virtual router, up to 15 characters. All interfaces belonging to a virtual router specify the same virtual router name in the ip-interface or connection profile.
active	Activate the virtual router.
vrouter-ip-address	System IP address for the virtual router.
pool-base-address	Base address of a pool of contiguous addresses on a local network or subnet. The pool will be exclusively for use by the virtual router. For details about defining address pools, see Chapter 1, "IP Routing."
assign-count	Number of addresses in the pool. The pool will be exclusively for use by the virtual router. For details about defining address pools, see Chapter 1, "IP Routing."
pool-name	A pool name, required only when TACACS+ authentication is in use. The pool will be exclusively for use by the virtual router. For details about defining address pools, see Chapter 1, "IP Routing."
pool-summary	Set/clear the pool summary flag to specify that the address pools will be summarized. For details about defining address pools that can be summarized, see Chapter 1, "IP Routing."
share-global-pool	Enable/disable the virtual router to share the address pools defined in the ip-global profile.
rip-policy	Policy for the virtual router to use when sending update packets that include routes received on the same interface. For details about RIP policy, see Chapter 1, "IP Routing."
summarize-rip-routes	Whether the virtual router summarizes subnet information in RIP-v1 advertisements. For details about this feature, see Chapter 1, "IP Routing."
rip-trigger	Enable/disable RIP triggering for the virtual router. For details about RIP triggering, see Chapter 1, "IP Routing."

Note For details about domain-name and other DNS parameters, see "Configuring virtual router DNS servers" on page 4-14.

Example of defining a virtual router

The following commands create a virtual router for Corporation A:

```
admin> new vrouter corpa
VROUTER/corpa read

admin> set vrouter-ip-addr = 130.200.200.100
```

```
admin> write
VROUTER/corpa written
```

Displaying the virtual router netstat information

The new virtual router defined for Corporation A in "Example of defining a virtual router" maintains the following minimal routing and interface tables at this point:

```
admin> netstat corpa -rn
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
127.0.0.0/8	-	bh0_corpa	CP	0	0	0	8172
127.0.0.1/32	-	local	CP	0	0	0	8172
127.0.0.2/32	-	rj0_corpa	CP	0	0	0	8172

```
admin> netstat corpa -in
```

Name	MTU	Net/Dest	Address	Ipkts	Ierr	Opkts	Oerr
vr0_corpa	1500	130.2.2.2/32	130.2.2.2	0	0	0	0
lo0_corpa	1500	127.0.0.1/32	127.0.0.1	0	0	0	0
local	65535	127.0.0.1/32	127.0.0.1	0	0	0	0
rj0_corpa	1500	127.0.0.2/32	127.0.0.2	0	0	0	0
bh0_corpa	1500	127.0.0.3/32	127.0.0.3	0	0	0	0

The virtual router also maintains its own IP, TCP, UDP, and ICMP statistics. For example:

```
admin> netstat corpa -s
udp:
       1442 packets received
       0 packets received with no ports
       0 packets received with errors
       0 packets dropped
       32 packets transmitted
tcp:
       0 active opens
       1 passive opens
       0 connect attempts failed
       0 connections were reset
       1 connections currently established
       858 segments received
       0 segments received out of order
       548 segments transmitted
       0 segments retransmitted
       0 active closes
       0 passive closes
       0 disconnects while awaiting retransmission
icmp:
       31 packets received
```

```
        0 packets received with errors
        Input histogram:
                30 echo requests
                1 netmask requests

        31 packets transmitted
        0 packets not transmitted due to lack of resources
        Output histogram:
                30 echo replies
                1 netmask replies

ip:
        0 packets received
        0 packets received with header errors
        0 packets received with address errors
        0 packets received forwarded
        0 packets received with unknown protocols
        0 inbound packets discarded
        0 packets delivered to upper layers
        0 transmit requests
        0 discarded transmit packets
        0 outbound packets with no route
        0 reassemblies timeout
        0 reassemblies required
        0 reassemblies succeeded
        0 reassemblies failed
        0 fragmentation succeeded
        0 fragmentation failed
        0 fragmented packets created
        0 route discards due to lack of memory
        64 default ttl
igmp:
        0 packets received
        0 bad checksum packets received
        0 bad version packets received
        0 query packets received
        0 leave packets received
        0 packets transmitted
        0 query packets sent
        0 resonse packets sent
        0 leave packets sent
mcast:
        0 packets received
        0 packets forwarded
        0 packets in error
        0 packets dropped
        0 packets transmitted
```

Note There is no support for IP multicast on a per-virtual router basis, so the IGMP and multicast statistics relate only to the global virtual router.

Defining address pools for a virtual router

The following commands define an address pool for the Corporation A virtual router defined in “Example of defining a virtual router” on page 4-6:

```
admin> read vrouter corpa
VROUTER/corpa read

admin> set pool-base 1 = 130.100.100.128

admin> set assign-count 1 = 127

admin> write
VROUTER/corpa written
```

Following is a comparable RADIUS pool definition:

```
pools-taos01 Password = "ascend", Service-Type = Outbound-User
   Ascend-IP-Pool-Definition = "1 130.100.100.128 127 corpa"
```

The Corporation A virtual router is now maintaining the following pool of addresses:

```
admin> ip-pools corpa

Pool#                Base         Count      InUse
   1          130.100.100.128      127          0

 Number of remaining allocated addresses:    0
```

Note The Ascend-IP-Pool-Definition attribute supports a virtual router name as the last syntax element in a pool definition. The value of Ascend-IP-Pool-Definition uses the following syntax:

"*pool-num base-addr assign-count* [*vrouter-name*]"

For background information about address pools, see “Configuring and using address pools” on page 1-60. The process of defining address pools for a virtual router is the same as described in that section.

Assigning interfaces to a virtual router

To assign virtual router membership to an interface, you specify a virtual router name in the interface profile. For a virtual router to be active, at least one IP interface (LAN or WAN) must specify its name.

Overview of settings

To assign virtual router membership to an interface in local profiles, set the `vrouter` parameter. For example:

```
[in IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 } ]
vrouter = corpa

[in CONNECTION/corpa-client]
vrouter = corpa
```

Parameter	Setting
vrouter	Name of a defined virtual router. Specifying the virtual router name groups the interface with the virtual router. The default null value specifies the global virtual router.

RADIUS uses the following attribute-value pair to support the use of a virtual router:

RADIUS attribute	Value
Ascend-VRouter-Name (102)	Name of a defined virtual router. Specifying the virtual router name groups the interface with the virtual router. The default null value specifies the global virtual router.

Examples of assigning virtual router membership to interfaces

The following commands group three WAN interfaces with the corpa virtual router:

```
admin> new connection dialin-1
CONNECTION/dialin-1 read
admin> set active = yes
admin> set vrouter = corpa
admin> set ip-options remote-address = 10.1.1.1/24
admin> write
CONNECTION/dialin-1 written
admin> new connection dialin-2
CONNECTION/dialin-2 read
admin> set active = yes
admin> set vrouter = corpa
admin> set ip-options remote-address = 11.1.1.1/24
admin> write
CONNECTION/dialin-2 written
admin> new connection dialin-3
CONNECTION/dialin-3 read
admin> set active = yes
admin> set vrouter = corpa
admin> set ip-options remote-address = 12.1.1.1/24
admin> write
CONNECTION/dialin-3 written
```

Following are comparable settings in RADIUS profiles:

```
dialin-1 Password = "pwd3"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
```

```
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.255.255.0,
   Ascend-Vrouter-Name = "corpa"

dialin-2 Password = "pwd2"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
   Framed-IP-Address = 11.1.1.1,
   Framed-IP-Netmask = 255.255.255.0,
   Ascend-Vrouter-Name = "corpa"

dialin-3 Password = "pwd1"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
   Framed-IP-Address = 12.1.1.1,
   Framed-IP-Netmask = 255.255.255.0,
   Ascend-Vrouter-Name = "corpa"
```

Displaying assigned virtual router interfaces

After interfaces have been assigned, as described in "Examples of assigning virtual router membership to interfaces" on page 4-10, new interfaces appear in the virtual router's routing and interface tables when the interfaces become active. For example:

```
admin> netstat corpa -rn
Destination    Gateway       IF          Flg   Pref Met    Use        Age
10.0.0.0/24    10.1.1.1      wan30       SG     120   7      0        215
10.1.1.1/32    10.1.1.1      wan30       S      120   7      1        215
11.0.0.0/24    11.1.1.1      wan31       SG     120   7      0        215
11.1.1.1/32    11.1.1.1      wan31       S      120   7      1        215
12.0.0.0/24    12.1.1.1      wan32       SG     120   7      0        215
12.1.1.1/32    12.1.1.1      wan32       S      120   7      1        215
127.0.0.0/8    -             bh0_corpa   CP       0   0      0       1193
127.0.0.1/32   -             local       CP       0   0      0       1193
127.0.0.2/32   -             rj0_corpa   CP       0   0      0       1193

admin> netstat corpa -in
Name       MTU   Net/Dest        Address       Ipkts Ierr Opkts  Oerr
vr0_corpa  1500 130.2.2.2/32    130.2.2.2         0    0      0    0
lo0_corpa  1500 127.0.0.1/32    127.0.0.1         0    0      0    0
local     65535 127.0.0.1/32    127.0.0.1         0    0      0    0
rj0_corpa  1500 127.0.0.2/32    127.0.0.2         0    0      0    0
bh0_corpa  1500 127.0.0.3/32    127.0.0.3         0    0      0    0
wan30      1500 10.1.1.1        130.2.2.2         0    0      0    0
wan31      1500 11.1.1.1        130.2.2.2         0    0      0    0
wan32      1500 12.1.1.1        130.2.2.2         0    0      0    0
```

Defining virtual router static routes

You specify a static route associated with a virtual router for one of the following reasons:

- To define a route on a per-virtual router basis
- To specify an inter-virtual router route

Overview of static route settings

Following are the virtual router-related parameters (shown here with default values) in ip-route profiles:

```
[in IP-ROUTE/""]
vrouter = ""
inter-vrouter = ""
```

Parameter	Setting
vrouter	Name of the virtual router that will own this route. The route will be part of that virtual router's routing table. If no name is specified (the default), the global virtual router is assumed.
inter-vrouter	Name of a virtual router to use as the route's next hop. All packets to the static route's destination network are sent to the specified virtual router for a routing decision. The gateway-address parameter must be set to the zero address for this parameter to apply.

In a RADIUS profile, the value of the Framed-Route (22) attribute can specify a virtual router name in the following syntax:

"*dest-addr* [*/prefix*] *gateway-addr metric* [private] [*profile*] [*preference*] [*vrouter-name*]"

Note The fields within the value of the Framed-Route attribute are positional. With the exception of the optional prefix-length specification, if any of the optional fields are specified, the optional fields to the left of that setting must also be specified.

Examples of defining a route on a per-virtual router basis

Following is an example of defining a static route to Corporation B. This route is within the Corporation A virtual router domain (the virtual router named corpa will own this route).

```
admin> new ip-route corpa-east
IP-ROUTE/corpa-east read

admin> set dest = 10.5.6.7/28

admin> set gateway = 10.1.1.1

admin> set vrouter = corpa

admin> write
IP-ROUTE/corpa-east written
```

Following is a comparable RADIUS profile:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
   Framed-Route = "10.5.6.7/28 10.1.1.1 7 n corpa-east 60 corpa"
```

Displaying virtual router static routes

The following sample output shows the new static route that was added to the Corporation A virtual router's routing table in "Examples of defining a route on a per-virtual router basis" :

```
admin> netstat corpa -rn
Destination     Gateway      IF          Flg   Pref Met    Use        Age
10.1.1.0/24     10.1.1.1     wan30       SG     120   7      0          9
10.1.1.1/32     10.1.1.1     wan30       S      120   7      2          9
10.5.6.0/28     10.1.1.1     wan30       SG      60   8      0          9
11.1.1.0/24     11.1.1.1     wan31       SG     120   7      0          9
11.1.1.1/32     11.1.1.1     wan31       S      120   7      1          9
12.1.1.0/24     12.1.1.1     wan32       SG     120   7      0          9
12.1.1.1/32     12.1.1.1     wan32       S      120   7      1          9
127.0.0.0/8     -            bh0_corpa   CP       0   0      0       2274
127.0.0.1/32    -            local       CP       0   0      0       2274
127.0.0.2/32    -            rj0_corpa   CP       0   0      0       2274
```

Specifying an inter-virtual router route

In the following example, the static route specifies the Corporation A virtual router as the route's next hop. All packets to the specified destination network are sent to the specified virtual router for a routing decision.

```
admin> new ip-route corpb
IP-ROUTE/corpb read
admin> set dest-address = 11.0.0.0/24
admin> set inter-vrouter = corpa
admin> write
IP-ROUTE/corpb written
```

Following is a comparable RADIUS route profile:

```
route-taos-1 Password = "ascend", Service-Type = Outbound-User
   Framed-Route = "11.0.0.0/28 0.0.0.0 corpa"
```

Displaying the inter-virtual router route in the global table

In the following example, the route has been added to the global virtual router's routing table, not to that of the Corporation A virtual router:

```
admin> netstat -rn
Destination      Gateway     IF         Flg  Pref Met    Use     Age
0.0.0.0/0        10.1.6.1    ie0        SGP   60   1      59       4
11.0.0.0/24       -          vr0_corpa   S    60   8       0       4
20.0.0.0/8        -          ie1-12-1    C     0   0      12     234
20.1.1.2/32       -          local      CP     0   0       0    2347
127.0.0.0/8       -          bh0        CP     0   0       0    2378
127.0.0.1/32      -          local      CP     0   0       0    2378
127.0.0.2/32      -          rj0        CP     0   0       0    2378
130.1.1.1/32      -          sip0        C     0   0       0    2378
130.1.1.252/30    -          rj0         C     0   0       0    2378
100.1.6.0/24     100.1.6.221 wanabe     SG    60   1       0       4
101.1.6.0/24      -          ie0         C     0   0    2531    2378
```

```
101.1.6.234/32    -          local      CP     0   0    4152     2378
224.0.0.0/4       -          mcast      CP     0   0       0     2378
224.0.0.1/32      -          local      CP     0   0       0     2378
224.0.0.2/32      -          local      CP     0   0       0     2378
224.0.0.5/32      -          local      CP     0   0     732     2378
224.0.0.6/32      -          local      CP     0   0       0     2378
255.255.255.255/32 -          ie0         P     0   0     422     2378
```

Configuring virtual router DNS servers

Virtual router DNS configuration includes settings for primary and secondary DNS servers, domain names, and client DNS servers. The settings direct connections that belong to the virtual router to a particular DNS service. To completely segment the virtual router's DNS information from any other hosts, you can configure and manage DNS information separately for each virtual router. The addresses configured for client DNS servers are presented to dial-in users during IP Control Protocol (IPCP) negotiation.

If DNS information is not found in the `vrouter` profile, the system uses the DNS information in the `ip-global` profile. The DNS list and the local DNS table maintained in RAM are systemwide DNS configurations that are not supported separately for each virtual router.

Overview of virtual router DNS settings

Following are the virtual router-specific DNS parameters (shown with their default settings):

```
[in VROUTER/""]
domain-name = ""
sec-domain-name = ""
dns-primary-server = 0.0.0.0
dns-secondary-server = 0.0.0.0
client-primary-dns-server = 0.0.0.0
client-secondary-dns-server = 0.0.0.0
allow-as-client-dns-info = True
```

Parameter	Setting
`domain-name`	Primary domain name (up to 63 characters) to use for DNS lookups for this virtual router. The system appends this domain name to hostnames when performing lookups.
`sec-domain-name`	Secondary domain name to use for DNS lookups for this virtual router if the hostname is not found in the primary domain.
`dns-primary-server`	Address of the primary local DNS server to use for lookups for this virtual router.
`dns-secondary-server`	Address of the secondary local DNS server to use for lookups for this virtual router. Used only if the primary server is not found.

Parameter	Setting
client-dns-primary-server	Address of a client DNS server for dial-in clients of this virtual router.
client-dns-secondary-server	Address of a secondary DNS server for dial-in clients of this virtual router.
allow-as-client-dns-info	Enable/disable use of main (local) DNS information if the client DNS servers are not found. To isolate local network information for this virtual router, set to false.

Example of a typical virtual router DNS configuration

The following commands specify a primary and secondary domain name for DNS lookups for a virtual router named xyz:

```
admin> read vrouter xyz
VROUTER/xyz read

admin> set domain-name = xyz.com

admin> set sec-domain-name = eng.xyz.com

admin> write
VROUTER/xyz written
```

If a lookup fails in the first domain, the router tries again with the secondary domain name. To enable the system to use DNS to look up addresses, specify DNS server addresses, as shown in the following example:

```
admin> read vrouter xyz
VROUTER/xyz read

admin> set dns-primary-server = 1.2.2.2

admin> set dns-secondary-server = 1.3.3.3

admin> write
VROUTER/xyz written
```

If the primary server is unavailable, the system attempts a lookup on the secondary server. The following commands configure a client DNS server for this virtual router:

```
admin> read vrouter xyz
VROUTER/xyz read

admin> set client-dns-primary-server = 1.2.2.2

admin> set client-dns-secondary-server = 1.2.2.96

admin> set allow-as-client-dns-info = false

admin> write
VROUTER/xyz written
```

The secondary server is accessed only if the primary one is inaccessible. If both of these client DNS servers are not accessible, the system does not allow the client to access local DNS servers.

Configuring virtual routers for tunneled connections

Virtual routers can be used to tunneled connections that use Layer 2 Tunnel Protocol (L2TP) or Ascend Tunnel Management Protocol (ATMP). When you specify a virtual router for a tunnel, L2TP or ATMP packets (control channel and encapsulated data) for that tunnel are sent only via that virtual router. For information about L2TP tunnels, see "L2TP Tunneling" on page 5-1. For information about ATMP, see "ATMP Tunneling" on page 6-1.

Because each virtual router maintains its own routing table and knows about only those interfaces that explicitly specify the same virtual router, you can use virtual routers to direct tunneled traffic to a specific L2TP network server (LNS) or ATMP Home Agent, even if multiple server end points exist with duplicate IP addresses.

For example, Figure 4-2 shows two dial-in clients, MC-1 and MC-2. Each client tunnels to a different LNS, but both LNS systems have the IP address 1.1.1.1. Because the tunnels are built on separate virtual routers, the traffic is kept separate and directed to the appropriate server end point.

Figure 4-2. L2TP tunnels built on separate virtual routers

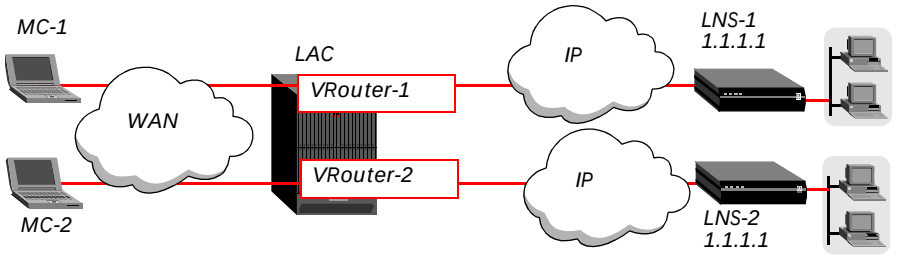

Connection profile setting for tunneling

Following is the parameter (shown with its default value) for specifying a virtual router name for a tunneled connection:

```
[in CONNECTION/MC-1:tunnel-options]
vrouter = ""
```

Parameter	Setting
vrouter	Name of a virtual router to use for establishing the tunnel. The specified virtual router must exist in the Stinger unit operating as LAC or Foreign Agent. With the default null value, the global virtual router is used.

For example, the following commands configure a mobile-client profile for an L2TP session that belongs to a virtual router named VRouter-1:

```
admin> new connection MC-1
CONNECTION/MC-1 read
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = localpw
```

```
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options tunneling-protocol = l2tp
admin> set tunnel-options vrouter = VRouter-1
admin> write
CONNECTION/MC-1 written
```

With this sample profile, the system authenticates the caller before building a tunnel to the LNS at 1.1.1.1 on the specified virtual router.

RADIUS profile setting for tunneling

RADIUS uses the following attribute-value pair to specify a virtual router name for a tunneled connection:

RADIUS attribute	Value
Ascend-Tunnel-VRouter-Name (31)	Name of a virtual router to use for establishing the tunnel. The specified virtual router must exist in the Stinger unit operating as LAC or Foreign Agent. With the default null value, the global virtual router is used. This attribute supports tagging.

For example, the following mobile-client profile specifies an L2TP session that belongs to a virtual router named VRouter-2:

```
MC-2 Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Tunnel-Server-Endpoint = "1.1.1.1",
    Tunnel-Type = L2TP,
    Ascend-Tunnel-VRouter-Name = "VRouter-2"
```

Following is a sample RADIUS profile for an ATMP connection:

```
commgroup Password = "123"
    User-Service = Framed-User,
    Framed-IP-Address = 199.199.199.200,
    Framed-IP-Netmask = 255.255.255.0,
    Framed-Protocol = PPP,
    Ascend-Route-IP = 1,
    Tunnel-Type = ATMP
    Tunnel-Server-Endpoint = 10.5.7.2,
    Tunnel-Password = "atmp",
    Ascend-Home-Agent-UDP-Port = 5150,
    Ascend-Home-Network-Name = "HOMNET",
    Ascend-Tunnel-VRouter-Name = "companyvr"
```

Deleting a virtual router

Deleting a `vrouter` profile deletes the virtual router. For example:

```
admin> delete vrouter corpa
```

Lucent Technologies recommends that you reset the system after deleting a virtual router with active connections. If a system reset is not possible, the recommended course of action before deleting the virtual router is to manually tear down its active connections, and then modify the local `connection`, `ip-interface`, and `ip-route` profiles that point to the virtual router to point instead to the global virtual router or another existing virtual router.

L2TP Tunneling

5

The T1000 module supports Layer 2 Tunneling Protocol (L2TP) for virtual private network (VPN) connectivity. For information about using Ascend Tunnel Management Protocol (ATMP) as an alternate method for VPN connectivity, see Chapter 6, “ATMP Tunneling.”

Note RADIUS-authenticated PPP sessions can make use of some L2TP tunnel features, such as tunnel tags, that are not supported in the local command-line interface. For details about those features, see the *TAOS RADIUS Guide and Reference*.

Introduction to L2TP tunneling

L2TP provides tunneling at OSI Layer 2 (at the HDLC layer of a PPP connection). A T1000 module operates as an L2TP access concentrator (LAC), which is the tunnel initiator. The T1000 module initiates a tunnel when it receives PPP session requests and begins the process of establishing a connection to an L2TP network server (LNS).

Figure 5-1 shows the elements of an L2TP tunnel. A PPP client (referred to as the *mobile client*) initiates a session request to the T1000 module. The module establishes the session and passes the data stream to the LNS across an IP network.

Figure 5-1. L2TP tunneling

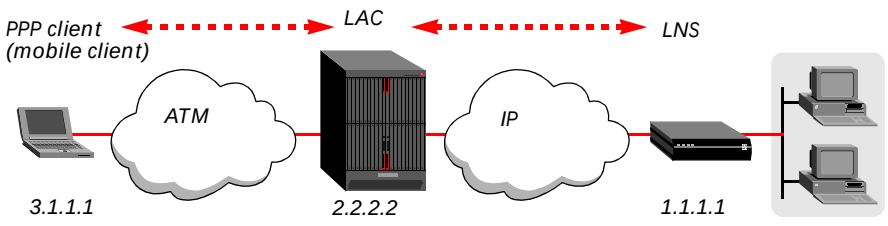

The mobile client can be any PPP client. For example, it could be a bridging modem supporting PPP over Ethernet (PPPoE) clients on a remote LAN, or an IP-routed PPP over ATM client.

The link between the LAC and the LNS can traverse a WAN interface, or it can be an Ethernet link. The connection to the LNS is an IP link, which consists of a control link and zero or more data links. Both the control and data links are encapsulated in UDP.

The control link carries information that is used both to query whether the LNS can accept the current call and to establish a tunnel. L2TP implements a Hello mechanism by which the LAC and LNS verify that both are operational by sending each other a control message every minute or so. If the Hello message does not arrive for several minutes, the tunnel and all the tunneled connections are torn down.

Data links carry the client data, which consists of PPP frames. Each tunneled client connection supports one data link.

Network settings for L2TP

Network settings for L2TP include settings for the IP connection between the LAC and LNS, and settings for the UDP communication required to establish tunnels.

System reset requirement

When you change the setting of the `base-udp-port` parameter in the `l2-tunnel-global` profile, a system reset is required for the L2TP subsystem to recognize the new UDP port number. All other parameter settings in the `l2-tunnel-global` profile take effect as soon as possible after the profile is written.

System IP address recommendation

Lucent Technologies recommends that you set the `system-ip-addr` parameter in the `ip-global` profile of a Stinger unit that is operating as a LAC, particularly if the unit has multiple interfaces into the IP cloud. This setting is not required if RIP is enabled on the interfaces between the LAC and LNS, but it is recommended, because it helps to simplify configurations.

Specifying a system name

Although L2TP provides several ways for the LAC to send a name to the tunnel server for authenticating a tunnel, the LAC sends the system name if none of the L2TP-specific names are specified. Keep this consideration in mind when configuring tunnel authentication and if necessary, specify a value for the `name` parameter in the `system` profile. For details about L2TP-specific names used in tunnel authentication, see “L2TP-specific IDs for tunnel authentication” on page 5-13.

Configuring LAC settings for all L2TP tunnels

To enable the T1000 module to operate as a LAC, you must first enable L2TP and configure some global tunneling parameters. The parameters in the `l2-tunnel-global` profile are used globally for all LAC operations. They are not specific to one LNS.

Top-level LAC operations

Following are the parameters, shown with default values, that enable the T1000 module to operate as LAC:

```
[in L2-TUNNEL-GLOBAL]
l2tp-mode = disabled
l2tp-auth-enabled = no
l2tp-rx-window = 0
l2tp-system-name = ""
```

Parameter	Setting
`l2tp-mode`	Enable/disable L2TP operations. With the default `disabled` value, L2TP is disabled. Set the parameter to `lac` to enable L2TP LAC operations.
`l2tp-auth-enabled`	Enable/disable L2TP tunnel authentication. With the `yes` setting, you can configure a variety of authentication methods as described in "L2TP-specific IDs for tunnel authentication" on page 5-13.
`l2tp-rx-window`	Advertised L2TP receive window size for data channels. The default, 0 (zero), specifies that the LAC asks for no flow control for inbound L2TP payloads. The valid range is from 0 to 63. **Note** A nonzero value enables behavior that predates RFC 2661. Not all L2TP implementations support a nonzero value.
`l2tp-system-name`	Name (up to 31 characters) that can be sent to the LNS during tunnel authentication. For details about L2TP-specific names used in tunnel authentication, see "L2TP-specific IDs for tunnel authentication" on page 5-13.

Enabling L2TP authentication

By default, L2TP authentication is disabled. With that setting, the LAC does not send a shared-secret value to the LNS, nor does it employ any of the authentication methods described in "L2TP-specific IDs for tunnel authentication" on page 5-13. When you enable L2TP authentication, you can choose the level of authentication you need to verify each LNS before bringing up a control channel.

L2TP timers and other variables

Following are the parameters, shown with default values, for configuring timers and other variables that affect tunnel retry attempts and other control message variables:

```
[in L2-TUNNEL-GLOBAL:l2tp-config]
first-retry-timer = 1000
retry-count = 10
hello-timer = 60
control-connect-establish-timer = 60
lac-incoming-call-timer = 60
```

```
base-udp-port = 0
dialout-auth-lns = no
dialout-send-profile-name = no
verify-remote-host-name = no
acct-tunnel-connection-encoding = normal
tunnel-server-pre-sccrq-lookup = n
```

Parameter	Setting
`first-retry-timer`	Initial interval, in milliseconds, that the system waits before retransmitting control packets in the attempt to establish an L2TP tunnel with an LNS system. The valid range is from 100 to 5000. The default is 1000. For details about how this timer works with the `retry-count` parameter in establishing and maintaining tunnel sessions, see "Specifying tunnel retry limits" on page 5-5.
`retry-count`	Number of times TAOS retransmits control packets in the attempt to establish or reestablish a tunnel. The valid range is from 1 to 10. The default value is 10.
`hello-timer`	Interval, in seconds, between Hello messages sent to the LNS. Specify a number from 0 to 600. The default is 60. The 0 setting specifies that no Hello messages are sent.
`control-connect-establish-timer`	Number of seconds during which the T1000 module can establish an L2TP tunnel with another host. Enter an integer from 0 to 600. The default is 60.
`lac-incoming-call-timer`	Number of seconds the T1000 module waits for an incoming session request to complete. Specify an integer from 1 to 600. The default is 60.
`base-udp-port`	UDP port number. The default zero value causes the system to dynamically assign a nonprivate port for exchanging control information while establishing a tunnel. This avoids the possibility of requesting a UDP port that is already in use. Nonzero values from `10,000` to `60,000` can be configured to use a known port, which is sometimes a firewall requirement. The configured value is used to assign a port number by using the following formula: *base-udp-port* + (*shelf-number* × `100`) + *slot-number*
`dialout-auth-lns`	If set to `yes`, restricts the LAC to accept dialout requests only from the LNS which has authenticated during the tunnel setup.
`dialout-send-profile-name`	Enable/disable the LNS to send the connection profile name vendor-specific attribute along with the dialout request.

Parameter	Setting
verify-remote-host-name	Enable/disable verification of the host name returned by the LNS. When enabled, the LAC compares the hostname returned by the LNS in the Start-Control-Connection-Reply (SCCRP) packet to the server-auth-id value configured in the local tunnel-server profile or the Tunnel-Server-Auth-ID value in a RADIUS profile. If the values do not match, the LAC terminates the tunnel request.
acct-tunnel-connection-encoding	Encoding method for the value of the RADIUS Acct-Tunnel-Connection attribute. The default normal setting specifies that the value consists of the source and destination IP addresses, tunnel ID, and connection ID. This value is expected by NavisRadius. The decimal-call-serial-number setting specifies that the attribute value consists of the 32-bit L2TP call serial number (CSN) encoded as a decimal string. The hexadecimal-call-serial-number setting specifies that the attribute value consists of the L2TP CSN encoded as a hexadecimal string. See "Setting appropriate RADIUS tunnel accounting values" on page 5-6 for related information.
tunnel-server-pre-sccrq-lookup	Enables/disables a lookup for a tunnel-server profile when a password is not available for a tunnel request. With the default no value, a lookup occurs for a matching tunnel-server profile occurs after the system receives an L2TP Start-Control-Connection-Reply (SCCRP) packet from the LNS. If it is set to yes and the tunnel password is not available, the system looks for a matching tunnel-server profile before the LAC sends a Start-Control-Connection-Request (SCCRQ) packet.

Specifying tunnel retry limits

RFC 2661, *Layer Two Tunneling Protocol "L2TP"* specifies that each retransmission of an L2TP control message must employ an exponential backoff interval—if the first retransmission occurs after 1 second, the next retransmission should occur after 2 seconds, the next after 4 seconds, and so forth. The intention of this requirement is to restore normal traffic quickly after a control packet is lost. However, the exponential backoff algorithm can take many seconds to restore traffic if several control packets are dropped.

TAOS uses a more aggressive retransmission strategy than required by RFC 2661. In TAOS, the retransmission delay is calculated on the basis of the values specified in the l2-tunnel-global profile of the first-retry-timer and retry-count parameters to provide a more linear backoff interval. This algorithm allows for a quicker recovery in the event of lost control packets. For example, with the following default values, retransmission occurs for up to 27 seconds:

```
[in L2-TUNNEL-GLOBAL:l2tp-config]
first-retry-timer = 1000
retry-count = 10
```

The initial transmission and 10 retransmissions (as specified by the retry count) occur at the following intervals:

1 1 second
2 1 second
3 1 second
4 1 second
5 2 seconds
6 2 seconds
7 3 seconds
8 3 seconds
9 4 seconds
10 4 seconds
11 5 seconds

With the following configured values, retransmission occurs for up to 11 seconds:

```
[in L2-TUNNEL-GLOBAL:l2tp-config]
first-retry-timer = 1000
retry-count = 6
```

The initial transmission and six retransmissions occur at the following intervals:

1 1 second
2 1 second
3 1 second
4 1 second
5 2 seconds
6 2 seconds
7 3 seconds

Setting appropriate RADIUS tunnel accounting values

For details about setting up the Stinger unit to support RADIUS accounting, see the *TAOS RADIUS Guide and Reference*. TAOS supports L2TP tunnel accounting in accordance with RFC 2867, *RADIUS Accounting Modifications for Tunnel Protocol Support.*

Note There is no guarantee that the L2TP call serial number (CSN) is unique at all times. Even in a single tunnel, identical CSN values can occur if tunnel links are initiated from both the LAC and LNS sides. Therefore, use caution when specifying CSN encoding in the `acct-tunnel-connection-encoding` parameter.

For an L2TP tunnel itself, no CSN exists. With CSN encoding, the 32-bit value that is encoded for the tunnel itself represents the initiator tunnel ID in the low 16 bits and the server tunnel ID in the high 16 bits.

Sample global L2TP tunneling configuration

The following commands enable the system to operate as LAC for L2TP tunnels. Tunnel authentication is enabled, and the system name `stinger-lac` is used for authentication purposes if the `tunnel-server` profile do not specify a name.

```
admin> read l2-tunnel
L2-TUNNEL-GLOBAL read
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = yes
admin> set l2tp-system-name = stinger-lac
admin> set l2tp-config retry-count = 6
admin> write
L2-TUNNEL-GLOBAL written
```

This configuration limits the retransmission period for failed (or terminated) tunnel attempts to 11 seconds by reducing the retry count to 6. For details, see “Specifying tunnel retry limits” on page 5-5.

Configuring LNS end points

When L2TP LAC operations have been enabled globally (as described in the preceding sections), the Stinger unit can initiate tunnel requests to LNS systems for which it has a valid tunnel-server configuration.

Overview of tunnel server settings

Following are the parameters, shown with default values, that enable the T1000 module to interoperate with the specified LNS:

```
[in TUNNEL-SERVER/""]
server-endpoint* = ""
enabled = yes
shared-secret = ""
client-auth-id = ""
server-auth-id = ""
```

Parameter	Setting
`server-endpoint`	Hostname or IP address that identifies the LNS. Usually, this is the same value as the Tunnel-Server-Endpoint RADIUS attribute, but it can differ. If a hostname is specified, the Stinger unit performs a DNS lookup for the host’s address.
`enabled`	Enable/disable tunnels to the specified LNS.
`shared-secret`	Value (up to 21 characters) used by both the LAC and the LNS ends of the tunnel to authenticate tunnel requests initiated by local `connection` profiles. For related RADIUS information, see “Shared secrets and secure exchanges” on page A-5.

Parameter	Setting
client-auth-id	LAC system name used for tunnel authentication. This name is sent in the Start-Control-Connection-Request (SCCRQ) message. For details about L2TP-specific names used in tunnel authentication, see "L2TP-specific IDs for tunnel authentication" on page 5-13.
server-auth-id	LNS system name used for tunnel authentication. This name is sent in the SCCRP. For details about L2TP-specific names used in tunnel authentication, see "L2TP-specific IDs for tunnel authentication" on page 5-13.

Shared secret tunnel authentication

The following RADIUS attribute-value pair provides the same capability as the shared-secret parameter in a tunnel-server profile:

RADIUS attribute	Value
Tunnel-Password (69)	Password sent to the LNS for authenticating a tunnel request initiated from a RADIUS profile. The Tunnel-Password value must be encrypted by the RADIUS server. Otherwise, tunnel authentication fails. You can specify the attribute in individual user profiles or in a pseudo-user profile used for authentication only.

For example, the following commands specify a shared secret to be used for authenticating tunnels initiated from local connection profiles to an LNS at lns-sys.domain.org:

```
admin> read tunnel-server lns-sys.domain.org
TUNNEL-SERVER/lns-sys.domain.org read

admin> set enabled = yes

admin> set shared-secret = tunnel-secret

admin> write
TUNNEL-SERVER/lns-sys.domain.org read
```

Following is a RADIUS user profile that uses the Tunnel-Password attribute for authentication with the same LNS:

```
l2tp-client Password = "my-password"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Framed-IP-Address = 10.50.1.1,
   Framed-IP-Netmask = 255.255.0.0,
   Tunnel-Type = L2TP,
   Tunnel-Medium-Type = IP,
   Tunnel-Server-Endpoint = "lns-sys.domain.org",
   Tunnel-Password = "tunnel-secret"
```

If you prefer, you can remove the Tunnel-Password attribute from calling clients' profiles and create a profile whose sole purpose is to authenticate L2TP tunnels. This

approach causes an extra RADIUS lookup the first time the tunnel is created, but it simplifies administration when shared secrets change. The RADIUS profile for tunnel authentication must specify the L2TP peer's name, a null password (""), and the Outbound-User setting for Service-Type. For example:

```
lns-sys.domain.org Password = "", Service-Type = Outbound-User
   Tunnel-Password = "tunnel-secret"
```

When an L2TP tunnel is initially established, both the LNS and the LAC issue a RADIUS lookup based on the peer's name. If the system finds a profile such as the one shown in the preceding example, it uses the Tunnel-Password value to authenticate the tunnel.

Note The password in the pseudo-user profile must be null (""). Because the null password represents a security risk, the profile must specify the Outbound-User setting for Service-Type.

Typical primary and secondary tunnel server configuration

Many sites configure two LNS end points to enable the system to automatically initiate a session with the second server if a tunnel request to the initial server fails.

```
admin> read tunnel-server l2tp-primary
TUNNEL-SERVER/l2tp-primary read
admin> set enabled = yes
admin> set shared-secret = secret1
admin> set client-auth-id = stinger-lac
admin> write
TUNNEL-SERVER/l2tp-primary read
admin> read tunnel-server l2tp-secondary
TUNNEL-SERVER/l2tp-secondary read
admin> set enabled = yes
admin> set shared-secret = secret2
admin> set client-auth-id = example.abc.com
admin> write
TUNNEL-SERVER/l2tp-secondary read
```

For details about configuring client profiles to attempt a tunnel with the secondary server if the primary server is unavailable, see "Typical connection to two tunnel end points" on page 5-12.

Configuring client connections

If a PPP client's profile is configured to initiate an L2TP tunnel, the LAC attempts to open a tunnel after it authenticates the client's session request by means of a name and password. After the initial authentication, the LAC negotiates Link Control Protocol (LCP) with the client and then opens the PPP Auth state. The LAC then forwards relevant LCP information (*proxy LCP*) as well as the caller's name and password (*proxy authentication*) to the LNS, so the LNS does not need to restart negotiation.

With proxy LCP, the LAC sends the following information to the LNS:

- The first LCP Config Request packet received from the client
- The last LCP Config Request packet received from the client
- The last LCP Config Request packet the LAC sent to the client

With proxy authentication, the LAC initiates PPP authentication of the session request and then sends the caller's name and password to the LNS in the appropriate L2TP attribute-value pairs. The LNS can then complete PPP authentication.

Overview of client profile settings

Following are the parameters, shown with default values, for configuring L2TP tunneling connections in local profiles:

```
[in CONNECTION/"":tunnel-options]
profile-type = disabled
tunneling-protocol = atmp-protocol
primary-tunnel-server = ""
secondary-tunnel-server = ""
password = ""
client-auth-id = ""
server-auth-id = ""
assignment-id = ""
```

Parameter	Setting
`profile-type`	Type of tunneling profile. Set to `mobile-client` for PPP clients using L2TP tunneling.
`tunneling-protocol`	Protocol used to establish the tunnel. For L2TP tunneling, specify `l2tp-protocol`.
`primary-tunnel-server`	IP address or hostname of a tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn. If the primary server is unavailable, the system attempts to establish a tunnel to the secondary server.
`secondary-tunnel-server`	IP address or hostname of a tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn. If the primary server is unavailable, the system attempts to establish a tunnel to the secondary server.
`password`	Password used for authenticating the tunnel.
`client-auth-id`	Name sent to the tunnel server for authenticating the tunnel. The name can contain up to 31 characters. For details about L2TP-specific names used in tunnel authentication, see "L2TP-specific IDs for tunnel authentication" on page 5-13.

Parameter	Setting
`server-auth-id`	Name sent from the tunnel server to the system initiating the tunnel during the tunnel authentication phase. The name can contain up to 31 characters. Note that this field is currently ignored if it is specified in a `connection` profile. For details about L2TP-specific names used in tunnel authentication, see "L2TP-specific IDs for tunnel authentication" on page 5-13.
`assignment-id`	Identification (name) assigned to tunnels to allow grouping sessions, a text string of up to 31 characters. The value has local significance only. It is not transmitted to the remote tunnel end point.

RADIUS profiles use the following attribute-value pairs to specify L2TP tunnels::

RADIUS attribute	Value
Tunnel-Type (64)	Tunneling protocol to be used. Set to L2TP (3) for L2TP tunneling.
Tunnel-Medium-Type (65)	Media to be used for the tunnel. Only IP (1) is supported at this time.
Tunnel-Server-Endpoint (67)	IP address or hostname of the tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn.
Tunnel-Password (69)	Shared secret for authenticating the tunnel.
Tunnel-Client-Auth-ID (90)	Name sent to the tunnel end point by the system requesting the tunnel (the NAS or LAC) during the tunnel authentication phase. The name can contain up to 31 characters. See "How the system name is selected" on page 5-18.
Tunnel-Server-Auth-ID (91)	Name sent from the tunnel end point (the gateway or LNS) to the system initiating the tunnel during the tunnel authentication phase. The name can contain up to 31 characters. The attribute can be specified in Access Response packets and is generated in Accounting Request packets.

Typical client connection configuration

In these examples, the Stinger unit negotiates the PPP session, including password authentication, and then opens the L2TP tunnel. The following commands create a `connection` profile that includes a PPP password. The system authenticates the caller before bringing up a tunnel to an LNS at 1.1.1.1.

```
admin> read conn l2test
CONNECTION/l2test read
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp recv-password = localpw
```

```
admin> set tunnel profile-type = mobile-client
admin> set tunnel primary-tunnel-server = 1.1.1.1
admin> set tunnel tunneling-protocol = l2tp
admin> write
CONNECTION/l2test written
```

Following is a comparable RADIUS profile:

```
l2test Password = "localpw"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Tunnel-Server-Endpoint = "1.1.1.1",
   Tunnel-Type = L2TP,
   Tunnel-Medium-Type = IP
```

Typical connection to two tunnel end points

Figure 5-2 shows a Stinger unit that can connect to one of two possible LNS end points to create an L2TP tunnel for the dial-in client. In this example, the LNS end points are on remote networks, so the system requires a `connection` or RADIUS profile to establish a connection to one of the end-point systems.

Figure 5-2. Primary and secondary L2TP tunnel end points

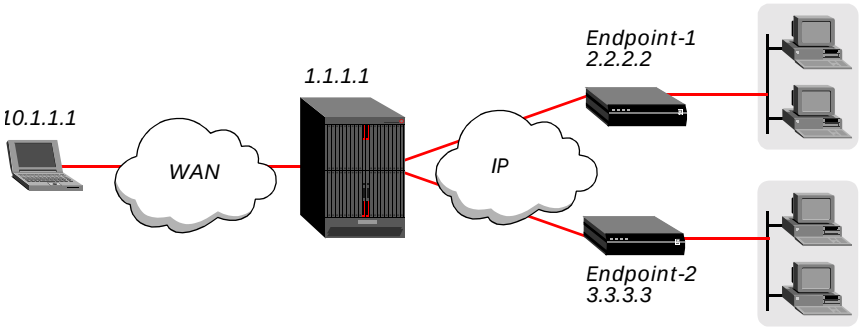

The following commands configure the Stinger unit's system IP address:

```
admin> read ip-global
IP-GLOBAL read
admin> set system-ip-addr = 1.1.1.1
admin> write
IP-GLOBAL written
```

The following commands configure `connection` profiles to the two LNS systems:

```
admin> read connection endpoint-1
CONNECTION/endpoint-1 read
admin> set active = yes
admin> set dial-number = 9-1-333-555-1212
admin> set ppp-options send-password = lns-pw
admin> set ppp-options recv-password = lac-pw
admin> set ip-options remote = 2.2.2.2
```

```
admin> write
CONNECTION/endpoint-1 written
admin> read connection endpoint-2
CONNECTION/endpoint-2 read
admin> set active = yes
admin> set dial-number = 9-1-123-555-1234
admin> set ppp-options send-password = lns-pw
admin> set ppp-options recv-password = lac-pw
admin> set ip-options remote = 3.3.3.3
admin> write
CONNECTION/endpoint-2 written
```

The following commands create a connection profile for the dial-in client:

```
admin> new connection dialin-1
CONNECTION/dialin-1 read
admin> set active = yes
admin> set clid = 555-1000
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp
admin> set tunnel-options primary-tunnel-server = 2.2.2.2
admin> set tunnel-options secondary-tunnel-server = 3.3.3.3
admin> write
CONNECTION/dialin-1 written
```

L2TP-specific IDs for tunnel authentication

In earlier tunneling implementations, because of constraints caused by L2TP and RADIUS protocol requirements, tunnel authentication required that every network access server (NAS) in the network used the same system name, even when the network spanned multiple administrative domains. These requirements now allow each NAS to send a unique system name for tunnel authentication purposes. The name can be specified on a per-connection or per-server basis. If RADIUS accounting is enabled, the system reports the names used for tunnel authentication in the Stop record.

Note Tunnel authentication occurs before a tunnel is established between two end points. It is negotiated between the LAC and a tunnel server and is independent of user authentication. If tunnel authentication fails, all pending calls associated with the tunnel are dropped.

For L2TP tunnels, because the LAC can now specify its name on a per-connection basis, you can configure profiles to create parallel tunnels to the same destination. For example, some sites use parallel tunnels to separate data streams that are directed to the same LNS but destined for different networks.

Summary of profile settings

For details about how the system uses these settings to determine whether to use an existing tunnel or start a new one, see “How the system finds a matching tunnel” on page 5-17. For details about how the system determines which name to use for tunnel authentication, see “How the system name is selected” on page 5-18.

Following are the parameters (shown with default values) for configuring the Stinger unit to use the enhanced tunnel authentication provided by this feature:

```
[in CONNECTION/"":tunnel-options]
profile-type = disabled
tunneling-protocol = atmp-protocol
primary-tunnel-server = ""
secondary-tunnel-server = ""
password = ""
client-auth-id = ""
server-auth-id = ""

[in TUNNEL-SERVER/""]
shared-secret = ""
client-auth-id = ""
server-auth-id = ""

[in L2-TUNNEL-GLOBAL]
l2tp-system-name = ""

[in SYSTEM]
name = ""
```

Parameter	Setting
`profile-type`	Type of tunneling profile. Must be set to `mobile-client` for PPP clients using L2TP tunneling.
`tunneling-protocol`	Protocol used to establish the tunnel. Must be set to `l2tp-protocol` for L2TP tunneling.
`primary-tunnel-server`	IP address or hostname of a tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn. If the primary server is unavailable, the system attempts to establish a tunnel to the secondary server.
`secondary-tunnel-server`	IP address or hostname of a tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn. If the primary server is unavailable, the system attempts to establish a tunnel to the secondary server.
`password`	Password used for authenticating the tunnel.
`client-auth-id`	Name sent to the tunnel server for authenticating the tunnel. The name can contain up to 31 characters.
`server-auth-id`	Name sent from the tunnel server to the system initiating the tunnel during the tunnel authentication phase. The name can contain up to 31 characters. Note that this field is currently ignored if it is specified in a `connection` profile.

Parameter	Setting
`shared-secret`	Shared secret for authenticating the tunnel.
`l2tp-system-name`	Name sent to the tunnel server for authenticating the tunnel if `client-auth-id` is not specified. See "How the system name is selected" on page 5-18.
`name`	Name sent to the tunnel server for authenticating the tunnel if `client-auth-id` is not specified, and `l2tp-system-name` is not specified. See "How the system name is selected" on page 5-18. If the domain name is configured in the `ip-global` profile, the specified system name is concatenated with the domain name.

RADIUS supports this feature by using the following attribute-value pairs. These attribute-value pairs support tag fields, as described in the Internet Draft `draft-ietf-radius-tunnel-auth-09.txt`. Each tag value (from 1 to 31) defines an independent tunnel attempt description. The Tunnel-Client-Auth-ID and Tunnel-Server-Auth-ID attributes can be specified in Access-Response packets and are generated in Accounting-Request packets.

RADIUS attribute	Value
Tunnel-Type (64)	Tunneling protocol(s) to be used. Must be set to L2TP (3) to use this feature.
Tunnel-Server-Endpoint (67)	IP address or hostname of the tunnel end point. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn.
Tunnel-Password (69)	Shared secret for authenticating the tunnel.
Tunnel-Client-Auth-ID (90)	Name sent to the tunnel end point by the system requesting the tunnel (the NAS or LAC) during the tunnel authentication phase. The name can contain up to 31 characters. See "How the system name is selected" on page 5-18.
Tunnel-Server-Auth-ID (91)	Name sent from the tunnel end point (the gateway or LNS) to the system initiating the tunnel during the tunnel authentication phase. The name can contain up to 31 characters. This attribute does not apply unless the protocol used to establish the tunnel is L2TP. The attribute can be specified in Access Response packets and is generated in Accounting Request packets.
Tunnel-Assignment-ID (82)	Identification (name) assigned to tunnels to allow grouping sessions. A text string of up to 31 characters. The value has local significance only. It is not transmitted to the remote tunnel end point.

Example of connection-based tunnel authentication

In the example shown in Figure 5-3, a PPP client dials into a Stinger unit to tunnel into its home network across the Internet.

Figure 5-3. Example of L2TP tunnel authentication

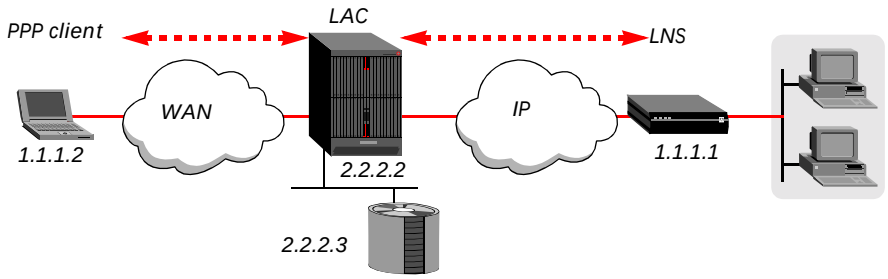

The following commands configure a local `connection` profile for the PPP client and specify a `client-auth-id` name:

```
admin> new connection ppp-user
CONNECTION/ppp-user read
admin> set ppp-options recv-password = test
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options password = conn-pass
admin> set tunnel-options client-auth-id = conn-lac
admin> write -f
CONNECTION/ppp-user written
```

Note that you need not assign an IP address because it is assigned by the LNS. Following is a comparable RADIUS profile:

```
ppp-user User-Password = "test", Service-Type = Dialout
   Tunnel-Type = L2TP,
   Tunnel-Server-Endpoint = 1.1.1.1,
   Tunnel-Password = conn-pass,
   Tunnel-Client-Auth-ID = conn-lac
```

With the sample profiles, the LAC authenticates the PPP client's session request. It then initiates a tunnel to the LNS if a tunnel does not already exist to that end-point address. When the Stinger unit requests the tunnel, it passes the LNS the string `conn-lac` as its local system name, and uses `conn-pass` as the password to authenticate the tunnel. The LNS uses the same strings to authenticate the LAC before establishing the tunnel.

Example of server-based tunnel authentication

The following commands configure a local `connection` profile for the PPP client and do not specify a password or `client-auth-id` name:

```
admin> new connection ppp-user
CONNECTION/ppp-user read
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = lns.example.com
```

```
admin> write -f
CONNECTION/ppp-user written
```

Following is a comparable RADIUS profile:

```
ppp-user User-Password = "", Service-Type = Dialout
   Tunnel-Type = L2TP,
   Tunnel-Server-Endpoint = lns.example.com
```

With the sample profiles, the LAC authenticates the PPP client's session request. It then initiates a tunnel to the LNS if a tunnel does not already exist to that end-point address. If tunnel authentication is enabled and no tunnel password is specified in the `connection` profile, the system searches for a `tunnel-server` profile before requesting the tunnel. If it finds a `tunnel-server` profile for the LNS, the system sends the `client-auth-id` to the LNS and the end points use the tunnel password (the shared secret) to authenticate the tunnel.

Following is a sample `tunnel-server` profile that specifies a password and local system name for use in tunnel authentication:

```
admin> new tunnel-server lns.example.com
TUNNEL-SERVER/lns.example.com read
admin> set shared-secret = ts-pass
admin> set client-auth-id = ts-lac
admin> write
TUNNEL-SERVER/lns.example.com written
```

Following is a comparable RADIUS profile:

```
lns.example.com Password = "", Service-Type = Dialout
   Tunnel-Password = ts-pass,
   Tunnel-Client-Auth-ID = ts-lac
```

Note If no `tunnel-server` profile exists, the LAC proceeds as described in "How the system name is selected" on page 5-18.

Examples of parallel L2TP tunnels to the same end point

After the LAC has authenticated a PPP client's session request, it looks for an existing tunnel that matches both the tunnel server end point and the `client-auth-id` specified in the client's profile. If the LAC finds an established tunnel that matches these values, it uses the tunnel. If it does not find a matching tunnel, it initiates a tunnel request. This process can be used to create parallel L2TP tunnels by specifying different `client-auth-id` values in profiles.

How the system finds a matching tunnel

If the client's profile specifies a hostname as the tunnel-server end point, the system must match both the hostname and the server's actual IP address to allow the client to use an established tunnel.

If `client-auth-id` is specified in the caller's profile, the system attempts to match the caller to an existing tunnel by using the following values:

- Tunnel server's IP address (and hostname, if specified)
- `client-auth-id`

If `client-auth-id` is *not* specified in the caller's profile, the system attempts to match the caller to an existing tunnel by using only the tunnel server's IP address (and hostname, if specified).

If the system finds a match on the basis of the values, it uses the tunnel. If the system does not find a matching tunnel entry, it initiates a new tunnel request.

How the system name is selected

If tunnel authentication is enabled and the Stinger unit is requesting a new tunnel, it looks for a system name to send to the LNS as follows:

1. Uses the `client-auth-id` if specified in the caller's `connection` profile. If `client-auth-id` is not specified in the `connection` profile, the system goes on to the next alternative.
2. Uses the `client-auth-id` if specified in the `tunnel-server` profile for the LNS. If `client-auth-id` is not specified in a `tunnel-server` profile, the system goes on to the next alternative.
3. Uses the `l2tp-system-name` if specified in the `l2-tunnel-global` profile. If `l2tp-system-name` is not specified in that profile, the system goes on to the next alternative.
4. Uses the `name` if specified in the `system` profile. If `name` is not specified in that profile, the system goes on to the next alternative.
5. Sends the string `noname`.

Examples of how **`client-auth-id`** settings create parallel tunnels

In this example, the LNS system's DNS hostname is `a.example.com` (a fully qualified domain name), which resolves to two IP addresses, 1.1.1.1 and 1.1.1.2. The hostname `b.example.net` also resolves to the 1.1.1.1 address. Table 5-1 shows existing tunnels to the LNS, which were authenticated by different `client-auth-id` strings.

Table 5-1. Existing tunnels to the same LNS

Address	client-auth-id	tunnel-server	Tunnel ID
1.1.1.1	a1	a.example.com	102
1.1.1.1	a2	a.example.com	103

Table 5-2 shows how the system matches the values in the clients' profiles as it receives incoming calls, and the resulting action the system takes in terms of using an existing tunnel or creating a new one.

Table 5-2. Tunnels created for incoming callers based on profile settings

Values used to match tunnel:			**Resulting action**	**Tunnel ID**
Address	**client-auth-id**	**tunnel-server**		
1.1.1.1	a1	a.example.com	Reuse tunnel	102
1.1.1.1	a2	a.example.com	Reuse tunnel	103
1.1.1.1	b	b.example.net	Establish new tunnel	104
1.1.1.1	b	a.example.com	Establish new tunnel	105
1.1.1.1		a.example.com	Reuse tunnel	102 or 103
1.1.1.1	a2	b.example.net	Establish new tunnel	106
1.1.1.2	a1	a.example.com	Establish new tunnel	107

Note The caller that does not supply a `client-auth-id` string matches the tunnel-server end point, so the existing tunnel to that end point (Tunnel ID `102`) is reused.

Examples of configuration errors causing multiple tunnels

Configuration errors can lead to unintentional parallel tunnels to the same tunnel end point. For this reason, you must either use the `client-auth-id` setting for all user profiles to a particular LNS or decide *not* to use that setting for callers tunneling to that LNS.

For example, suppose your RADIUS users file contains the following two user profiles and tunnel server profile:

```
user1 Password = userpass
   Tunnel-Type = L2TP,
   Tunnel-Server-Endpoint = lns.example.com,
   Tunnel-Client-Auth-ID = A-LAC,
   ...
user2 Password = userpass
   Tunnel-Type = L2TP,
   Tunnel-Server-Endpoint = lns.example.com,
   ...
```

```
lns.example.com User-Password = "", Service-Type = Dialout
   Tunnel-Password = tunpass,
   Tunnel-Client-Auth-ID = AllMyLACs
```

If `user1` calls in first and establishes a tunnel, `user2` can reuse that tunnel, as shown in Table 5-3.

Table 5-3. Tunnels created when user1 dials in first (configuration error not detected)

Values used to match tunnels:			Resulting action	Tunnel ID
Address	**client-auth-id**	**tunnel-server**		
2.2.2.2	A-LAC	lns.example.com	Create new tunnel	88
2.2.2.2		lns.example.com	Reuse tunnel	88

However, if `user2` calls in first and establishes a tunnel, the system obtains a system name for authentication from the `tunnel-server` profile. When `user1` dials in, the caller is unable to reuse the tunnel, because the authentication names do not match. This result is shown in Table 5-4.

Table 5-4. Tunnels created when user2 dials in first (configuration error shown)

Values used to match tunnels:			Resulting action	Tunnel ID
Address	**client-auth-id**	**tunnel-server**		
2.2.2.2	AllMyLACs	lns.example.com	Create new tunnel	40
2.2.2.2	A-LAC	lns.example.com	Create new tunnel	42

Tunnel assignment IDs

Like the `client-auth-id` described in "L2TP-specific IDs for tunnel authentication" on page 5-13, the tunnel assignment ID helps the LAC determine whether to assign a client session to an existing tunnel or to create a new one. This value is also used for grouping client sessions into specific tunnels. For details, see `draft-ietf-radius-tunnel-auth-09.txt`.

How tunnel assignment IDs affect tunnel matching

When selecting an existing tunnel or deciding to create a new one, the system first checks the criteria described in "How the system finds a matching tunnel" on page 5-17, and then performs an additional, final check for a tunnel assignment ID.

After comparing the tunnel transport address and, if specified in the client's profile, the tunnel server's hostname (`server-endpoint`) against existing tunnels, the system begins comparing the following optional parameters, in the order shown:

- `client-auth-id` specified in the client's profile and the `client-auth-id` used for existing tunnels
- `assignment-id` specified in the client's profile and the tunnel assignment ID of existing tunnels

The client profile matches existing tunnels only if both the `client-auth-id` and the tunnel assignment ID match. A null value in any one of these fields in an existing tunnel matches only a null value in the corresponding parameter in the client profile. If the system does not find a matching tunnel entry, it initiates a new tunnel request.

Overview of local profile settings

The following new parameters (shown with default values) have been added to local profiles:

```
[CONNECTION/"":tunnel-options]
assignment-id = ""
```

Parameter	Setting
`assignment-id`	Identification (name) assigned to tunnels to allow grouping sessions. A text string of up to 31 characters. The value has local significance only. It is not transmitted to the remote tunnel end point.

Overview of RADIUS attribute-value pair

RADIUS supports this feature by means of the following attribute-value pair:

RADIUS attribute	Value
Tunnel-Assignment-ID (82)	Identification (name) assigned to tunnels to allow grouping sessions. A text string of up to 31 characters. The value has local significance only. It is not transmitted to the remote tunnel end point.

Example of configuring a tunnel assignment ID

In this example, the Stinger unit is configured to perform tunnel authentication for L2TP tunnels. The two PPP clients shown in Figure 5-4 are configured to use different tunnels to the LNS on the basis of their tunnel assignment IDs. (The same clients can be configured to use the same multiplexed tunnel if their tunnel assignment IDs are set to the same string.)

Figure 5-4. L2TP tunnel setup that uses tunnel assignment IDs

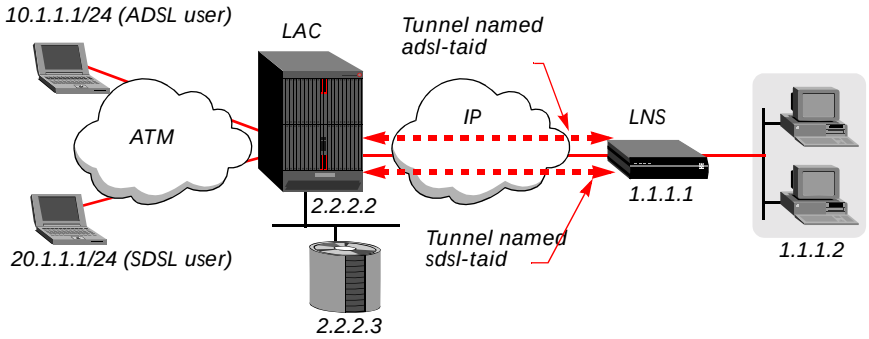

The following set of commands enables tunnel authentication:

```
admin> read l2-tunnel-global
L2-TUNNEL-GLOBAL read
admin> set l2tp-mode = lac
admin> set l2tp-auth-enabled = yes
admin> write
L2-TUNNEL-GLOBAL written
```

The following set of commands creates local connection profiles for the two mobile clients:

```
admin> new connection adsluser
CONNECTION/adsluser read
admin> set ppp-options recv-password = test
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options password = shared
admin> set tunnel-options client-auth-id = taos-unit
admin> set tunnel-options assignment-id = adsl-taid
admin> write
CONNECTION/adsluser written
admin> new connection sdsluser
CONNECTION/sdsluser read
admin> set ppp-options recv-password = test
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options tunneling-protocol = l2tp-protocol
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options password = shared
admin> set tunnel-options client-auth-id = taos-unit
admin> set tunnel-options assignment-id = sdsl-taid
admin> write
CONNECTION/sdsluser written
```

Following are comparable RADIUS profiles:

```
adsluser Password = "test"
    User-Service = Framed-User,
    Framed-Protocol = PPP,
    Test-Idle-Limit = 0,
    Tunnel-Type = L2TP :1,
    Tunnel-Server-Endpoint = 1.1.1.1 :1,
    Tunnel-Client-Auth-ID = taos-unit: 1,
    Tunnel-Password = shared,
    Tunnel-Assignment-ID = adsl-taid:1
sdsluser Password = "test"
    User-Service = Framed-User,
    Framed-Protocol = PPP,
    Test-Idle-Limit = 0,
```

```
Tunnel-Type = L2TP :1,
Tunnel-Server-Endpoint = 1.1.1.1 :1,
Tunnel-Client-Auth-ID = taos-unit: 1,
Tunnel-Password = shared,
Tunnel-Assignment-ID = sdsl-taid:1
```

RADIUS accounting support

RADIUS accounting Stop records display the Tunnel-Assignment-ID used for the user session. For example:

```
Tue May 2 15:58:08 2000
        User-Name = "adsluser"
        NAS-Identifier = 2.2.2.2
        NAS-Port = 11313
        NAS-Port-Type = Async
        Acct-Status-Type = Stop
        Acct-Delay-Time = 0
        Acct-Session-Id = "317658341"
        Acct-Authentic = Local
        Acct-Session-Time = 112
        Acct-Input-Octets = 2155
        Acct-Output-Octets = 513
        Acct-Input-Packets = 23
        Acct-Output-Packets = 14
        Ascend-Disconnect-Cause = 185
        Ascend-Connect-Progress = 60
        Ascend-Xmit-Rate = 28800
        Ascend-Data-Rate = 33600
        Ascend-PreSession-Time = 19
        Ascend-Pre-Input-Octets = 0
        Ascend-Pre-Output-Octets = 0
        Ascend-Pre-Input-Packets = 0
        Ascend-Pre-Output-Packets = 0
        Ascend-Modem-PortNo = 1
        Ascend-Modem-SlotNo = 7
        Ascend-Modem-ShelfNo = 1
        Tunnel-Type = L2TP
        Tunnel-Server-Endpoint = "1.1.1.1"
        Tunnel-Client-Auth-ID = "taos-unit"
        Tunnel-Server-Auth-ID = "max6k-lns"
        Tunnel-Assignment-ID = "adsl-taid"
```

ATMP Tunneling

6

The T1000 module supports Ascend Tunnel Management Protocol (ATMP) for virtual private network (VPN) connectivity. For information about using L2TP tunneling protocols as an alternate method for VPN connectivity, see Chapter 5, “L2TP Tunneling.”

Introduction to ATMP

ATMP is a UDP/IP-based protocol for tunneling between two TAOS units across an IP network. Data is transported through the tunnel in Generic Routing Encapsulation (GRE), as described in RFC 1701. (For a complete description of ATMP, see RFC 2107, *Ascend Tunnel Management Protocol - ATMP.*)

Figure 6-1 shows one use for ATMP tunneling: Mobile clients dial in to a local ISP to log in to a distant LAN across the Internet. ATMP creates and tears down a cross-Internet tunnel between the two TAOS units. In effect, the tunnel collapses the IP cloud and provides what looks like direct access to a home network.

Figure 6-1. ATMP tunnel from an ISP to a corporate home network

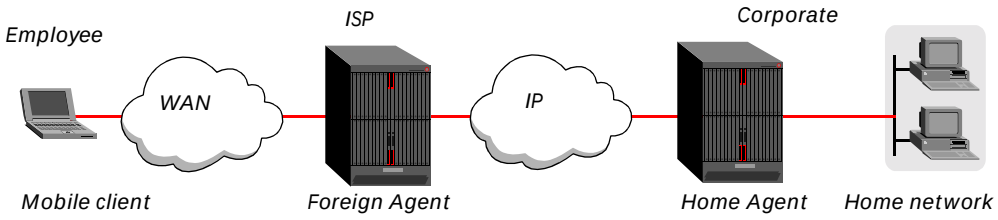

A mobile client dials in to the *Foreign Agent*, which authenticates the `connection` profile (or RADIUS profile) and initiates an IP connection to the specified *Home Agent*.

The Foreign Agent then requests a tunnel for the connected mobile client. The Home Agent authenticates the tunnel request (by password), and then registers the tunnel

and assigns it an ID. If the Home Agent refuses the tunnel, the Foreign Agent disconnects the mobile client.

If the tunnel is successfully established, the Home Agent forwards or routes tunneled data to the home network. If the mobile client has a multichannel MP+ or MP connection, the tunnel remains active when the connection adds or subtracts channels, and is not torn down until the final channel of the call is disconnected.

The Home Agent must be able to access the home network either as an ATMP gateway or by routing the packets. For a description of how the Home Agent operates as a gateway or router, see "Home Agent atmp profile settings" on page 6-18.

If an ATMP client disconnects because of an ATMP error, ATMP disconnect codes can help you diagnose the exact cause of the problem. Each code can appear in a `syslog` record or as the value of Ascend-Disconnect-Cause (195) in a RADIUS accounting record. For additional information about disconnect codes, see the *TAOS RADIUS Guide and Reference*.

Network settings for ATMP

Network settings for ATMP include settings related to the IP connection between TAOS units, settings related to the UDP communication required to establish tunnels, and settings related to packet fragmentation and reassembly.

System reset requirement

When you change the setting of the `udp-port` parameter in the `atmp` profile of a Home Agent, a system reset is required for the ATMP subsystem to recognize the new UDP port number.

When you change the `agent-mode` parameter from its default `tunnel-disabled` setting to any other setting, a system reset is required for the new value to take effect.

All other parameter settings in the `atmp` profile take effect as soon as possible after the profile is written.

System IP address recommendation

Lucent Technologies recommends that you set the `system-ip-addr` parameter in the `ip-global` profile, on a Stinger unit that is operating as an ATMP agent, particularly if the unit has multiple interfaces into the IP cloud that separates it from other ATMP agents. This recommendation has two aspects:

- On a Foreign Agent, in the `connection` profile for mobile clients, specify the system IP address of a Home Agent rather than the interface address on which the Home Agent accepts tunneled data. This setting helps to avoid communication problems if a route changes within the IP cloud.
- On both a Foreign Agent and a Home Agent, the link to the other agent can specify the unit's system IP address. This setting is not required if RIP is enabled on the interfaces between the two agents, but it is recommended, because it helps to simplify configurations.

Figure 6-2 shows a Home Agent and Foreign Agent, with two Ethernet interfaces connecting them. (The principle is the same as if there were two WAN connections between the units.)

Figure 6-2. System IP addresses and routes between ATMP agents

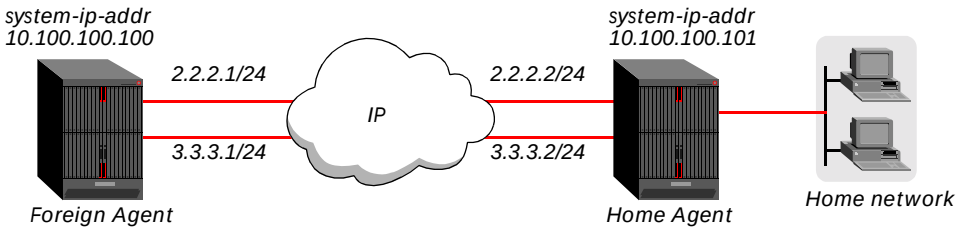

When RIP is enabled on the IP interfaces between the two units, it advertises the system address on both ports. For example, suppose a Foreign Agent has the following system IP address and IP interface configuration:

```
[in IP-GLOBAL]
system-ip-addr = 10.100.100.100

[in IP-INTERFACE { {shelf-1 slot-3 1} 0 } ]
ip-address = 2.2.2.1/24
rip = both-v2

[in IP-INTERFACE { {shelf-1 slot-3 2} 0 } ]
ip-address = 3.3.3.1/24
rip = both-v2
```

Supposing a Home Agent has the following system IP address and IP interface configuration:

```
[in IP-GLOBAL]
system-ip-addr = 10.100.100.101

[in IP-INTERFACE { {shelf-1 slot-3 1} 0 } ]
ip-address = 2.2.2.2/24
rip = both-v2

[in IP-INTERFACE { {shelf-1 slot-3 2} 0 } ]
ip-address = 3.3.3.2/24
rip = both-v2
```

With this configuration, the Foreign Agent advertises, on both of its Ethernet ports, a route to its own system address, 10.100.100.100. Similarly, the Home Agent advertises, on both of its Ethernet ports, a route to its own system address, 10.100.100.101.

When the Home Agent receives the advertisements for 10.100.100.100, it selects one of the ports advertising the route and adds that route to its routing table. The next time the Home Agent establishes a connection with the Foreign Agent, it uses the port indicated in the routing table. If that port becomes unavailable (for example, if the cable is disconnected), the Home Agent soon updates its routing table to use the other port to connect to the Foreign Agent.

Setting the UDP port

By default, ATMP agents use UDP port 5150 to exchange control information while establishing a tunnel. If the Home Agent `atmp` profile specifies a different UDP port number, all tunnel requests to that Home Agent must specify the same UDP port.

Note A system reset is required for the ATMP subsystem to recognize the new UDP port number.

Specifying tunnel retry limits

The `retry-timeout` and `retry-limit` parameters in the `atmp` profile work together to limit how many tunnel RegisterRequest messages (to open a tunnel) and DeregisterRequest messages (to close a tunnel) are sent and the number of seconds between each message. If a tunnel request fails, the Foreign Agent times out, logs a message, and disconnects the mobile client. When a tunnel request succeeds, the Home Agent assigns a tunnel ID, and the UDP port is no longer used for that tunnel. The actual data transfer uses the IP connection with GRE.

The `retry-timeout` and `retry-limit` parameters have default settings that are appropriate for most sites, but you might want to increase or decrease the values on the basis of what type of link connects the Foreign Agent and Home Agent. For example, if the Foreign Agent and the Home Agent are on the same Ethernet segment, you might want to reduce the values to provide a quicker response to the mobile client when the Home Agent is unavailable.

If you increase the `retry-timeout` and `retry-limit` values, keep in mind that the values determine response time to mobile clients when the Home Agent is unavailable. If a tunnel request reaches a secondary Home Agent that is also unavailable, the mobile client waits for twice the specified period before being informed that the connection failed.

Setting an MTU limit

The type of link that connects a Foreign Agent and Home Agent determines the maximum transmission unit (MTU). The link can be a WAN or Ethernet link, and it can be a local network or routed through multiple hops. If the link between devices is multihop (if it traverses more than one network segment), the path MTU is the *minimum* MTU of the intervening segments.

Figure 6-3 shows an ATMP setup across a 100BaseT Ethernet segment, which limits the path MTU to 1500 bytes.

Figure 6-3. Path MTU on an Ethernet segment

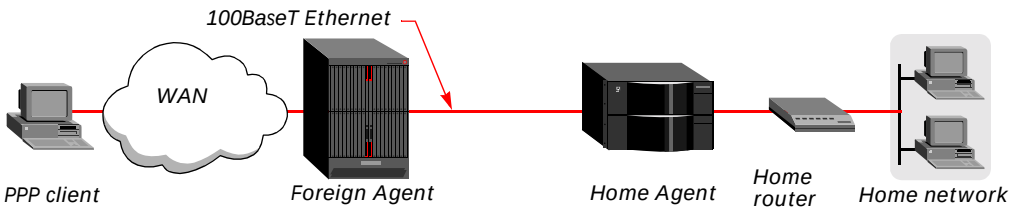

If any segment of the link between the agents has an MTU smaller than 1528, some packet fragmentation and reassembly will occur. You can push fragmentation and reassembly tasks to connection end points (a mobile client and a device on the home

network) by setting an MTU limit. Client software then uses MTU discovery mechanisms to determine the maximum packet size, and fragments packets before sending them.

How link compression affects the MTU

Compression affects which packets must be fragmented, because compressed packets are shorter than their original counterparts. If any kind of compression is enabled (such as VJ header or link compression), the connection can transfer larger packets without exceeding a link's maximum receive unit (MRU). If a compressed packet is smaller than the MRU, it can be sent across the connection, whereas the same packet without compression cannot.

How ATMP tunneling causes fragmentation

To transmit packets through an ATMP tunnel, the Stinger unit adds an 8-byte GRE header and a 20-byte IP header to the frames it receives. The addition of these packet headers can make the packet larger than the MTU of the tunneled link, in which case the unit must either fragment the packet after encapsulating it or reject the packet.

Fragmenting packets after encapsulating them has several disadvantages for the Foreign Agent and Home Agent. For example, it causes a performance degradation, because both agents have extra overhead. It also means that the Home Agent device cannot be a GRF® switch. (To maintain its very high aggregate throughput, a GRF® switch does not perform reassembly.)

Pushing the fragmentation task to connection end points

To avoid the extra overhead incurred when ATMP agents perform fragmentation, you can either set up a link between the two units that has an MTU greater than 1528 (which means it cannot include Ethernet segments), or you can set the `mtu-limit` parameter in the `atmp` profile to a value that is 28 bytes less than the path MTU.

If `mtu-limit` is set to zero (the default), the Stinger unit might have to fragment encapsulated packets before transmission. The other ATMP agent must then reassemble the packets.

If `mtu-limit` is set to a nonzero value, the unit reports that value to the client software as the path MTU, causing the client to send packets of the specified size. This setting pushes the task of fragmentation and reassembly out to the connection end points, lowering the overhead of the ATMP agents.

For example, if the Stinger unit is communicating with another ATMP agent across an Ethernet segment, you can set the `mtu-limit` parameter to a value 28 bytes smaller than 1500 bytes, as shown in the following example, to enable the unit to send unfragmented packets that include the 8-byte GRE header and a 20-byte IP header:

```
admin> read atmp
ATMP read

admin> set mtu-limit = 1472

admin> write
ATMP written
```

With this setting, the connection end point sends packets with a maximum size of 1472 bytes. When the Stinger unit encapsulates them, adding 28 bytes to the size, the packets still do not violate the 1500-byte Ethernet MTU.

Forcing fragmentation for interoperation with outdated clients

To discover the path MTU, some clients normally send packets that are larger than the negotiated maximum receive unit (MRU) and that have the Don't Fragment (DF) bit set. Such packets are returned to the client with an ICMP message informing the client that the host is unreachable without fragmentation. This standard, expected behavior improves end-to-end performance by enabling the connection end points to perform any required fragmentation and reassembly.

However, some outdated client software does not handle this process correctly and continues to send packets that are larger than the specified `mtu-limit`. To enable the Stinger unit to interoperate with these clients, you can configure the unit to ignore the DF bit and perform the fragmentation that normally should be performed by the client software. This function is referred to as *prefragmentation*.

When the `mtu-limit` parameter is set to a nonzero value, you can set the `force-fragmentation` parameter to `yes` to enable the Stinger unit to prefragment packets it receives that are larger than the negotiated MRU with the DF bit set. It prefragments those packets, and then adds the GRE and IP headers.

Note Setting the `force-fragmentation` parameter to `yes` causes the Stinger unit to bypass the standard MTU discovery mechanism and fragment larger packets before encapsulating them in GRE. Because this changes expected behavior, it is not recommended except for interoperation with outdated client software that does not handle fragmentation properly.

Mobile clients with duplicate IP addresses

A Foreign Agent accepts multiple mobile-client connections with duplicate IP addresses, as long as they request different Home Agents or home networks. This behavior allows the use of unregistered IP addresses by multiple independent private networks.

A Home Agent does not accept multiple mobile-client connections to the same home network with duplicate IP addresses or overlapping subnet ranges. If a mobile client attempts to connect to a Home Agent with an address that duplicates or is within the same subnet of an established mobile-client connection, the Home Agent immediately terminates the *existing* client connection. This behavior allows a mobile client to reconnect if its connection is lost because a Foreign Agent became unavailable.

Network isolation and duplicate IP addresses

ATMP isolates home networks from each other as well as from other IP networks between the Foreign Agent and Home Agents. A Foreign Agent can therefore accept multiple client connections that have the same IP address. For example, Figure 6-4 shows two mobile clients with the same IP address tunneling to two different home networks. The home networks are isolated from each other and from the IP cloud between the tunnel end points.

Figure 6-4. Foreign Agent supporting duplicate client IP addresses

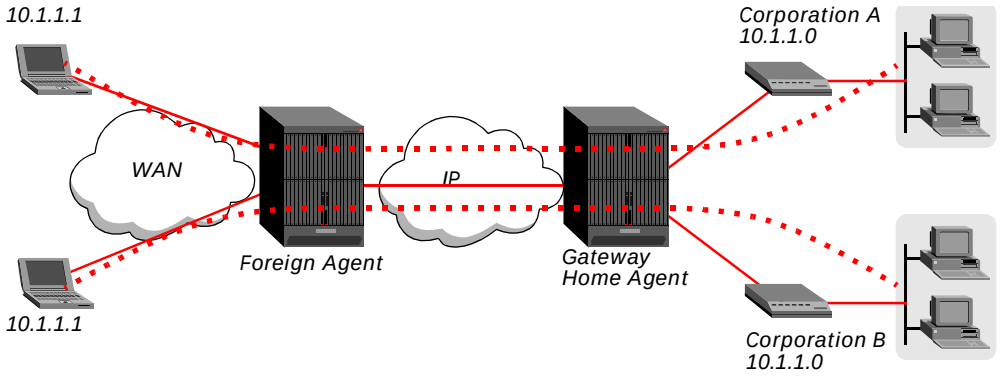

To provide network isolation, a Foreign Agent does not create routes for mobile clients. Similarly, gateway Home Agents do not create routes for ATMP gateway connections or for registered mobile clients. (However, router Home Agents *do* create routes for registered mobile clients.) Network isolation is also the reason why a mobile client or a home network router does not receive a response when attempting to ping a Foreign Agent or Home Agent.

Duplicate addresses connecting to the same home network

If a mobile client attempts to connect to a home network with an address that duplicates or is within the same subnet of an established mobile-client connection, the Home Agent immediately terminates the established connection and negotiates the incoming request. This behavior is required to enable a mobile client to reconnect if its connection is terminated when a Foreign Agent becomes unavailable.

For example, supposing a mobile client is connected to a home network with the following address:

```
10.10.10.10/24
```

The client's subnet range includes the addresses from 10.10.10.0 to 10.10.10.255. Supposing a second mobile client attempts to connect with the following address, which occupies the same subnet range as the first client:

```
10.10.10.199/24
```

The Home Agent terminates the first connection and allows the second mobile client to connect.

Configuring the agent-to-agent connection

The link between a Foreign Agent and Home Agent can be any kind of connection (switched, nailed, frame relay, and so forth) or an Ethernet link. It can be on a local network or routed through multiple hops. The only requirement is that the two units must communicate over an IP network.

For example, the following commands on a Home Agent configure an IP connection to a Foreign Agent. In this case, the Home Agent uses the `atmpfa` profile to authenticate the Foreign Agent dialing in.

```
admin> new connection atmpfa
CONNECTION/atmpfa read

admin> set active = yes
```

```
admin> set ppp send-auth = chap-ppp-auth
admin> set ppp send-password = remotepw
admin> set ppp recv-password = localpw
admin> set ip-options remote-address = 1.1.1.1
admin> write
CONNECTION/atmpfa written
```

For details about IP connections, see Chapter 1, "IP Routing."

Note If the Foreign Agent and Home Agent reside on the same Ethernet and use RADIUS authentication, you must use separate RADIUS servers for the tunnel end points to avoid session loopbacks.

Configuring a Foreign Agent

To configure a Foreign Agent, you must set parameters in the `atmp` profile, configure a `connection` or RADIUS profile for the connection to the Home Agent, and configure mobile-client `connection` or RADIUS profiles.

For information about configuring a connection to the Home Agent, see "Configuring the agent-to-agent connection" on page 6-7.

Foreign Agent atmp profile settings

The `atmp` profile contains the following parameters (shown with sample values) related to a Foreign Agent configuration:

```
[in ATMP]
agent-mode = foreign-agent
retry-timeout = 3
retry-limit = 10
mtu-limit = 0
force-fragmentation = no
```

Parameter	Setting for Foreign Agent configuration
`agent-mode`	Must specify `foreign-agent`.
`retry-timeout` `retry-limit`	Together, these parameters specify how many tunnel RegisterRequest and DeregisterRequest messages are sent and the number of seconds between each message. They have default settings that are appropriate for most sites. For details, see "Specifying tunnel retry limits" on page 6-4.
`mtu-limit`	Maximum transmission unit (MTU) for the path between the foreign and Home Agents. For details, see "Setting an MTU limit" on page 6-4.
`force-fragmentation`	If outdated client software sends large packets with the DF bit set, you can set this parameter to force the Stinger unit to fragment the packets anyway. For details, see "Forcing fragmentation for interoperation with outdated clients" on page 6-6.

Mobile client profile settings

All mobile-client profiles reside on the Foreign Agent side of the ATMP tunnel. A Foreign Agent can authenticate a mobile client locally in a `connection` profile or externally in a RADIUS profile.

Settings in connection profiles

The `tunnel-options` subprofile of a local `connection` profile contains the following parameters (shown with sample values) related to a mobile-client connection:

```
[in CONNECTION/mclient-1:tunnel-options]
profile-type = mobile-client
primary-tunnel-server = 2.2.2.2:8877
secondary-tunnel-server = 3.3.3.3:1555
udp-port = 5150
password = tunnel-password
home-network-name = ""
```

Parameter	Setting for mobile client configuration
`profile-type`	Must specify `mobile-client`.
`primary-tunnel-server`	Must specify the system IP address or hostname of a Home Agent.
`secondary-tunnel-server`	Specifies the system IP address or hostname of a secondary Home Agent. If a tunnel request to the first Home Agent fails, the Foreign Agent tries again with this host.
`udp-port`	Specifies a UDP port for one or both of the specified Home Agents. If the Home Agent specification includes a port number, that value overrides this parameter.
`password`	Must specify the password, if any, that is in the `atmp` profile of the Home Agent (up to 21 characters).
`home-network-name`	If the Home Agent is operating in gateway mode, must specify the name of the gateway profile that defines the connection to the home network.

Settings in RADIUS profiles

RADIUS uses the following attribute-value pairs to specify mobile-client connections:

RADIUS attribute	Value
Tunnel-Type (64)	Type of protocol used for the tunnel. To ensure forward compatibility, the TAOS-specific Tunneling-Protocol (127) attribute is converted into Tunnel-Type (value 4 means ATMP). To maintain backward compatibility, RADIUS accounting still generates the Tunneling-Protocol attribute.

RADIUS attribute	Value
Tunnel-Server-Endpoint (67)	System IP address or hostname of a Home Agent. The string can be followed by a colon and the UDP port number used on the ATMP Home Agent. To ensure forward compatibility, the Ascend-specific Ascend-Primary-Home-Agent (129) attribute is converted into Tunnel-Server-Endpoint.
Ascend-Secondary-Home-Agent (130)	System IP address or hostname of a secondary Home Agent. If a tunnel request fails with the first Home Agent, the Foreign Agent tries again with this host.
Ascend-Home-Agent-UDP-Port (186)	UDP port for one or both of the specified Home Agents. If the Home Agent specification includes a port number, that value overrides this parameter.
Tunnel-Password (69)	Password, if any, in the `atmp` profile of the Home Agent (up to 21 characters). To ensure forward compatibility, the Ascend-specific Home-Agent-Password (184) attribute is converted into Tunnel-Password. For more details, see "Tunnel password authentication" on page 6-10.
Tunnel-Private-Group-ID (81)	If the Home Agent is operating in gateway mode, you must use this attribute or the vendor-specific Ascend-Home-Network-Name (185) attribute to specify the name of the gateway profile that defines the connection to the home network.

When a standard RADIUS attribute for tunneling is available, you can specify either the standard attribute or the Ascend vendor attribute. For example, the following RADIUS profiles are equivalent:

```
user1 Password = "pass1"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.255.255.255,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "atmp-ha1.example.com",
   Tunnel-Password = "tunnel-password"

user1 Password = "pass1"
   Service-Type = Framed-User,
   Framed-Protocol = PPP,
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.255.255.255,
   Tunneling-Protocol = ATMP,
   Ascend-Primary-Home-Agent = "atmp-ha1.example.com",
   Ascend-Home-Agent-Password = "tunnel-password"
```

Tunnel password authentication

The Home Agent `atmp` profile contains a `home-agent-password` parameter. If it is not null, mobile client profiles must supply the password to initiate a tunnel. If the Foreign Agent supplies the proper password when requesting a tunnel, the Home

Agent returns a RegisterReply message with a number that identifies the tunnel, and the mobile client's tunnel is established. If the password does not match, the Home Agent rejects the tunnel, and the Foreign Agent logs a message and disconnects the mobile client. The following sample commands configure the Home Agent `atmp` profile to require tunnel authentication:

```
admin> read atmp
ATMP read

admin> set password = tunnel-password

admin> write
ATMP written
```

The `password` parameter in the mobile-client's `connection` profile must specify the same value. For example:

```
admin> read connection mobile-client
CONNECTION/mobile-client read

admin> set tunnel profile-type = mobile-client

admin> set tunnel primary-tunnel-server = 3.3.3.3:8877

admin> set tunnel password = tunnel-password

admin> write
CONNECTION/mobile-client written
```

Following is a comparable RADIUS profile:

```
mobile-client Password = "my-password",
   Service-Type = Framed-User
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "3.3.3.3:8877",
   Tunnel-Password = "tunnel-password"
```

Many RADIUS servers encrypt tunnel passwords before sending them to the Home Agent if the mobile-client profile uses the Tunnel-Password (69) attribute to specify the password. If the profile specifies a value for Tunnel-Password and the RADIUS server does not encrypt the password, tunnel authentication will fail.

If, instead, the mobile-client profile uses the Ascend-Home-Agent-Password (184) attribute to specify the password, the RADIUS server performs no encryption before sending the password to the Home Agent. This option might be required if you are using a RADIUS server that does not encrypt the Tunnel-Password value.

Note Unless you are using a RADIUS server that does not support tunnel password encryption (or encryption is not required), use of the Tunnel-Password attribute is recommended in place of Ascend-Home-Agent-Password to protect against local sniffers detecting tunnel passwords.

Specifying Home Agent addresses and port numbers

When a mobile client connects to a Foreign Agent, the Foreign Agent sends an ATMP RegisterRequest message to the IP address of the primary Home Agent. (If the Home Agent is specified as a hostname, the Foreign Agent first performs a DNS lookup.) Depending on the network configuration, the Foreign Agent might dial a connection to reach the Home Agent.

If the Foreign Agent does not receive a response to its request, it tries again. The number of retries is controlled by the `retry-limit` setting in the Foreign Agent's `atmp` profile.

If the Foreign Agent still does not receive a response or if it receives a negative response (such as Home Network Unreachable), it attempts to repeat the procedure with the secondary Home Agent address. If there is no secondary Home Agent address specified or if the registration with the secondary Home Agent also fails, the mobile client is disconnected.

If the Home Agent `atmp` profile specifies a UDP port number other than the default of 5150, you can specify that port number as part of the Home Agent address by appending a colon character (:) followed by the port number. The following commands specify the system IP address followed by a UDP port number for a primary and secondary Home Agent:

```
admin> read connection user1
CONNECTION/user1 read
admin> set ip-options remote-address = 10.1.1.1/32
admin> set tunnel profile-type = mobile-client
admin> set primary-tunnel-server = 2.2.2.2:8877
admin> set secondary-home-agent = 3.3.3.3:4000
admin> write
CONNECTION/user1 read
```

Or, in a RADIUS profile:

```
user1 Password = "pass1"
   Service-Type = Framed-User,
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.255.255.255,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "2.2.2.2:8877",
   Ascend-Secondary-Home-Agent = "3.3.3.3",
   Ascend-Home-Agent-UDP-Port = 4000
```

In this case, the Foreign Agent dials the connection to the primary Home Agent and requests a tunnel on port 8877. If that attempt fails, it dials the connection to the secondary Home Agent and requests a tunnel on port 4000. (If the address does not specify a port number, the Foreign Agent uses the value of the `udp-port` parameter in the mobile client `connection` profile.) For example, with the following settings, the Foreign Agent dials the connection to the Primary tunnel server and requests a tunnel on port 8877:

```
admin> set primary-tunnel-server = 2.2.2.2
admin> set secondary-tunnel-server = ha2.company.com:6789
admin> set udp-port = 8877
```

If that attempt fails, the Foreign Agent dials the connection to the secondary tunnel server and requests a tunnel on port 6789.

Specifying the home network name

For definitions of gateway and router Home Agents, see "Home Agent atmp profile settings" on page 6-18. For a mobile client tunnel to a gateway Home Agent, you

must specify the name of the gateway profile for connection to the home network. For example, suppose you are creating the following gateway profile on a Home Agent:

```
admin> new connection homenet
CONNECTION/homenet read
admin> set active = yes
admin> set tunnel profile-type = gateway-profile
admin> set telco call-type = ft1
admin> set telco nailed-groups = 7
admin> write
CONNECTION/homenet written
```

In the mobile client's profile, you would specify the following home network name:

```
admin> set home-network-name = homenet
```

Or you would include one of the following settings in a RADIUS profile:

```
Tunnel-Private-Group-ID = "homenet"
Ascend-Home-Network-Name = "homenet"
```

Note If the mobile client tunnels to a router Home Agent, you must, in the mobile-client profiles, leave the `home-network` parameter blank or omit the Tunnel-Private-Group-ID or Ascend-Home-Network-Name attributes.

Typical Foreign Agent configuration

Figure 6-5 shows a Foreign Agent that connects to two Home Agents across IP WAN connections. One is a gateway Home Agent and the other is a router Home Agent. The illustration also shows two mobile-client connections, one to each of the Home Agents.

Figure 6-5. Foreign Agent tunneling to two Home Agents

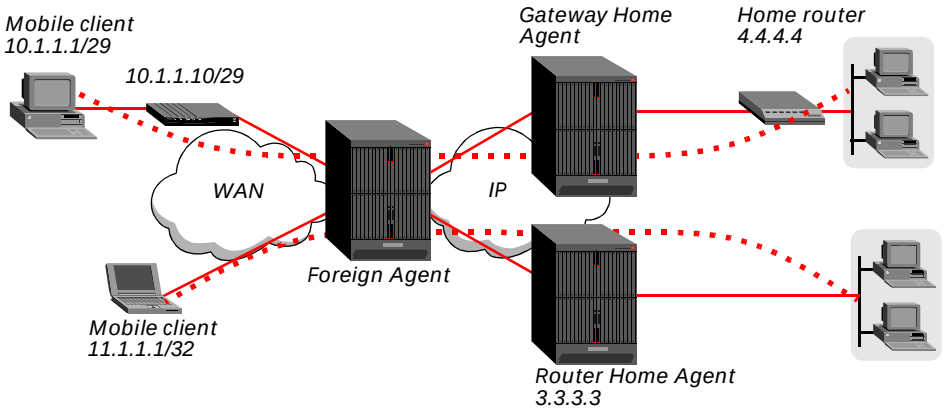

In this example, the WAN connections are multichannel PPP connections, which typically negotiate a path MTU of 1500 bytes. The agents set the `mtu-limit` to 1472, to enable the connection end points to fragment packets at that size. For background information, see "Setting an MTU limit" on page 6-4.

Setting the Foreign Agent system address

The following commands set the Foreign Agent's system IP address:

```
admin> read ip-global
IP-GLOBAL read
admin> set system-ip-addr = 1.1.1.1
admin> write
IP-GLOBAL written
```

Configuring the Foreign Agent atmp profile

The following commands configure a minimal atmp profile:

```
admin> read atmp
ATMP read
admin> set agent-mode = foreign-agent
admin> set mtu-limit = 1472
admin> write
ATMP written
admin> reset
```

Note When you change the agent-mode parameter from its default tunnel-disabled setting to any other setting, you must reset the system for the new value to take effect.

Configuring a connection to the gateway Home Agent

In this example, the gateway Home Agent has the following system-ip-addr setting:

```
[in IP-GLOBAL]
system-ip-addr = 2.2.2.2
```

The following commands configure a connection profile to the gateway Home Agent:

```
admin> read conn hagateway
CONNECTION/hagateway read
admin> set active = yes
admin> set dial-number = 9-1-333-555-1212
admin> set ppp send-auth = chap-ppp-auth
admin> set ppp send-password = remotepw
admin> set ip-options remote = 2.2.2.2
admin> write
CONNECTION/hagateway written
```

Following are comparable RADIUS profiles:

```
route-taos-1 Password = "ascend", Service-Type = Dialout-Framed-User
   Framed-Route = "2.0.0.0 2.2.2.2 1 n hagateway-out"
hagateway-out Password = "ascend", Service-Type = Dialout-Framed-User
   User-Name = "hagateway",
   Framed-Protocol = MPP,
   Ascend-Route-IP = Route-IP-Yes,
   Framed-IP-Address = 2.2.2.2,
```

```
   Ascend-Dial-Number = "9-1-333-555-1212",
   Ascend-Send-Auth = Send-Auth-CHAP,
   Ascend-Send-Password = "remotepw"
```

Configuring a connection to the router Home Agent

In this example, the router Home Agent has the following system-ip-addr setting:

```
[in IP-GLOBAL]
system-ip-addr = 3.3.3.3
```

The following commands configure a connection profile for the connection to the router Home Agent:

```
admin> read connection harouter
CONNECTION/harouter read

admin> set active = yes

admin> set dial-number = 9-1-888-555-1234

admin> set ppp send-auth = chap-ppp-auth

admin> set ppp send-password = remotepw

admin> set ip-options remote = 3.3.3.3

admin> write
CONNECTION/harouter written
```

Following are comparable RADIUS profiles:

```
route-taos-1 Password = "ascend", Service-Type = Dialout-Framed-User
   Framed-Route = "3.0.0.0 3.3.3.3 1 n harouter-out"

harouter-out Password = "ascend", Service-Type = Dialout-Framed-User
   User-Name = "harouter",
   Framed-Protocol = MPP,
   Ascend-Route-IP = Route-IP-Yes,
   Framed-IP-Address = 3.3.3.3,
   Ascend-Dial-Number = "9-1-888-555-1234",
   Ascend-Send-Auth = Send-Auth-CHAP,
   Ascend-Send-Password = "remotepw"
```

Configuring a mobile-client connection to the gateway Home Agent

For the purposes of this example, the gateway Home Agent has a nailed profile named home-router for connection to the home network. It also has the following settings in its atmp profile:

```
[in ATMP]
agent-mode = home-agent
agent-type = gateway-home-agent
udp-port = 1555
home-agent-password = tunnel-password
```

The following set of commands, entered on the Foreign Agent, configures a mobile-client connection to the gateway Home Agent:

```
admin> read connection mobile-client-1
CONNECTION/mobile-client-1 read

admin> set active = yes
```

```
admin> set ppp recv-password = my-password
admin> set ip-options remote-address= 10.1.1.1/29
admin> set tunnel profile-type = mobile-client
admin> set tunnel primary-tunnel-server = 2.2.2.2:1555
admin> set tunnel password = tunnel-password
admin> set tunnel home-network-name = home-router
admin> write
CONNECTION/mobile-client-1 written
```

Following is a comparable RADIUS profile:

```
mobile-client-1 Password = "my-password"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
   Ascend-IP-Route = Route-IP-Yes,
   Framed-IP-Address = 10.1.1.1,
   Framed-IP-Netmask = 255.255.255.248,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "2.2.2.2:1555",
   Tunnel-Password = "tunnel-password",
   Tunnel-Private-Group-ID = "home-router"
```

Configuring a mobile-client connection to the router Home Agent

For the purposes of this example, the router Home Agent has the following settings in its atmp profile:

```
[in ATMP]
agent-mode = home-agent
agent-type = router-home-agent
udp-port = 8877
home-agent-password = tunnel-password
```

The next set of commands, entered on the Foreign Agent, configures a mobile-client connection to the router Home Agent:

```
admin> read connection mobile-client-2
CONNECTION/mobile-client-2 read
admin> set active = yes
admin> set ppp recv-password = my-password
admin> set ip-options remote-address= 11.1.1.1/32
admin> set tunnel profile-type = mobile-client
admin> set tunnel primary-tunnel-server = 3.3.3.3:8877
admin> set tunnel password = tunnel-password
admin> write
CONNECTION/mobile-client-2 written
```

Following is a comparable RADIUS profile:

```
mobile-client-2 Password = "my-password", Service-Type= Framed-User
   Framed-Protocol = MPP,
   Ascend-IP-Route = Route-IP-Yes,
```

```
Framed-IP-Address = 11.1.1.1,
Framed-IP-Netmask = 255.255.255.255,
Tunnel-Type = ATMP,
Tunnel-Server-Endpoint = "3.3.3.3:8877",
Tunnel-Password = "tunnel-password"
```

Example of a Foreign Agent that tunnels to a GRF switch

When a Stinger unit is operating as a Foreign Agent tunneling to a GRF® switch Home Agent, setting the `mtu-limit` parameter becomes a requirement rather than a recommendation. To maintain its very high throughput, the GRF® does not perform packet reassembly. If an `mtu-limit` value is not specified and a mobile client sends a large packet, the Foreign Agent might be forced to fragment the packet before sending it to the Home Agent. The GRF® switch Home Agent drops such packets.

Figure 6-6 shows a Foreign Agent tunneling to a GRF® Home Agent across a 100BaseT Ethernet segment.

Figure 6-6. Foreign Agent tunneling to a GRF switch

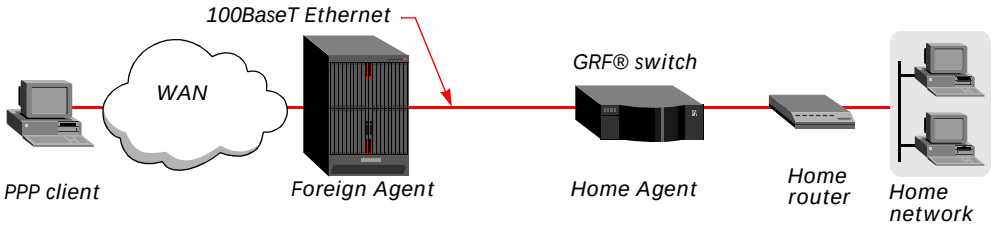

The following commands configure the Foreign Agent `atmp` profile for the Stinger unit in Figure 6-6:

```
admin> read atmp
ATMP read
admin> set agent-mode = foreign-agent
admin> set mtu-limit = 1472
admin> write
ATMP written
```

Note The GRF® switch ATMP configuration must specify the same `mtu-limit` value.

Configuring Home Agents

To configure an ATMP Home Agent, you must set parameters in the `atmp` profile, configure an IP connection to the Foreign Agent, and configure the connection to the home network.

For information about configuring a connection to the Foreign Agent, see “Configuring the agent-to-agent connection” on page 6-7.

Home Agent atmp profile settings

The atmp profile contains the following parameters (shown with sample values) related to a Home Agent:

```
[in ATMP]
agent-mode = home-agent
agent-type = gateway-home-agent
udp-port = 5150
home-agent-password = tunnel-password
retry-timeout = 3
retry-limit = 10
idle-timer = 30
mtu-limit = 0
force-fragmentation = no
```

Parameter	Setting for Home Agent configuration
agent-mode	Must specify home-agent.
agent-type	Specifies gateway-home-agent (the default) or router-home-agent, depending on how the Home Agent accesses the home network.
udp-port	Specifies the UDP port Foreign Agents must use to establish tunnels with the Home Agent, as described in "Setting the UDP port" on page 6-4.
home-agent-password	Specifies the password Foreign Agents must supply to establish a tunnel with this unit. You can specify up to 21 characters.
retry-timeout retry-limit	Together, these parameters specify how many tunnel RegisterRequest and DeregisterRequest messages are sent and the number of seconds between each message. The default settings are appropriate for most sites, as described in "Specifying tunnel retry limits" on page 6-4.
idle-timer	Specifies the number of minutes the Home Agent maintains an idle tunnel before disconnecting it.
mtu-limit	Specifies the maximum transmission unit (MTU) for the path between the foreign and Home Agents, as described in "Setting an MTU limit" on page 6-4.
force-fragmentation	Enable/disable prefragmentation of packets that have the DF bit set, as described in "Forcing fragmentation for interoperation with outdated clients" on page 6-6.

Specifying a gateway Home Agent

A gateway Home Agent delivers tunneled data to the home network without routing. A gateway Home Agent cannot ping or otherwise communicate with the home router. (The same restriction applies in the other direction.)

When the gateway Home Agent receives tunneled data, it removes the GRE header and forwards the packets to the home router, as shown in Figure 6-7.

Figure 6-7. How a gateway Home Agent works

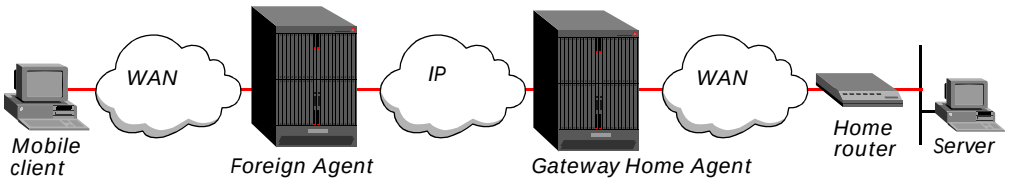

The link to the home network cannot be a regular switched dial-out connection, because the Home Agent will not dial the connection upon receipt of tunneled data. If the gateway connection is down when the Home Agent receives a tunnel request, it rejects the request. For more details about the gateway connection to the home network, see “Home network gateway profile settings” on page 6-20.

Following is an example of specifying a gateway Home Agent:

```
admin> read atmp
ATMP read
admin> set agent-mode = home-agent
admin> set agent-type = gateway-home-agent
admin> write
ATMP written
admin> reset
```

Note When you change the `agent-mode` parameter from its default `tunnel-disabled` setting to any other setting, you must reset the system for the new value to take effect.

Specifying a router Home Agent

A router Home Agent relies on packet routing to reach the home network, as shown in Figure 6-8.

Figure 6-8. How a router Home Agent works

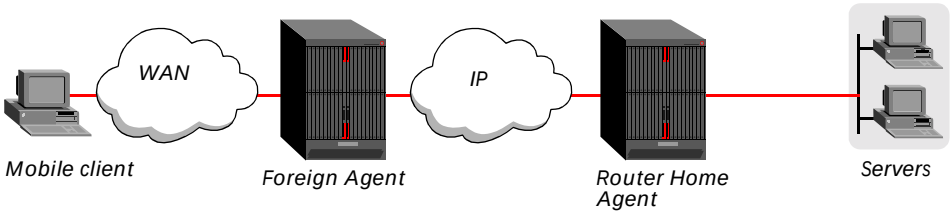

When the router Home Agent receives tunneled data, it removes the GRE encapsulation, passes the packets to its router software, and adds a route to the mobile client. If the mobile client is a PPP client, it adds a host route. If the mobile client is a router, such as a Pipeline® unit, the router Home Agent adds a regular route to the subnet addresses assigned to that router.

Following is an example of specifying a router Home Agent:

```
admin> read atmp
ATMP read
admin> set agent-mode = home-agent
```

```
admin> set agent-type = router-home-agent

admin> write
ATMP written

admin> reset
```

Note When you change the `agent-mode` parameter from its default `tunnel-disabled` setting to any other setting, you must reset the system for the new value to take effect.

Specifying the tunnel password

The Home Agent typically requests a password before establishing a tunnel. The Foreign Agent returns an encrypted version of the password found in the mobile-client profile. For details, see “Tunnel password authentication” on page 6-10.

Setting an idle timer for unused tunnels

When a mobile client disconnects normally, the Foreign Agent sends a request to the Home Agent to close down the tunnel. However, when a Foreign Agent restarts, tunnels that were established to a Home Agent are not normally cleared, because the Home Agent is not informed that the mobile clients are no longer connected. The unused tunnels continue to occupy memory on the Home Agent. To enable the Home Agent to reclaim the memory held by unused tunnels, you can now set an inactivity timer on a Home Agent by changing the default value of the following parameter:

```
[in ATMP]
idle-timer = 0
```

The inactivity timer runs only on the Home Agent side. Its value specifies the number of minutes (1 to 65535) that the Home Agent maintains an idle tunnel before disconnecting it. A value of 0 disables the timer, which means that idle tunnels remain connected forever. The setting affects only tunnels created after the timer was set. Tunnels that existed before the timer was set are not affected.

Home network gateway profile settings

When a gateway Home Agent receives a tunnel RegisterRequest message from the Foreign Agent, it checks the status of the connection to the home network. If the connection is down, the Home Agent rejects the tunnel request and does not attempt to dial the connection. If the connection terminates after a tunnel is established, all mobile clients that were using it are disconnected.

The gateway connection to the home network can be a nailed connection or a regular dial-in switched connection. Using an incoming connection from the home router enables the administrator of the home network to regulate when mobile clients can access the network. For example, the administrator of the home network could configure an access router to dial the Home Agent every weekday at 8:00 a.m. and disconnect at 5:00 p.m., thereby limiting mobile client access to those hours. In that case, the gateway connection must be operational before mobile clients dial in, or their tunnel requests will fail.

To configure a gateway profile, set up a nailed or dial-in connection and specify the following parameters (shown with sample settings) in the `connection` profile:

```
[in CONNECTION/gwprofile]
station* = gwprofile

[in CONNECTION/gwprofile:tunnel-options]
profile-type = gateway-profile
max-tunnels = 0
atmp-ha-rip = rip-send-v2
```

Parameter	Setting for gateway profile configuration
station	Name of the home router. The home-network-name value specified in the mobile-client profile on the Foreign Agent must specify the same name.
profile-type	Must specify gateway-profile.
max-tunnels	Maximum number of mobile clients that can use the connection, all at the same time, to tunnel into the home network. The default value of 0 sets no limit.
atmp-ha-rip	Enable/disable construction of mobile-client routes in RIP-2 responses to the home router. This parameter does not apply unless profile-type is set to gateway-profile. The parameter operates independently of the RIP parameter in the ip-options subprofile. For gateway profiles, the ip-options RIP parameter must be Off.

Limiting the maximum number of tunnels

If you decide to limit the maximum number of tunnels a gateway supports, be sure to consider the expected traffic per mobile-client connection, the bandwidth of the connection to the home network, and the availability of alternative Home Agents (if any). For example, the lower the amount of traffic generated by each mobile-client connection, the more tunnels a gateway connection will be able to handle.

Enabling RIP on the interface to the home router

The atmp-ha-rip setting enables the gateway Home Agent to inform the home router about routes to its mobile clients. Enabling RIP eliminates the requirement for the home router to maintain a static route for each ATMP mobile client. It also provides the basis for a resilient configuration, in which a secondary Home Agent can take over for a primary Home Agent if the primary agent becomes unavailable.

Informing the home router about routes to mobile clients

The router at the far end of the connection defined by the gateway profile must be able to route back to mobile clients. The easiest way to establish this capability is by setting the atmp-ha-rip parameter to rip-send-v2. With this setting, the gateway Home Agent constructs a RIP-2 Response(2) packet at every RIP interval and sends it to the home network from all tunnels using the gateway profile. For each tunnel, the Response packet contains the mobile-client IP address, the subnet mask, a next hop of 0.0.0.0, and a metric of 1. RIP-2 authentication and route tags are not supported.

The following commands enable `atmp-ha-rip` in the gateway profile for the connection to the home router:

```
admin> new connection home-router
CONNECTION/home-router read
admin> set tunnel profile-type = gateway-profile
admin> set tunnel atmp-ha-rip = rip-send-v2
admin> write
CONNECTION/home-router written
```

Note The Home Agent will not inspect RIP updates coming from the home network, regardless of the RIP setting in the `ip-options` subprofile. If the Home Agent receives RIP updates from the home network, it forwards the update packets to the mobile clients, as it would any other type of packet.

The alternative: Maintaining static routes in the home router

If the gateway profile does *not* set `atmp-ha-rip` to `rip-send-v2`, the administrator of the home network must configure a static route to each mobile client. A static route to a mobile client can be specific to the client, in which case the route's destination is the mobile-client IP address and the next-hop router is the Home Agent address. For example, in the following route the mobile client is a router (this is not a host route), and the Home Agent address is 2.2.2.2:

```
[in IP-ROUTE/mobile-client]
destination = 10.1.1.10/29
gateway = 2.2.2.2
```

Or, if the mobile clients have addresses allocated from the same address block (including router mobile-client addresses with subnet masks of less than 32 bits) and no addresses from that block are assigned to other hosts, the home network administrator can specify a single static route that encompasses all mobile clients that use the same Home Agent. For example, in the following route all mobile clients are allocated addresses from the 10.4.*n.n* block (and no other hosts are allocated addresses from that block), and the Home Agent address is 2.2.2.2:

```
[in IP-ROUTE/mobile-clients]
destination = 10.4.0.0/16
gateway = 2.2.2.2
```

Routing in a resilient installation

A resilient ATMP installation supports multiple ATMP paths to the same home network, providing resiliency in the event of Home Agent failure or failure of the link between a Home Agent and home router. The two Home Agents might connect to two home routers, as shown in Figure 6-9, or the Home Agents might connect to the same home router.

Figure 6-9. Resilient ATMP installation

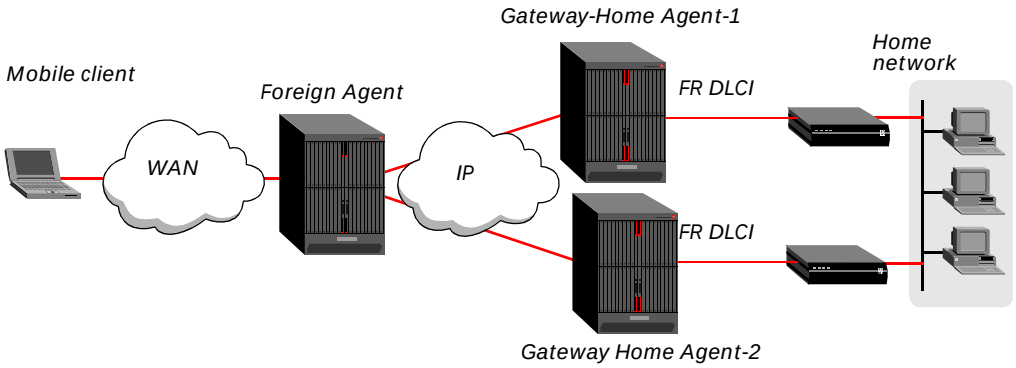

Mobile clients access the home network through one of the Home Agents, but not always the same Home Agent. Therefore, a static route maintained by the home router would not allow hosts on the home network to reliably send packets back to mobile clients. Setting the `atmp-ha-rip` parameter resolves the routing problems that can occur in a resilient configuration.

The following example shows a gateway profile that can reside in both Home Agents shown in Figure 6-9:

```
admin> new connection home-router
CONNECTION/home-router read
admin> set active = yes
admin> set tunnel profile-type = gateway-profile
admin> set tunnel max-tunnels = 120
admin> set tunnel atmp-ha-rip = rip-send-v2
admin> write
CONNECTION/home-router written
```

Typical gateway Home Agent configuration

Figure 6-10 shows a gateway Home Agent with a leased line to the home network.

Figure 6-10. Gateway Home Agent with a leased line to a home network

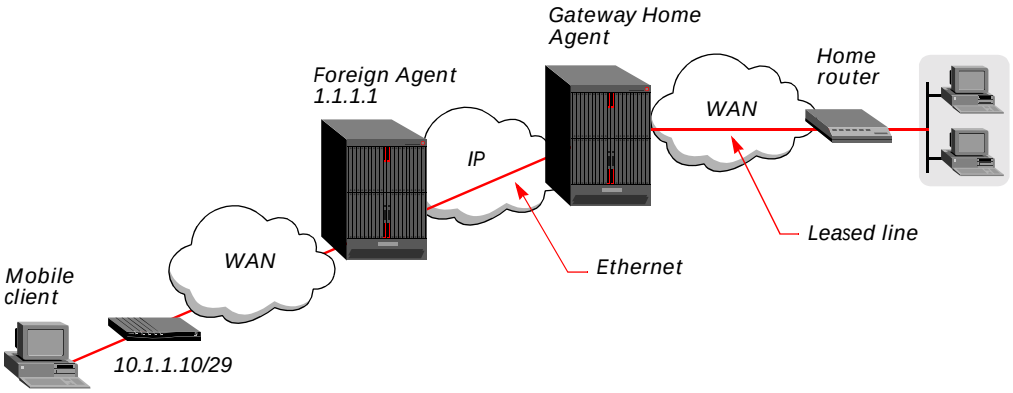

Note In the preceding example, the ATMP Foreign Agent and Home Agent are on the same Ethernet segment, so no `connection` profiles are required for communication.

Setting the Home Agent's system address

The following commands set the Home Agent's system IP address:

```
admin> read ip-global
IP-GLOBAL read
admin> set system-ip-addr = 2.2.2.2
admin> write
IP-GLOBAL written
```

Configuring the Home Agent ATMP profile

The following commands configure the Home Agent `atmp` profile, with the default setting of `gateway-home-agent` for the `agent-type` parameter:

```
admin> read atmp
ATMP read
admin> set agent-mode = home-agent
admin> set udp-port = 1234
admin> set password = tunnel-password
admin> set idle-timer = 30
admin> set mtu-limit = 1472
admin> write
ATMP written
admin> reset
```

Note When you change the `agent-mode` parameter from its default `tunnel-disabled` setting to any other setting, you must reset the system for the new value to take effect.

The Foreign Agent has an `atmp` profile such as the following:

```
[in ATMP]
agent-mode = foreign-agent
agent-type = gateway-home-agent
udp-port = 5150
home-agent-password = ""
retry-timeout = 3
retry-limit = 10
idle-timer = 0
mtu-limit = 1472
force-fragmentation = no
```

Configuring a gateway profile for connection to the home network

In the following set of commands, which configure the interface to the home network, `call-type` is set to FT1 (nailed) and a group of nailed channels (group number 7) is assigned to the link. The `atmp-ha-rip` parameter is enabled on the interface.

```
admin> new connection home-router
CONNECTION/home-router read
admin> set active = yes
admin> set tunnel profile-type = gateway-profile
admin> set tunnel atmp-ha-rip = rip-send-v2
admin> set telco call-type = ft1
admin> set telco nailed-groups = 7
admin> write
CONNECTION/home-router written
```

Configuring a mobile-client connection to the gateway Home Agent

Mobile-client connections on the Foreign Agent will require a tunnel configuration such as the following in a local connection profile:

```
[in CONNECTION/mclient:tunnel-options]
profile-type = mobile-client
primary-tunnel-server = 2.2.2.2:1234
password = tunnel-password
home-network-name = home-router
```

Or you can enable comparable settings in a RADIUS profile:

```
mclient Password = "local-password"
   Service-Type = Framed-User,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "2.2.2.2:1234",
   Tunnel-Password = "tunnel-password",
   Tunnel-Private-Group-ID = "home-router"
```

Typical router Home Agent configuration

Figure 6-11 shows a router Home Agent with an Ethernet connection to the home network. The ATMP Foreign Agent and Home Agent connect across a multichannel PPP link.

Figure 6-11. Router Home Agent on the home network

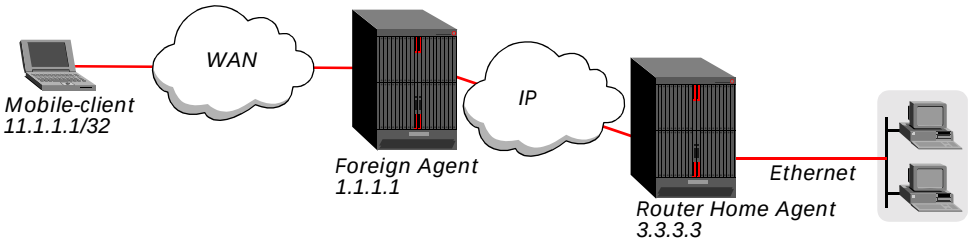

For information about configuring a connection to the Foreign Agent, see “Configuring the agent-to-agent connection” on page 6-7.

Setting the Home Agent's system address

The following commands set the router Home Agent's system IP address:

```
admin> read ip-global
IP-GLOBAL read
admin> set system-ip-addr = 3.3.3.3
admin> write
IP-GLOBAL written
```

Configuring the **ip-interface** profile for the connection to the home network

If you enable RIP on the interface that leads to the home network, other hosts and networks can route to the mobile client. Enabling RIP is particularly useful if the home network is one or more hops away. If RIP is turned off, intervening routers require static routes that specify the Home Agent as the route to mobile clients. You can also turn on proxy ARP to allow local hosts to ARP for mobile clients. For example:

```
admin> read ip-interface {{1 10 1}0}
IP-INTERFACE/{ { 1 10 1 } 0 } read
admin> set ip-address = 3.3.3.3
admin> set proxy-mode = always
admin> set rip-mode = routing-send-and-recv-v2
admin> write
IP-INTERFACE/{ { 1 10 1 } 0 }written
```

Configuring the Home Agent's ATMP profile

The following commands configure the Home Agent's `atmp` profile:

```
admin> read atmp
ATMP read
admin> set agent-mode = home-agent
admin> set agent-type = router
admin> set password = tunnel-password
admin> set idle-timer = 30
admin> set mtu-limit = 1472
admin> write
ATMP written
admin> reset
```

Note When you change the `agent-mode` parameter from its default `tunnel-disabled` setting to any other setting, you must reset the system for the new value to take effect.

The Foreign Agent has an `atmp` profile such as the following:

```
[in ATMP]
agent-mode = foreign-agent
agent-type = gateway-home-agent
udp-port = 5150
home-agent-password = ""
```

```
retry-timeout = 3
retry-limit = 10
idle-timer = 0
mtu-limit = 1472
force-fragmentation = no
```

Configuring a mobile-client connection to the router Home Agent

Mobile-client connections on the Foreign Agent require a tunnel configuration such as the following in a local `connection` profile:

```
[in CONNECTION/mclient:tunnel-options]
profile-type = mobile-client
primary-tunnel-server = 3.3.3.3
password = tunnel-password
```

Or this type of connection requires comparable tunnel settings in a RADIUS profile:

```
mclient Password = "local-password"
   Service-Type = Framed-User,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "3.3.3.3",
   Tunnel-Password = "tunnel-password"
```

Configuring a Home-and-Foreign Agent

In some configurations, a Stinger unit acts as a Home Agent for some mobile clients and as a Foreign Agent for others. The two configurations operate side-by-side without any conflict, provided that all requirements are met for each type of configuration.

Configuring the atmp profile

The `atmp` profile contains the following parameters (shown with sample values) related to the Home-and-Foreign Agent configuration:

```
[in ATMP]
agent-mode = home-and-foreign-agent
agent-type = gateway-home-agent
udp-port = 5150
home-agent-password = tunnel-password
retry-timeout = 3
retry-limit = 10
idle-timer = 0
mtu-limit = 1472
force-fragmentation = no
```

The `agent-mode` parameter must specify `home-and-foreign-agent`. For details about all of the other settings, see “Configuring Home Agents” on page 6-17 or “Configuring a Foreign Agent” on page 6-8.

Typical Home-and-Foreign Agent configuration

Figure 6-12 shows a Stinger unit operating as Home Agent for home network B and as Foreign Agent for mobile clients tunneling into home network A.

Figure 6-12. Stinger unit acting as both Home Agent and Foreign Agent

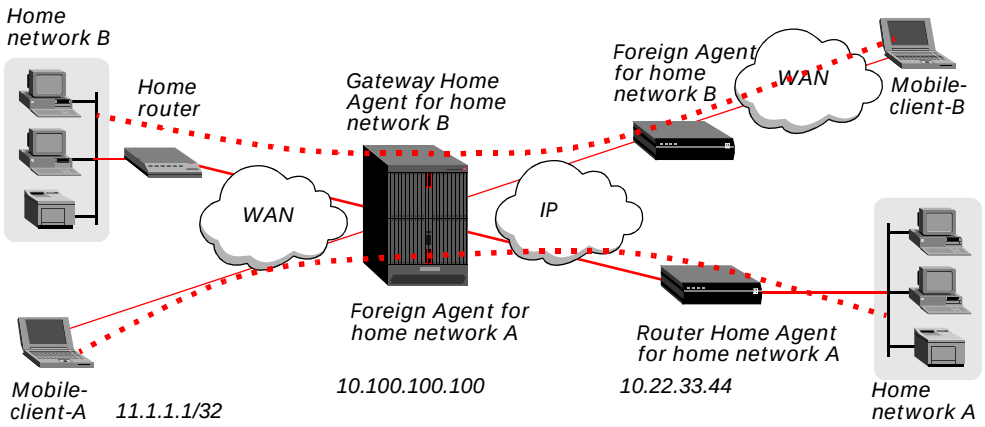

For information about configuring connections between Home Agents and Foreign Agents, see “Configuring the agent-to-agent connection” on page 6-7.

Setting the system address

The following commands set the Home-and-Foreign Agent’s system IP address:

```
admin> read ip-global
IP-GLOBAL read
admin> set system-ip-addr = 10.100.100.100
admin> write
IP-GLOBAL written
```

Configuring the atmp profile for a Home-and-Foreign Agent

The following set of commands configures the `atmp` profile:

```
admin> read atmp
ATMP read
admin> set agent-mode = home-and-foreign-agent
admin> set agent-type = gateway-home-agent
admin> set password = tunnel-password
admin> set udp-port = 1567
admin> set idle-timer = 30
admin> set mtu-limit = 1472
admin> write
ATMP written
admin> reset
```

Note When you change the `agent-mode` parameter from its default `tunnel-disabled` setting to any other setting, you must reset the system for the new value to take effect.

The Foreign Agent for home network B has an `atmp` profile such as the following:

```
[in ATMP]
agent-mode = foreign-agent
```

```
agent-type = gateway-home-agent
udp-port = 5150
home-agent-password = ""
retry-timeout = 3
retry-limit = 10
idle-timer = 0
mtu-limit = 1472
force-fragmentation = no
```

The Home Agent for home network A has an `atmp` profile such as the following:

```
[in ATMP]
agent-mode = home-agent
agent-type = router-home-agent
udp-port = 8877
home-agent-password = tunnel-password
retry-timeout = 3
retry-limit = 10
idle-timer = 0
mtu-limit = 1472
force-fragmentation = no
```

Configuring a mobile-client profile

The following set of commands configures a `connection` profile for mobile-client-A in Figure 6-12. For this profile, the Stinger unit is operating as Foreign Agent to enable the mobile client to tunnel to home network A:

```
admin> read connection mobile-client-A
CONNECTION/mobile-client-A read
admin> set active = yes
admin> set ip-options remote-address = 11.1.1.1/32
admin> set tunnel profile-type = mobile-client
admin> set tunnel primary-tunnel-server = 10.22.33.44:8877
admin> set tunnel password = tunnel-password
admin> write
CONNECTION/mobile-client-A written
```

Following is a comparable RADIUS profile:

```
mobile-client-A Password = "local-password"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
   Ascend-IP-Route = Route-IP-Yes,
   Framed-IP-Address = 11.1.1.1,
   Framed-IP-Netmask = 255.255.255.255,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "10.22.33.44",
   Ascend-UDP-Port = 8877,
   Tunnel-Password = "tunnel-password"
```

Another example of a Home-and-Foreign Agent configuration

Figure 6-13 shows another configuration that makes use of the Home-and-Foreign Agent setup. In this example, all three mobile clients want to tunnel to the home network, using TAOS2 as their Home Agent. The two ATMP units are geographically distant.

Figure 6-13. Enabling a mobile client to bypass the Foreign Agent connection

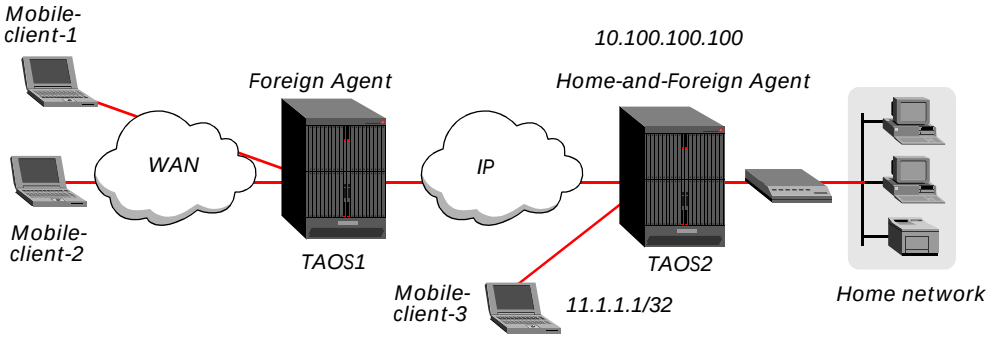

Mobile-client-1 and mobile-client-2 make local calls to dial in to the Foreign Agent (TAOS1) and then tunnel to the Home Agent. However, mobile-client-3 is geographically closer to TAOS2, and would prefer to dial directly in to TAOS2. In this case, TAOS2 is configured to provide both Home Agent and Foreign Agent functionality to mobile-client-3. There is no need to encapsulate data to and from mobile-client-3 in GRE. The data comes in on one of TAOS2's interfaces and it is sent to another interface without encapsulation processing, but with all of the network isolation benefits that ATMP provides.

Setting the system IP address

The following commands set the Home-and-Foreign Agent's system IP address:

```
admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 10.100.100.100

admin> write
IP-GLOBAL written
```

Configuring the ATMP profile for Home-and-Foreign Agent

The following commands configure the `atmp` profile in TAOS2:

```
admin> read atmp
ATMP read

admin> set agent-mode = home-and-foreign-agent

admin> set agent-type = gateway-home-agent

admin> set password = tunnel-password

admin> set udp-port = 6789

admin> set idle-timer = 30

admin> set mtu-limit = 1472
```

```
admin> write
ATMP written

admin> reset
```

Note When you change the agent-mode parameter from its default tunnel-disabled setting to any other setting, you must reset the system for the new value to take effect.

TAOS1 has an atmp profile such as the following:

```
[in ATMP]
agent-mode = foreign-agent
agent-type = gateway-home-agent
udp-port = 5150
home-agent-password = ""
retry-timeout = 3
retry-limit = 10
idle-timer = 0
mtu-limit = 1472
force-fragmentation = no
```

Configuring a profile for mobile-client-3

The next set of commands configures a connection profile for mobile-client-3 in Figure 6-13. For this profile, the Stinger unit is operating as both Foreign Agent and Home Agent.

```
admin> read connection mobile-client-3
CONNECTION/mobile-client-3 read

admin> set active = yes

admin> set ip-options remote-address = 11.1.1.1/32

admin> set tunnel profile-type = mobile-client

admin> set tunnel primary-home-agent = 10.100.100.100:6789

admin> set tunnel password = tunnel-password

admin> write
CONNECTION/mobile-client-3 written
```

Following is a comparable RADIUS profile:

```
mobile-client-3 Password = "local-password"
   Service-Type = Framed-User,
   Framed-Protocol = MPP,
   Ascend-IP-Route = Route-IP-Yes,
   Framed-IP-Address = 11.1.1.1,
   Framed-IP-Netmask = 255.255.255.255,
   Tunnel-Type = ATMP,
   Tunnel-Server-Endpoint = "10.100.100.100:6789",
   Tunnel-Password = "tunnel-password"
```

Packet Filters

7

Filter overview

A filter consists of specifications describing packets and actions to take upon packets that match the descriptions. After you apply a filter to an interface, the system monitors the data stream on that interface.

Depending on how you define a filter, it can apply to inbound packets, outbound packets, or both. In addition, filters are flexible enough to specify taking an action (such as forward or drop) on those packets that match the specifications, or on all packets *except* those that match the specifications.

Basic types of filters

Each `filter` profile contains up to 12 input filters (applied to inbound packets) and 12 output filters (applied to outbound packets). Each of the up to 24 specifications can be one of the following basic types of filters:

- Generic filters
- IP filters
- Type-of-service (TOS) filters
- Route filters (local `filter` profiles only)

Generic filters examine the byte-level or bit-level contents of any packet, comparing specified bytes or bits to a value defined in the filter. On the basis of this comparison, the filter specifies a forwarding action. To use generic filters effectively, you need to know the contents of certain bytes in the packets you wish to filter. Protocol specifications are usually the best source of such information.

IP filters apply only to IP-related packets. They specify a forwarding action on the basis of higher-level fields in IP packets (for example, the source or destination address, or the protocol number). IP filters operate on logical information, which is relatively easy to obtain.

Type-of-service (TOS) filters set priority bits in the TOS header of IP packets. Other routers can then use the information to prioritize and select links for particular data streams.

Route filters apply only to RIP update packets. They specify whether matching routes in a RIP packet will be accepted into the routing table, denied, or accepted with an increased metric. Route filters can also specify a source address, which means that they can take an action on all updates from that address.

Data and call filters

Data filters are commonly used for security, but they can apply to any purpose that requires the system to drop or forward specific packets. The focus is typically on keeping out traffic that you do not want on a LAN. For example, you can use data filters to drop packets addressed to particular hosts or to prevent broadcasts from going across the WAN. You can also use data filters to allow users to access only specific devices across the WAN.

When you apply a data filter (Figure 7-1), its forwarding action (forward or drop) affects the actual data stream by preventing certain packets from reaching the Ethernet from the WAN, or vice versa. Data filters do not affect the idle timer, and a data filter applied to a `connection` profile does not affect the answering process.

Figure 7-1. Data filters drop or forward certain packets.

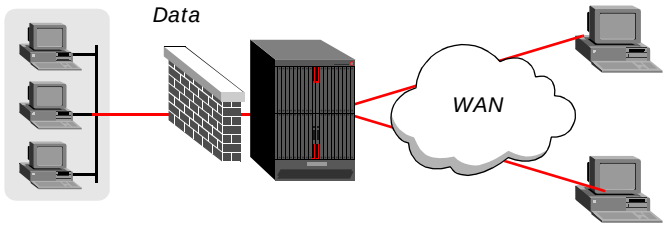

Call filters (Figure 7-2) prevent unnecessary connections and help the system distinguish active traffic from “noise.” By default, any traffic to a remote site triggers a call, and any traffic across an active connection resets the connection’s idle timer.

Figure 7-2. Call filters prevent certain packets from resetting the timer.

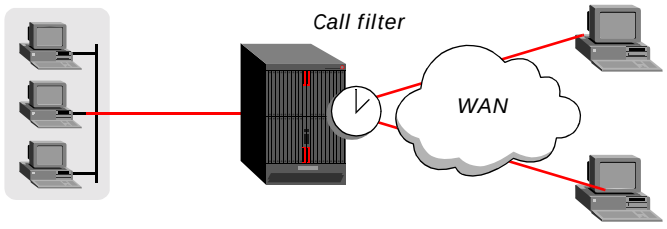

When you apply a call filter, its forwarding action (forward or drop) does *not* affect which packets are sent across an active connection. The forwarding action of a call filter determines which packets can either initiate a connection or reset a session’s

timer. When a session's idle timer expires, the session is terminated. With the default `idle-timer` setting of 120 seconds, the system terminates a connection that has been inactive for 2 minutes.

How filters work

When no filter is in use, the system forwards all packets. But applying a filter to an interface reverses this default.

A `filter` profile can include up to 12 input-filter and 12 output-filter specifications (filters). Each filter has its own forwarding action—forward or drop. The filters are applied in sequence. At the first successful comparison between a filter and the packet being examined, the filtering process stops and the forwarding action in that filter is applied to the packet. For route filters, the forwarding action has no effect, but another type of action in the filter is applied to the packet when a comparison succeeds.

If all comparisons fail, the packet does not match the filter. However, the packet is forwarded only if the filter explicitly allows such packets to pass. For security purposes, the unit does not automatically forward nonmatching packets. (For a sample input filter that forwards packets that did not match a previous filter, see "Examples of an IP filter to prevent local address spoofing" on page 7-15.)

Note For a call filter to prevent an interface from remaining active unnecessarily, you must define filters for both input and output packets. Otherwise, if only input filters are defined, output packets will keep a connection active, or vice versa.

Generic filters

In a generic filter, all the settings in a filter specification work together to specify a location in a packet and a number to be compared to that location. The type of comparison that constitutes a match (equal or not-equal) must also be specified. When a comparison fails, the packet undergoes the next comparison. When a comparison succeeds, the filtering process stops and the forwarding action in that filter is applied to the packet.

If a generic filter is applied as a call filter and a comparison succeeds, the forwarding action is either to reset the idle timer or not, depending on how the filter is defined. If a generic filter is applied as a data filter, the forwarding action is either to forward the packet or drop it.

IP filters

In an IP filter, each filter specification includes a set of comparisons that are made in a defined order. When a comparison fails, the packet undergoes the next comparison. When a comparison succeeds, the filtering process stops and the forwarding action in that filter is applied to the packet. The IP filter tests proceed in the following order:

1 Apply the `source-address-mask` value to the `source-address` value and compare the result to the source address of the packet. If they are not equal, the comparison fails.

2 Apply the `dest-address-mask` value to the `dest-address` value and compare the result to the destination address in the packet. If they are not equal, the comparison fails.

3 If the protocol parameter is zero (which matches any protocol), the comparison succeeds. If it is nonzero and not equal to the Protocol field in the packet, the comparison fails.

4 If the src-port-cmp parameter is not set to none, compare the source-port number to the source port number of the packet. If they do not match as specified by the src-port-cmp parameter, the comparison fails.

5 If the dst-port-cmp parameter is not set to none, compare the dest-port number to the destination port number of the packet. If they do not match as specified by the dst-port-cmp parameter, the comparison fails.

If an IP filter is applied as a call filter and a comparison succeeds, the forwarding action is either to reset the idle timer or not, depending on how the filter is defined. If an IP filter is applied as a data filter, the forwarding action is either to forward the packet or drop it.

Type-of-service (TOS) filters

In an IP TOS filter, each filter specification includes a set of comparisons that are made in a defined order. When a comparison fails, the packet undergoes the next comparison. When a comparison succeeds, the filtering process stops and the action specified in that filter is applied to the packet. The TOS filter tests proceed in the following order:

1 Apply the source-address-mask value to the source-address value and compare the result to the source address of the packet. If they are not equal, the comparison fails.

2 Apply the dest-address-mask value to the dest-address value and compare the result to the destination address in the packet. If they are not equal, the comparison fails.

3 If the protocol parameter is zero (which matches any protocol), the comparison succeeds. If it is nonzero and not equal to the Protocol field in the packet, the comparison fails.

4 If the src-port-cmp parameter is not set to none, compare the source-port number to the source port number of the packet. If they do not match as specified by the src-port-cmp parameter, the comparison fails.

5 If the dst-port-cmp parameter is not set to none, compare the dest-port number to the destination port number of the packet. If they do not match as specified by the dst-port-cmp parameter, the comparison fails.

If a comparison succeeds, the system sets the precedence bits and class of service (depending on how the filter is defined) in the TOS header of the packet.

Route filters

In a route filter, each filter specification includes a set of comparisons that are made in a defined order. When a comparison fails, the RIP packet undergoes the next comparison. When a comparison succeeds, the filtering process stops and the action specified in that filter is applied to the matching route or packet. The route filter tests proceed in the following order:

1 Apply the source-address-mask value to the source-address value and compare the result to the source address of the packet. If they are not equal, the comparison fails.

2 Apply the route-mask value to the route-address value and compare the result to the routes in the packet. If there is no match, the comparison fails.

If a comparison succeeds, the system performs one of the following actions, depending on how the filter is defined:

- If action is set to add, increase the metric field of the matching routes by the add-metric value and then add them to the routing table.
- If action is set to accept, add the matching routes to the routing table.
- If action is set to deny, reject the matching routes (do not add them to the routing table).

Specifying a filter's name and direction

A filter profile can define up to 12 input-filter specifications and 12 output-filter specifications. You apply the filter to an interface by referring to the profile by name. Only filters with valid specifications apply. Following are the relevant parameters (shown with their default settings):

```
[in FILTER/""]
filter-name* = ""
input-filters = [ { no no generic-filter { 0 0 no no 00:00:00:00:00:0+
output-filters = [ { no no generic-filter { 0 0 no no 00:00:00:00:00:+

[in FILTER/"":input-filters[n]]
valid-entry = no

[in FILTER/"":output-filters[n]]
valid-entry = no
```

Parameter	Setting
filter-name	Name of the filter profile, up to 36 characters. You apply a filter to an interface by referring to this name.
input-filters	Each filter can contain an array of up to 12 input-filter specifications, which are defined individually and applied in order (1–12) to the inbound packet stream. The order in which the input filters are defined is significant.
output-filters	Each filter can contain an array of up to 12 output-filter specifications, which are defined individually and applied in order (1–12) to the outbound packet stream. The order in which the output filters are defined is significant.
valid-entry	Enable/disable the filter specification. With a setting of no (the default), the system skips the specification when filtering the data stream. Set this parameter to yes for each defined filter you intend to use.

In a RADIUS profile, each filter is specified separately through the use of the Ascend-Data-Filter and Ascend-Call-Filter attributes. As is always the case with filters, the order in which they are applied within the user profile is significant.

In a RADIUS filter definition, you specify the direction in which to monitor the data stream as `in` or `out`. This setting provides the same function as the `input-filters` and `output-filters` parameters in a local profile. The following example shows an input-filter definition in RADIUS:

```
test-user Password = "test-pw"
    Ascend-Data-Filter = "ip in forward tcp dstport > 1023"
```

Specifying a filter's forwarding action

For generic or IP filters, each input or output filter in a local `filter` profile specifies a forwarding action for packets that match the filter. Following is the relevant parameter (shown with its default settings):

```
[in FILTER/"":input-filters[n]]
forward = no
```

```
[in FILTER/"":output-filters[n]]
forward = no
```

Parameter	Setting
`forward`	Forwarding action for the filter. When no filters are in use, the system forwards all packets by default. When a filter is in use, the default is to discard matching packets (`forward = no`).

Note For route filters and type-of-service (TOS) filters, the forwarding action has no effect. Those filters perform a different type of action on matching packets. See "Type-of-service (TOS) filters" on page 7-4.

In a RADIUS definition, you specify the action a filter takes as `forward` or `drop`. This setting provides the same function as the `forward` parameter in a local profile. The following example shows an input filter whose forwarding action is to drop matching packets:

```
test-user Password = "test-pw"
    Ascend-Data-Filter = "ip in drop tcp dstport > 1023"
```

Defining generic filters

Generic filters can match any packet, regardless of its protocol type or header fields. The filter specifications operate together to define a location in a packet and a hexadecimal value to compare to it.

Generic filter settings in a local filter profile

In a local `filter` profile, a generic filter uses the following parameters. The parameters are shown with their default values for input filters. The same values apply for output filter specifications.

```
[in FILTER/"":input-filters[n]]
type = generic-filter
```

```
[in FILTER/"":input-filters[n]:gen-filter]
offset = 0
len = 0
```

```
more = no
comp-neq = no
mask = 00:00:00:00:00:00:00:00:00:00:00:00
value = 00:00:00:00:00:00:00:00:00:00:00:00
```

If you set the parameters in an input filter, only inbound packets are examined. If you set them in an output filter, only outbound packets are examined.

Parameter	Setting
`type`	Type of filter. Leave the default value `generic-filter` for a generic filter. Only the parameters in the corresponding subprofile are applicable.
`offset`	Byte-offset at which to start comparing packet contents to the `value` setting specified in the filter. For details, see "Specifying the offset to the bytes to be examined" on page 7-8.
`len`	Number of bytes to test in a packet, starting with the byte specified by the `offset` parameter. For details, see "Specifying the number of bytes to test" on page 7-9.
`more`	Enable/disable application of the next filter before determining whether the packet matches the specification. If `more` is set to `yes`, the current specification is linked to the one immediately following. The match occurs only if *both* specifications are matched. (The subsequent specification must be enabled, or the system ignores the filter specification in which `more` is set to `yes`.) The `more` parameter enables you to create a filter that examines multiple noncontiguous bytes within a packet before the forwarding decision is made.
`comp-neq`	Type of comparison to perform. If `comp-neq` (Compare-Not-Equals) is set to `yes`, the comparison succeeds (the filter matches) if the contents do not equal the specified value. For a filter that requires the packet contents to equal the specified value, leave `comp-neq` set to `no`.
`mask`	Binary mask. The system applies the mask to the value specified by the `value` parameter before comparing it to the bytes in a packet specified by the `offset` parameter. For details, see "Masking the value before comparison" on page 7-9.
`value`	Hexadecimal number to be compared to the packet data identified by the `offset`, `len`, and `mask` calculations. After you have entered the number, the system enters a colon at the byte boundaries.

Generic filter settings in a RADIUS profile

In RADIUS, a generic filter entry is a value of the Ascend-Call-Filter or Ascend-Data-Filter attribute. To specify a generic filter value, use the following format:

"generic *dir action offset mask value compare* [more]"

Keyword or argument	Value
generic	Type of filter. Valid types specified by the Ascend-Data-Filter and Ascend-Call-Filter attributes are generic (the default) and ip.
dir	Specifies direction of the packets. You can specify in (to filter packets coming in to the Stinger unit or out (to filter packets going out of the Stinger unit).
action	Specifies the action that the system takes with a packet that matches the filter. Specify either forward or drop.
offset	Byte-offset in a packet at which to start comparing packet contents to the value specified in the filter. For details, see "Specifying the offset to the bytes to be examined" on page 7-8.
mask	Binary mask. The system applies the mask to the specified value before comparing it to the bytes specified by offset. For details, see "Masking the value before comparison" on page 7-9.
value	A hexadecimal number to compare to the packet contents at the specified offset. The length of the number must be the same as the length of the mask (up to 12 bytes).
compare	A comparison operator that determines how the system compares packet contents to the filter value. You can specify = (equal to) or != (not equal to). The default is = (equal to).
more	If the more flag is present, the system applies the next filter definition in the profile to the current packet before deciding whether to forward or drop the packet. The direction and forwarding action of the next filter must be the same as the current filter, or the system ignores this flag.

Specifying the offset to the bytes to be examined

The offset in a generic filter is a byte offset from the start of a packet to the start of the data in the packet to be tested. For example, assume a filter with the following filter specification:

```
[in FILTER/"":input-filters[n]:gen-filter]
offset = 2
len = 8
more = no
```

```
comp-neq = no
mask = 0f:ff:ff:ff:00:00:00:f0:00:00:00:00
value = 07:fe:45:70:00:00:00:90:00:00:00:00
```

Or assume a filter with comparable RADIUS filter definition:

```
Ascend-Data-Filter = "generic in drop 2 0fffffff000000f 07fe45700000009"
```

Then assume a packet with the following contents:

```
2A 31 97 FE 45 70 12 22 33 99 B4 80 75
```

When the filter is applied, the first 2 byes in the packet (2A and 31) are ignored because of the 2-byte offset.

Specifying the number of bytes to test

In a RADIUS profile, the length of the mask must equal the length of the *value* setting. The system tests that number of bytes in the packet, starting at the specified offset. In a local `filter` profile, the `len` setting specifies the number of bytes to test in a packet, starting with the byte specified by the `offset` parameter. The `mask` setting is assumed to have the same number of octets as the data specified by the `len` parameter.

For example, assume a filter with the following filter specification:

```
[in FILTER/"":input-filters[n]:gen-filter]
offset = 2
len = 8
more = no
comp-neq = no
mask = 0f:ff:ff:ff:00:00:00:f0:00:00:00:00
value = 07:fe:45:70:00:00:00:90:00:00:00:00
```

Then assume a packet with the following contents:

```
2A 31 97 FE 45 70 12 22 33 99 B4 80 75
```

The filter tests the value of bytes three (97) through ten (99).

Masking the value before comparison

A generic filter can include a mask to apply to the value specified by the `value` parameter before the system compares it to the bytes starting at the specified offset. You can use the mask to specify exactly which bits you want to compare. The mask is assumed to have the same number of octets as the data specified by the `len` parameter.

The system translates both the mask and the value specified by the `value` parameter into binary format and then applies a logical AND to the results. Each binary 0 (zero) in the mask hides the bit in the corresponding position in the value. A mask of all ones (FF:FF:FF:FF:FF:FF:FF:FF) masks no bits, so the full specified value must match the packet contents. For example, assume a filter with the following specification:

```
[in FILTER/"":input-filters[n]:gen-filter]
offset = 2
len = 8
more = no
comp-neq = no
```

```
mask = 0f:ff:ff:ff:00:00:00:f0:00:00:00:00
value = 07:fe:45:70:00:00:00:90:00:00:00:00
```

Or assume a filter with comparable RADIUS definition:

```
Ascend-Data-Filter = "generic in drop 2 0fffffff000000f 07fe45700000009"
```

Then assume a packet with the following contents:

```
2A 31 97 FE 45 70 12 22 33 99 B4 80 75
```

The value setting matches the packet data after application of the mask.

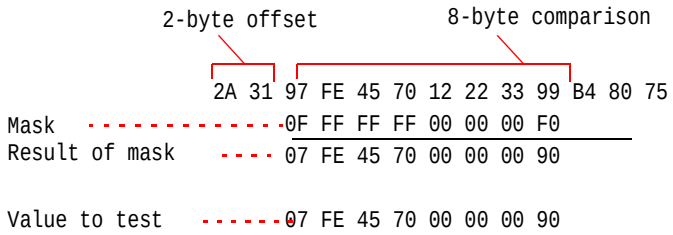

Assuming that the forward parameter is set to no, the packet is dropped because it matches this filter. The byte comparison works as follows:

- The system ignores 2A and 31 because of the 2-byte offset.
- The 9 in the third byte is also ignored, because the mask has a 0 (zero) in its place. The 7 in the third byte matches the value parameter's 7 for that byte.
- In the fourth byte, F and E match the fourth byte specified by the value parameter.
- In the fifth byte, 4 and 5 match the fifth byte specified by the value parameter.
- In the sixth byte, 7 and 0 match the sixth byte specified by the value parameter.
- The seventh (12), eighth (22), and ninth (33) bytes are ignored because the mask has zeroes in those places.
- In the tenth byte, 9 matches the value parameter's 9 for that byte. The second 9 in the packet's tenth byte is ignored because the mask has a 0 (zero) in its place.

Examples of a generic call filter

The following example shows how to define a generic call filter. The filter's purpose is to prevent inbound packets from resetting the session timer.

In the input filter, the default values are left unchanged in the gen-filter subprofile, so all packets are matched. Also, the forwarding action is left at its default of no. In the output filter, the default values again match all packets, but the forwarding action is set to yes. Therefore, the filter does not prevent outbound packets from resetting the timer or placing a call.

```
admin> new filter out-only
FILTER/out-only read

admin> set input 1 valid = yes

admin> set output 1 valid = yes

admin> set output 1 forward = yes
```

```
admin> write
FILTER/out-only written
```

Following is a comparable RADIUS filter definition:

```
test-user Password = "test-pw"
    Ascend-Call-Filter = "generic in drop"
    Ascend-Call-Filter = "generic out forward"
```

Defining IP filters

IP filters affect only IP and related packets. They make use of high-level information in packets (for example, protocol numbers, logical addresses, and TCP or UDP ports).

IP filter settings in a local filter profile

The `ip-filter` subprofile contains the following parameters. The parameters are shown with their default values for input filters. The same values apply for output filter specifications.

```
[in FILTER/"":input-filters[n]]
type = ip-filter
```

```
[in FILTER/"":input-filters[n]:ip-filter]
protocol = 0
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
dest-address-mask = 0.0.0.0
dest-address = 0.0.0.0
Src-Port-Cmp = none
source-port = 0
Dst-Port-Cmp = none
dest-port = 0
tcp-estab = no
```

If you set the parameters in an input filter, only inbound packets are examined. If you set them in an output filter, only outbound packets are examined.

Parameter	Setting
`type`	Type of filter. Use the default value `generic-filter` for a generic filter. Only the parameters in the corresponding subprofile are applicable.
`protocol`	Protocol number. A number of 0 (zero) matches all protocols. If you specify a nonzero number, the system compares it to the Protocol field in each packet. For a list of assigned protocol numbers, see RFC 1700, *Assigned Numbers*, by Reynolds, J. and Postel, J., October 1994.
`source-address-mask`	Mask to be applied to the `source-address` value before the system compares that value to the source address of a packet.

Parameter	Setting
source-address	IP address. After applying the source-address-mask value, the system compares the result to the source address in a packet. For details, see "Filtering by source or destination IP address" on page 7-14.
dest-address-mask	A mask to be applied to the dest-address value before the system compares that value to the destination address of a packet.
dest-address	IP address. After applying the dest-address-mask value, the system compares the result to the destination address in a packet. For details, see "Filtering by source or destination IP address" on page 7-14.
src-port-cmp	Type of comparison to perform when comparing source port numbers. With a setting of none (the default), no comparison is made. You can specify that the filter matches the packet if the packet's source port number is less (less than), eql (equal to), gtr (greater than), or neq (not equal to) the source-port value.
source-port	Port number to be compared with the source port of a packet. TCP and UDP port numbers are typically assigned to services. For more details, see "Filtering by port numbers" on page 7-14.
dst-port-cmp	Type of comparison to perform when comparing destination port numbers. With a setting of none (the default), no comparison is made. You can specify that the filter matches the packet if the packet's destination port number is less (less than), eql (equal to), gtr (greater than), or neq (not equal to) the dest-port value.
dest-port	Port number to be compared with the destination port of a packet. TCP and UDP port numbers are typically assigned to services. For more details, see "Filtering by port numbers" on page 7-14.
tcp-estab	Enable/disable application of the filter only to packets in an established TCP session. Applicable only if the protocol number has been set to 6 (TCP).

IP filter settings in a RADIUS profile

In RADIUS, an IP filter entry is a value of the Ascend-Call-Filter or Ascend-Data-Filter attribute. To specify an IP filter value, use the following format:

"ip *dir action* [dstip *n.n.n.n/nn*] [srcip *n.n.n.n/nn*][[*proto*] [destport *cmp value*] [srcport *cmp value*] [est]]"

Note A filter definition cannot contain newline indicators. The syntax is shown here on two lines for printing purposes only.

Keyword or argument	Value
`ip`	Type of filter. Valid types specified by the Ascend-Data-Filter and Ascend-Call-Filter attributes are `generic` (the default) and `ip`.
`dir`	Specifies direction of the packets. You can specify `in` (to filter packets coming in to the Stinger unit or `out` (to filter packets going out of the Stinger unit).
`action`	Specifies the action that the system takes with a packet that matches the filter. You can specify either `forward` or `drop`.
`dstip` *`n.n.n.n/nn`*	If the `dstip` keyword is followed by a valid IP address, the filter will match only packets with that destination address. If a subnet mask portion of the address is present, the system compares only the masked bits. If the `dstip` keyword is followed by the zero address (0.0.0.0), or if this keyword and its IP address specification are not present, the filter matches all IP packets. For more details, see "Filtering by source or destination IP address" on page 7-14.
`srcip` *`n.n.n.n/nn`*	If the `srcip` keyword is followed by a valid IP address, the filter will match only packets with that source address. If a subnet mask portion of the address is present, the system compares only the masked bits. If the `srcip` keyword is followed by the zero address (0.0.0.0), or if this keyword and its IP address specification are not present, the filter matches all IP packets. For more details, see "Filtering by source or destination IP address" on page 7-14.
`proto`	A protocol number. A value of zero matches all protocols. If you specify a nonzero number, the system compares it to the Protocol field in packets. For a list of protocol numbers, see RFC 1700.
`dstport` *`cmp value`*	If the `dstport` keyword is followed by a comparison symbol and a number, the number is compared to the destination port of a packet. The comparison symbol can be < (less than), = (equal), > (greater than), or != (not equal). The port value can be one of the following names or numbers: `ftp-data (20)`, `ftp (21)`, `telnet (23)`, `smtp (25)`, `nameserver (42)`, `domain (53)`, `tftp (69)`, `gopher (70)`, `finger (79)`, `www (80)`, `kerberos (88)`, `hostname (101)`, `nntp (119)`, `ntp (123)`, `exec (512)`, `login (513)`, `cmd (514)`, or `talk (517)`. For more details, see "Filtering by port numbers" on page 7-14.

Keyword or argument	Value
`srcport cmp value`	If the `srcport` keyword is followed by a comparison symbol and a number, the number is compared to the source port of a packet. The comparison symbol can be < (less than), = (equal), > (greater than), or != (not equal to). The port value can be one of the following names or numbers: `ftp-data (20)`, `ftp (21)`, `telnet (23)`, `smtp (25)`, `nameserver (42)`, `domain (53)`, `tftp (69)`, `gopher (70)`, `finger (79)`, `www (80)`, `kerberos (88)`, `hostname (101)`, `nntp (119)`, `ntp (123)`, `exec (512)`, `login (513)`, `cmd (514)`, or `talk (517)`. For more details, see "Filtering by port numbers" on page 7-14.
`est`	If the `est` flag is present, it restricts application of the filter to packets in an established TCP session. The protocol number must be set to 6 (TCP), or the flag is ignored.

Filtering by source or destination IP address

When you specify a source or destination address in an IP filter, the system applies the filter's forwarding action to packets received from or sent to that address. If you also specify a subnet mask, the system applies the mask to the address value before comparing the resulting value to the source or destination address in a packet.

To apply the mask, the system translates both the mask and address values into binary format and then uses a logical AND to apply the mask to the address. The mask hides the bits whose positions match those of the binary zeroes in the mask. A mask of all zeros (the default) masks all bits. If the address value itself is also all zeros (the default), the filter matches any source or destination address. A mask of all ones (255.255.255.255) masks no bits, so the full address for a single host is compared to the address value.

You can use the address mask to mask out the host portion of an address, for example, or the host and subnet portion, so the specification matches the address to or from any host on a given network.

Filtering by port numbers

IP filters can specify a port number to be compared to the source or destination port (or both) in a packet. A port number of zero matches nothing. TCP and UDP port numbers are typically assigned to services. For a list of well-known port assignments, see RFC 1700, *Assigned Numbers*.

Note For security purposes, Lucent Technologies recommends that you filter all services from outside your domain that are not required. UDP-based services make you network particularly vulnerable to certain types of security attacks.

The specified type of comparison determines when a match occurs. If no comparison operator is specified in the filter, no comparison is made. You can specify that the filter matches the packet if the packet's port number is less than, equal to, greater than, or not equal to the port number specified in the filter.

Examples of an IP filter to prevent local address spoofing

IP-address spoofing typically occurs when a remote device illegally acquires a local address and uses it to try to break through a data filter. This section presents an example of a data filter that prevents IP-address spoofing. For related information, see “Example of per-session source address checking” on page 1-22.

The sample filter first defines three input filters. The first filter drops packets whose source address is on the local IP network. The second filter drops packets whose source address is the loopback address (127.0.0.0). The third input filter accepts all remaining source addresses (by specifying a source address of 0.0.0.0) and forwards them to the local network.

In this example, the local IP network has an IP address of 192.100.50.128, with a subnet mask of 255.255.255.192. These values are just arbitrary examples.

Note If you apply this filter to the Ethernet interface, the system drops IP packets it receives from the local LAN, and you are not able to telnet to the unit.

The following set of commands creates the first input filter, setting the type to `ip-filter`. The first filter specifies the source mask and address for the local network. If an incoming packet has the local address, the system drops it instead of forwarding it to the Ethernet network, because `forward` is set to No (the default).

```
admin> new filter ip-spoof
FILTER/ip-spoof read

admin> set input 1 valid = yes

admin> set input 1 type = ip-filter

admin> set input 1 ip-filter source-address-mask = 255.255.255.192

admin> set input 1 ip-filter source-address = 192.100.50.128
```

The next set of commands creates the second input filter, setting the type to `ip-filter`. The second filter specifies the loopback source address. If an incoming packet has the loopback address, the system drops it instead of forwarding it to the Ethernet network, because `forward` is set to No.

```
admin> set input 2 valid = yes

admin> set input 2 type = ip-filter

admin> set input 2 ip-filter source-address-mask = 255.0.0.0

admin> set input 2 ip-filter source-address = 127.0.0.0
```

The next set of commands creates the third input filter, setting the type to `ip-filter` and setting `forward` to `yes`. Except for `forward` = `yes`, the third filter uses all default values. Because `forward` is set to `yes`, the system forwards all remaining packets (those with nonlocal source addresses) to the Ethernet network.

```
admin> set input 3 valid = yes

admin> set input 3 forward = yes

admin> set input 3 type = ip-filter
```

The next set of commands creates an output filter, setting the type to `ip-filter` and the forwarding action to `yes`. This filter specifies the source mask and address for the local network. (Packets originating on the local network should be forwarded across the WAN.)

```
admin> set output 1 valid = yes
admin> set output 1 type = ip-filter
admin> set output 1 forward = yes
admin> set output 1 ip-filter source-address-mask = 255.255.255.192
admin> set output 1 ip-filter source-address = 192.100.50.128
admin> write
FILTER/ip-spoof written
```

Following is a comparable RADIUS filter definition:

```
test-user Password = "test-pw"
    Ascend-Data-Filter = "ip in drop srcip 192.100.50.128/26"
    Ascend-Data-Filter = "ip in drop srcip 127.0.0.0/8"
    Ascend-Data-Filter = "ip in forward"
    Ascend-Data-Filter = "ip out forward srcip 192.100.50.128/26"
```

Examples of an IP filter for more complex security issues

This section illustrates some of the issues you might need to consider when writing your own IP filters. However, the sample filter presented here does not address the fine points of network security. You might want to use this filter as a starting point and augment it to address your security requirements.

In this example, the local network supports a Web server, and the administrator needs to carry out the following tasks:

- Provide client access to the server's IP address.
- Restrict ingress traffic to all other hosts on the local network.

However, many local IP hosts need to access the Internet and use IP-based applications such as Telnet or FTP, so their response packets need to be directed appropriately to the originating host. In this example, the Web server's IP address is 192.9.250.5. The filter will be applied in `connection` profiles as a data filter.

The following set of commands creates the first input filter, setting the type to `ip-filter` and `forward` to `yes`, and configures the first filter to allow packets to reach the Web server's destination address at a destination TCP port that can be used for Telnet or FTP:

```
admin> new filter web-access
FILTER/web-access read
admin> set input 1 valid = yes
admin> set input 1 forward = yes
admin> set input 1 type = ip-filter
admin> set input 1 ip-filter protocol = 6
admin> set input 1 ip-filter dest-address-mask = 255.255.255.255
admin> set input 1 ip-filter dest-address = 192.9.250.5
admin> set input 1 ip-filter dst-port-cmp = eql
admin> set input 1 ip-filter dest-port = 80
```

The next set of commands creates the second input filter, with `type` set to `ip-filter` and `forward` set to `yes`. This filter allows inbound TCP packets in response to a local

user's outbound Telnet request, by specifying that TCP packets whose destination port number is higher than that of the source port are forwarded. (Telnet requests go out on port 23, and responses come back on some random port above port 1023.)

```
admin> set input 2 valid = yes
admin> set input 2 forward = yes
admin> set input 2 type = ip-filter
admin> set input 2 ip-filter protocol = 6
admin> set input 2 ip-filter dst-port-cmp = gtr
admin> set input 2 ip-filter dest-port = 1023
```

The next set of commands creates the third input filter, with `type` set to `ip-filter` and `forward` set to `yes`. This filter allows inbound RIP updates, by specifying that inbound UDP packets are forwarded if the destination port number is higher than that of the source port. (For example, suppose a RIP packet goes out as a UDP packet to destination port 520. The response to this request goes to a random destination port above port 1023.)

```
admin> set input 3 valid = yes
admin> set input 3 forward = yes
admin> set input 3 type = ip-filter
admin> set input 3 ip-filter protocol = 17
admin> set input 3 ip-filter dst-port-cmp = gtr
admin> set input 3 ip-filter dest-port = 1023
```

The following commands create the fourth input filter, setting the type to `ip-filter` and `forward` to `yes`. The fourth filter uses all default values, which allows unrestricted use of Ping and Traceroute. Unlike TCP and UDP, ICMP does not use ports, so a port comparison is unnecessary.

```
admin> set input 4 valid = yes
admin> set input 4 forward = yes
admin> set input 4 type = ip-filter
admin> write
FILTER/web-access written
```

Following are comparable RADIUS filter definitions:

```
Ascend-Data-Filter="ip in forward dstip 192.9.250.5/32 dstport = 80 proto 6"
Ascend-Data-Filter="ip in forward dstport > 1023 proto 6"
Ascend-Data-Filter="ip in forward dstport > 1023 proto 6"
Ascend-Data-Filter="ip in forward"
```

Defining TOS filters

Type-of-service (TOS) filters are used to enable proxy quality of service (QoS) for packets that match the filter specification. For TOS filters, the forwarding action in the filter has no effect.

TOS filter settings in a local filter profile

Defining a local TOS filter involves setting the following parameters. The parameters are shown with their default values for input filters. The same values apply for output filter specifications.

```
[in FILTER/"":input-filters[n]]
type = tos-filter
```

```
[in FILTER/"":input-filters[n]:tos-filter]
protocol = 0
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
dest-address-mask = 0.0.0.0
dest-address = 0.0.0.0
Src-Port-Cmp = none
source-port = 0
Dst-Port-Cmp = none
dest-port = 0
precedence = 000
type-of-service = normal
```

Parameter	Setting
type	Type of filter. Use the `tos-filter` value for TOS filters. Only the parameters in the corresponding subprofile are applicable.
`protocol`	Protocol number. A value of zero matches all protocols. If you specify a nonzero number, the system compares it to the Protocol field in each packet. For a list of protocol numbers, see RFC 1700.
`source-address-mask`	Mask to be applied to the `source-address` value before the system compares that value to the source address of a packet.
`source-address`	IP address. After applying the `source-address-mask` value, the system compares the result to the source address in a packet. For details, see "Filtering by source or destination IP address" on page 7-14.
`dest-address-mask`	Mask to be applied to the `dest-address` value before the system compares that value to the destination address of a packet.
`dest-address`	An IP address. After applying the `dest-address-mask` value, the system compares the result to the destination address in a packet. For details, see "Filtering by source or destination IP address" on page 7-14.
`src-port-cmp`	Type of comparison to perform when comparing source port numbers. With a setting of `none` (the default), no comparison is made. You can specify that the filter matches the packet if the packet's source port number is `less` (less than), `eql` (equal to), `gtr` (greater than), or `neq` (not equal to) the `source-port` value.

Parameter	Setting
source-port	Port number to be compared with the source port of a packet. TCP and UDP port numbers are typically assigned to services. For more details, see "Filtering by port numbers" on page 7-14.
dst-port-cmp	Type of comparison to perform when comparing destination port numbers. With a setting of none (the default), no comparison is made. You can specify that the filter matches the packet if the packet's destination port number is less (less than), eql (equal to), gtr (greater than), or neq (not equal to) the dest-port value.
dest-port	Port number to be compared with the destination port of a packet. TCP and UDP port numbers are typically assigned to services. For more details, see "Filtering by port numbers" on page 7-14.
precedence	Priority level of the data stream. The three most significant bits of the TOS byte are priority bits used to set precedence for priority queuing. When TOS is enabled and the packet matches the filter, the bits can be set to one of the following values (most significant bit first): ■ 000—Normal priority ■ 001—Priority level 1 ■ 010—Priority level 2 ■ 011—Priority level 3 ■ 100—Priority level 4 ■ 101—Priority level 5 ■ 110—Priority level 6 ■ 111—Priority level 7 (the highest priority)
type-of-service	Type of service of the data stream. The value of this parameter sets the four bits following the three most significant bits of the TOS byte. The four bits are used to choose a link according to the type of service. When TOS is enabled and the packet matches the filter, specify one of the following values: ■ normal—Normal service ■ cost—Minimize monetary cost ■ reliability—Maximize reliability ■ throughput—Maximize throughput ■ latency—Minimize delay

TOS filter settings in a RADIUS profile

In RADIUS, a TOS filter entry is a value of the Ascend-Filter attribute. To specify a TOS filter value, use the following format:

`iptos` *`dir`* `[ dstip` *`n.n.n.n/nn`* `] [ srcip` *`n.n.n.n/nn`* `][` *`proto`* `] [ destport` *`cmp value`* `] [ srcport` *`cmp value`* `][ precedence` *`value`* `] [ type-of-service` *`value`* `]`

Note A filter definition cannot contain newline indicators. The syntax is shown here on multiple lines for printing purposes only.

Keyword or argument	Description
`iptos`	Specifies an IP TOS filter.
`dir`	Specifies direction of the packets. You can specify `in` (to filter packets coming in to the Stinger unit or `out` (to filter packets going out of the Stinger unit).
`dstip` *`n.n.n.n/nn`*	If the `dstip` keyword is followed by a valid IP address, the TOS filter will set bytes only in packets with that destination address. If a subnet mask portion of the address is present, the system compares only the masked bits. If the `dstip` keyword is followed by the zero address (0.0.0.0), or if this keyword and its IP address specification are not present, the filter matches all IP packets. For more details, see "Filtering by source or destination IP address" on page 7-14.
`srcip` *`n.n.n.n/nn`*	If the `srcip` keyword is followed by a valid IP address, the TOS filter will set bytes only in packets with that source address. If a subnet mask portion of the address is present, the system compares only the masked bits. If the `srcip` keyword is followed by the zero address (0.0.0.0), or if this keyword and its IP address specification are not present, the filter matches all IP packets. For more details, see "Filtering by source or destination IP address" on page 7-14.
`proto`	A protocol number. A value of zero matches all protocols. If you specify a nonzero number, the system compares it to the Protocol field in packets. For a list of protocol numbers, see RFC 1700.
`dstport` *`cmp value`*	If the `dstport` keyword is followed by a comparison symbol and a port, the port is compared to the destination port of a packet. The comparison symbol can be < (less than), = (equal to), > (greater than), or != (not equal to). The port value can be one of the following names or numbers: `ftp-data (20)`, `ftp (21)`, `telnet (23)`, `smtp (25)`, `nameserver (42)`, `domain (53)`, `tftp (69)`, `gopher (70)`, `finger (79)`, `www (80)`, `kerberos (88)`, `hostname (101)`, `nntp (119)`, `ntp (123)`, `exec (512)`, `login (513)`, `cmd (514)`, or `talk (517)`. For more details, see "Filtering by port numbers" on page 7-14.

Keyword or argument	Description
`srcport` *`cmp value`*	If the `srcport` keyword is followed by a comparison symbol and a port, the port is compared to the source port of a packet. The comparison symbol can be < (less than), = (equal to), > (greater than), or != (not equal to). The port value can be one of the following names or numbers: `ftp-data (20)`, `ftp (21)`, `telnet (23)`, `smtp (25)`, `nameserver (42)`, `domain (53)`, `tftp (69)`, `gopher (70)`, `finger (79)`, `www (80)`, `kerberos (88)`, `hostname (101)`, `nntp (119)`, `ntp (123)`, `exec (512)`, `login (513)`, `cmd (514)`, or `talk (517)`. For more details, see "Filtering by port numbers" on page 7-14.
`precedence value`	Specifies the priority level of the data stream. The three most significant bits of the TOS byte are priority bits used to set precedence for priority queuing. If a packet matches the filter, the bits are set to the specified value (most significant bit first). Specify one of the following values: ■ `000`—Normal priority ■ `001`—Priority level 1 ■ `010`—Priority level 2 ■ `011`—Priority level 3 ■ `100`—Priority level 4 ■ `101`—Priority level 5 ■ `110`—Priority level 6 ■ `111`—Priority level 7 (the highest priority).
`type-of-service` *`value`*	Type of Service of the data stream. If a packet matches the filter, the system sets the four bits following the three most significant bits of the TOS byte to the specified value. The four bits are used to choose a link according to the type of service. Specify one of the following values: ■ `Normal (0)`—Normal service ■ `Disabled (1)`—Disable TOS ■ `Cost (2)`—Minimize monetary cost ■ `Reliability (4)`—Maximize reliability ■ `Throughput (8)`—Maximize throughput ■ `Latency (16)`—Minimize delay

Examples of defining a TOS filter

The following set of commands defines a TOS filter for TCP packets (protocol 6) that are destined for a single host at 10.168.6.24. The packets must be sent on TCP port 23. For incoming packets that match this filter, the priority is set at level 2. This relatively low priority means that an upstream router that implements priority queuing can drop these packets when it becomes loaded. The commands also set TOS

to prefer a low latency connection, which means that the upstream router will choose a fast connection if one is available, even if it has higher cost or lower bandwidth, or is less reliable than another available link.

```
admin> new filter jfans-tos-filter
FILTER/jfans-tos-filter read

admin> list input 1
[in FILTER/jfans-tos-filter:input-filters[1] (new)]
valid-entry = no
forward = no
Type = generic-filter
gen-filter = { 0 0 no no 00:00:00:00:00:00:00:00:00:00:00:00 00:00:00:0+
ip-filter = { 0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 none 0 none 0 no }
route-filter = { 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0 none }
tos-filter = { 0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 none 0 none 0 000 norma+

admin> set valid = yes

admin> set type = tos-filter

admin> list tos
[in FILTER/jfans-tos-filter:input-filters[1]:tos-filter (changed)]
protocol = 0
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
dest-address-mask = 0.0.0.0
dest-address = 0.0.0.0
Src-Port-Cmp = none
source-port = 0
Dst-Port-Cmp = none
dest-port = 0
precedence = 000
type-of-service = normal

admin> set protocol = 6

admin> set dest-address-mask = 255.255.255.0

admin> set dest-address = 10.168.6.24

admin> set dst-port-cmp = eql

admin> set dest-port = 23

admin> set precedence = 010

admin> set type-of-service = latency

admin> write
FILTER/jfans-tos-filter written
```

Following is a RADIUS user profile that contains a comparable filter definition:

```
jfan-pc Password = "secret"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.168.6.120,
    Framed-IP-Netmask = 255.255.255.0,
```

```
Ascend-Filter = "iptos in dstip 10.168.6.24/32 dstport = 23
   precedence 010 type-of-service latency"
```

Note Filter definitions cannot contain newline indicators. The preceding example shows the Ascend-Filter value on two lines for printing purposes only.

Defining route filters

Route filters are applied to RIP update packets to exclude routes from the local system's routing table, or to include routes in the table only after modifying their metrics. Route filter specifications are not supported in RADIUS. For route filters, the forwarding action in the filter has no effect.

In a local `filter` profile, the `route-filter` subprofile contains the following parameters. The parameters are shown with their default values for input filters. The same values apply for output filter specifications.

```
[in FILTER:input-filters[n]]
type = route-filter
```

```
[in FILTER:input-filters[n]:route-filter]
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
route-mask = 0.0.0.0
route-address = 0.0.0.0
add-metric = 0
action = none
```

If you set the parameters in an input filter, only inbound packets are examined. If you set them in an output filter, only outbound packets are examined.

Parameter	Setting
`type`	Type of filter. Use the default value `route-filter` for a route filter. Only the parameters in the corresponding subprofile are applicable.
`source-address-mask`	Mask to be applied to the `source-address` value before comparing that value to the source address of a RIP update packet.
`source-address`	IP address. After applying the `source-address-mask` value, the system compares the result to the source address in a RIP packet.
`route-mask`	Mask to be applied to the destination address of a route.
`route-address`	IP address. After applying the `route-mask` value, the system compares the result to routes in a RIP packet. If it finds a route with a matching destination, it takes the specified action.
`add-metric`	Number from 1 to 15, to be added to the metric value for a route that matches the filter specification, if the specified value for the `action` parameter is Add.

Parameter	Setting
action	An action to take on a route that matches the filter specification. Valid values are none (the default), accept (accept the route by allowing it to affect the routing table), deny (deny the route by not allowing it to affect the routing table), or add (add the value of the add-metric parameter to the route metric and accept the route).

Example of a filter that excludes a route

In this example, the defined input filters accept all inbound RIP packets except those with a destination of 90.0.0.0. Following are the commands entered to define the filter, and the system's responses:

```
admin> new filter route-test
FILTER/route-test read
admin> set input 1 valid = yes
admin> set input 1 type = route-filter
admin> set input 1 route route-mask = 255.0.0.0
admin> set input 1 route route-address = 90.0.0.0
admin> set input 1 route action = deny
admin> set input 2 valid = yes
admin> set input 2 type = route-filter
admin> set input 2 route action = accept
admin> write
FILTER/route-test written
```

In this sample route filter, any route that matches filter 1 is rejected, and all other routes are accepted (because they match filter 2).

Example of a filter that configures a route's metric

In this example, an output filter identifies the route 11.0.0.0 in outbound RIP packets and assigns a high metric to that route. Following are the commands entered and the system's responses:

```
admin> new filter metrics
FILTER/metrics read
admin> set output 1 valid = yes
admin> set output 1 type = route-filter
admin> set output 1 route route-mask = 255.0.0.0
admin> set output 1 route route-address = 11.0.0.0
admin> set output 1 route add-metric = 7
admin> set output 1 route action = add
admin> write
FILTER/metrics written
```

Defining dynamic remote filters

You can create RADIUS pseudo-user profiles that define data filters, and apply the filters to multiple local `connection` or RADIUS profiles by referring to the pseudo-user profile name.

When the Stinger unit receives a filter ID in an Access-Accept packet from RADIUS, it searches for a matching local filter. If it does not find one, the system requests the filter from the RADIUS server. You can specify how the system behaves if the filter referred to in a profile is not found. The system can either establish the session and log a message about the missing filter, or simply terminate the call.

Externally defined filters are cached locally for a configurable interval. The `filtcache` command displays statistics about each cached RADIUS filter profile, and enables you to flush profiles from the cache. For more information about the `filtcache` command, see the *Stinger Administration Guide*.

Current limitations on dynamic remote filters

In the current software version, the remote filter implementation is subject to the following limitations:

- Filters applied to dial-out calls are not supported in this release.
- Call filters, route filters, and TOS filters are not supported in this release. Only data filters are currently supported.

Overview of local profile settings

Following are the local parameters (shown with default values) related to dynamic remote filters:

```
[in ANSWER-DEFAULTS:session-info]
filter-required = no

[in CONNECTION:session-options]
filter-required = no
data-filter = ""

[in IP-GLOBAL]
default-filter-cache-time = 1440
```

Parameter	Setting
`filter-required`	Whether access to the filter is required for the session. With the default value of `no`, the system establishes the session even if the specified filter is not found. If the parameter is set to `yes`, the system disconnects the call if the filter is not found. This parameter does not apply if the profile does not refer to a filter by name. In the `answer-defaults` profile, this parameter is used for RADIUS user profiles that apply a filter and do not specify a value for Ascend-Filter-Required (50).

Parameter	Setting
`data-filter`	Name of a `filter` profile associated with the connection. The name can be that of a local profile or of a pseudo-user profile in RADIUS. However, if a local `connection` profile does not use authentication, it cannot specify a RADIUS filter profile.
`default-filter-cache-time`	Number of minutes to cache RADIUS filter profiles that do not include a value for Ascend-Cache-Time (57). The default is `1440` (24 hours). Once the cache timer expires, cached profiles are deleted from system memory. The next time a remote filter is needed, the system retrieves the profile from RADIUS and stores it in cache again. Keeping a profile in cache increases performance when establishing sessions that use the filter, at the cost of some system memory. If this parameter is set to `0` (zero), the default timer is disabled so that only RADIUS profiles that specify a cache time are cached.

Overview of RADIUS user profile settings

RADIUS user profile support for filter profiles is provided by the following vendor-specific attributes (VSAs):

RADIUS attribute	Value
Filter-ID (11)	Name of a local or remote filter profile associated with the connection.
Ascend-Filter-Required (50)	Whether access to the filter is required for the session. With the default value of Required-No (0), the system establishes the session even if the specified filter is not found. If the attribute is set to Required-Yes (1), the system disconnects the call if the filter is not found. This attribute does not apply if the profile does not refer to a filter by name. If no value is specified for this attribute, setting for the Filter-Required parameter in the `answer-defaults` profile is used to determine system behavior when the specified filter is not found.

A filter profile is a pseudo-user profile in which the first two lines have the following format:

`profile-name` `Password = "ascend" Service-Type = Outbound`

The *`profile-name`* value is any name you assign to the profile. Duplicate filter names are not allowed. If a local `filter` profile is already stored, the system does not retrieve a filter profile of the same name from the RADIUS server. The `filter` profile definitions can include the following attribute-value pairs:

RADIUS attribute	Value
Ascend-Data-Filter (242)	Filter definition using one of the following formats: "generic *dir action offset mask value compare* [more]" "ip *dir action* [dstip *n.n.n.n/nn*] [srcip *n.n.n.n/nn*][[*proto*] [destport *cmp value*] [srcport *cmp value*] [est]]" For details, see "Defining generic filters" on page 7-6 and "Defining IP filters" on page 7-11.
Ascend-Cache-Refresh (56)	Whether the timer for cached routes in this profile is reset each time a new session that refers to the pseudo-user profile becomes active. Refresh-No (0) does not reset the timer. Refresh-Yes (1) resets the cache timer when a session referring to the profile becomes active.
Ascend-Cache-Time (57)	Number of minutes to cache the profile. Once the cache timer expires for a RADIUS profile, the profile is deleted from system memory. The next time it is needed, the system retrieves it from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of route lookups, at the cost of some system memory. The minimum possible cache time is 0 minutes, which causes the system to retrieve the profile every time it is needed. This value is usually not desirable. If no value is specified for this attribute, the setting for the Default-Filter-Cache-Time parameter in the IP-Global profile is used.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the system must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius
```

```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *Stinger Reference*.

Examples of configuring a filter profile in RADIUS

Following is a sample RADIUS filter profile:

```
filter-c Password = "ascend", Service-Type = Outbound
   Ascend-Cache-Time = 20,
   Ascend-Cache-Refresh = Refresh-Yes,
   Ascend-Data-Filter = "ip out forward tcp dstip 10.1.1.3/16",
   Ascend-Data-Filter = "ip out drop"
```

The cache timer has been set to 20 minutes, and the timer is reset each time the filter is applied to a session.

The following commands configure a default cache time for RADIUS filter profiles:

```
admin> read ip-global
IP-GLOBAL read

admin> set default-filter-cache-time = 180

admin> write
IP-GLOBAL written
```

Following is a sample RADIUS filter profile that uses the default instead of specifying a value for Ascend-Cache-Time (57):

```
filter-e Password = "ascend", Service-Type = Outbound
   Ascend-Data-Filter = "ip out forward tcp dstip 10.2.2.2/28",
   Ascend-Data-Filter = "ip out drop"
```

Examples of applying remote filters

The following commands modify a `connection` profile so that the session uses a remote filter and the system disconnects the call if the filter is not found:

```
admin> read connection p50-v2
CONNECTION/p50-v2 read

admin> set session-options data-filter = filter-c

admin> set session-options filter-required = yes

admin> write
CONNECTION/p50-v2 written
```

The following RADIUS profile applies the same filter profile and has the same requirements. This profile also specifies how the filters must be cached for this connection.

```
p50-v2 Password = "my-password", Service-Type = Framed
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.1.1.1,
    Framed-IP-Netmask = 255.0.0.0,
    Filter-ID = "filter-c",
    Ascend-Filter-Required = Required-Yes
```

The following commands configure the system to reject incoming calls when the RADIUS user profile specifies a filter that is not found and the user profile does not explicitly state what to do if the filter is not found:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set session-info filter-required = yes

admin> write
ANSWER-DEFAULTS written
```

Following is a sample RADIUS profile that uses the default instead of specifying a value for Ascend-Filter-Required (55):

```
p50-v2 Password = "my-password", Service-Type = Framed
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.1.1.1,
```

```
Framed-IP-Netmask = 255.0.0.0,
Filter-ID = "filter-c"
```

Applying a filter to an interface

When you apply a filter to a WAN interface, it takes effect when the connection is established.

Packets can pass through both a data filter and call filter on a WAN interface. When both a data filter and call filter are applied to the same interface, the data filter is applied first.

Settings in local profiles

To apply a filter to an interface, set the following parameters (shown with their default settings):

```
[in ANSWER-DEFAULTS]
use-answer-for-all-defaults = yes

[in ANSWER-DEFAULTS:session-info]
call-filter = ""
data-filter = ""
filter-persistence = no

[in CONNECTION/"":session-options]
call-filter = ""
data-filter = ""
filter-persistence = no

[in CONNECTION/"":ip-options]
route-filter = ""
tos-filter = ""

IP-INTERFACE { { any-shelf any-slot 0 } 0}
route-filter = ""

ETHERNET { any-shelf any-slot 0 }
filter-name= ""
```

Parameter	Setting
use-answer-for-all-defaults	Whether values in the answer-defaults profile override values in the default Ethernet profile when the Stinger unit uses RADIUS to validate an incoming call.
call-filter	Name of a filter profile. For details, see "Examples of applying a call filter to a WAN interface" on page 7-32. The setting in the answer-defaults profile is used only for RADIUS-authenticated connections that do not include a call filter.
data-filter	Name of a filter profile. For details, see "Examples of applying a data filter to a WAN interface" on page 7-31. The setting in the answer-defaults profile is used only for RADIUS-authenticated connections that do not include a data filter.

Parameter	Setting
`filter-persistence`	Enable/disable filter persistence across connection state changes.
`route-filter`	Name of a `filter` profile. For details, see "Examples of applying a route filter to a WAN or LAN IP interface" on page 7-33.
`tos-filter`	Name of a `filter` profile. For details, see "Examples of applying a TOS filter to a WAN interface" on page 7-32.
`filter-name`	Name of a `filter` profile. For details, see "Example of applying a filter to a LAN interface" on page 7-34.

Settings in RADIUS profiles

The following RADIUS attribute-value pairs are used to apply a filter to a WAN connection:

RADIUS attribute	Value

Ascend-Call-Filter (243)	Filter definition using one of the following formats:
	`"generic` *`dir action offset mask value compare`* `[more]"`
	`"ip` *`dir action`* `[ dstip` *`n.n.n.n/nn`* `] [ srcip` *`n.n.n.n/nn`* `][[` *`proto`* `] [ destport` *`cmp value`* `] [ srcport` *`cmp value`* `] [est]]"`
	For details, see "Defining generic filters" on page 7-6 and "Defining IP filters" on page 7-11.

Ascend-Data-Filter (242)	Filter definition using one of the following formats:
	`"generic` *`dir action offset mask value compare`* `[more]"`
	`"ip` *`dir action`* `[ dstip` *`n.n.n.n/nn`* `] [ srcip` *`n.n.n.n/nn`* `][[` *`proto`* `] [ destport` *`cmp value`* `] [ srcport` *`cmp value`* `] [est]]"`
	For details, see "Defining generic filters" on page 7-6 and "Defining IP filters" on page 7-11.

Ascend-Filter (90)	String-format filter specification using the following format: `iptos` *`dir`* `[ dstip` *`n.n.n.n/nn`* `] [ srcip` *`n.n.n.n/nn`* `][` *`proto`* `] [ destport` *`cmp value`* `] [ srcport` *`cmp value`* `][ precedence` *`value`* `] [ type-of-service` *`value`* `]` For details, see "Defining TOS filters" on page 7-17.
Filter-ID (11)	Name of a local `filter` profile that defines a data filter. The next time the system accesses the RADIUS user profile in which this attribute appears, the referenced filter is applied to the connection.

How the system uses answer-defaults profile settings

When the use-answer-for-all-defaults parameter is set to yes (the default), the system uses the settings in the answer-defaults profile to create a baseline profile for RADIUS-authenticated calls. The unit uses the baseline values for settings that are not specified in the caller's RADIUS profile. For example, if the caller's RADIUS profile does not apply a data filter, a call filter, or both, and the use-answer-for-all-defaults parameter is set to yes, any filters applied in the answer-defaults profile are applied to the authenticated connection. But if the caller's profile does apply a data or call filter, filters applied in the answer-defaults profile are used.

Examples of applying a data filter to a WAN interface

When you apply a data filter, its forwarding action (forward or drop) affects the actual data stream by preventing certain packets from reaching the Ethernet network from the WAN, or vice versa. Data filters do not affect the idle timer, and a data filter applied to a connection profile does not affect the answering process. In the following examples, the Stinger unit supports the following two local filter profiles:

```
admin> dir filter
370  09/13/1998 15:04:31  ip-spoof
372  09/13/1998 15:04:43  web-access
```

Following is an example of applying a data filter:

```
admin> read conn tlynch
CONNECTION/tlynch read

admin> set session data-filter = ip-spoof

admin> write
CONNECTION/tlynch written
```

Following is a comparable RADIUS profile:

```
tlynch Password = "secret"
    Service-Type = Framed-User,
    Framed-Protocol = MPP,
    Framed-IP-Address = 10.10.10.64,
    Framed-IP-Netmask = 255.255.255.0,
    Filter-Id = "ip-spoof"
```

The following RADIUS profile refers to both local filters:

```
tlynch Password = "secret"
    Service-Type = Framed-User,
    Framed-Protocol = MPP,
    Framed-IP-Address = 10.10.10.64,
    Framed-IP-Netmask = 255.255.255.0,
    Filter-Id = "ip-spoof",
    Filter-Id = "web-access"
```

As is always the case with filters, the order in which they are applied within the user profile is significant. If the Stinger unit supports multiple filter profiles with similar names, it attempts to match the first filter profile to the characters specified in the user profile.

Following is an example of defining an antispoofing filter within the user's RADIUS profile:

```
tlynch Password = "secret"
    Service-Type = Framed-User,
    Framed-Protocol = MPP,
    Framed-IP-Address = 10.10.10.64,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-Data-Filter = "ip in drop srcip 192.100.50.128/26",
    Ascend-Data-Filter = "ip in drop srcip 127.0.0.0/8",
    Ascend-Data-Filter = "ip in forward",
    Ascend-Data-Filter = "ip out forward srcip 192.100.50.128/26"
```

Examples of applying a call filter to a WAN interface

Call filters prevent unnecessary connection time and help the system distinguish active traffic from "noise." By default, any traffic to a remote site triggers a call, and any traffic across an active connection resets the connection's idle timer.

The following commands apply a filter to a WAN connection and set the idle timer to 20 seconds. If no packets get through the call filter in either direction for 20 seconds, the connection is closed.

```
admin> read conn bob
CONNECTION/bob read

admin> set session call-filter = out-only

admin> set session idle-timer = 20

admin> write
CONNECTION/bob written
```

Following is a comparable RADIUS profile:

```
bob Password = "secret"
    Service-Type = Framed-User,
    Framed-Protocol = MPP,
    Framed-IP-Address = 10.10.10.23,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-Idle-Limit = 20,
    Ascend-Call-Filter = "generic in drop",
    Ascend-Call-Filter = "generic out forward"
```

Examples of applying a TOS filter to a WAN interface

TOS filters instruct the system to set priority bits and type-of-service (TOS) classes of service on behalf of customer applications. The Stinger unit does not implement priority queuing, but it does set information that can be used by upstream routers to prioritize and select links for particular data streams. TOS filters specify which bits to set in the TOS header of IP packets.

The following set of commands applies a TOS filter to a `connection` profile. When the incoming data stream contains packets that match the TOS filter specification, the proxy-QoS and TOS settings specified in the filter are set in those packets.

```
admin> read connection jfan-pc
CONNECTION/jfan-pc read

admin> set ip-options tos-filter = jfans-tos-filter
```

```
admin> write
CONNECTION/jfan-pc written
```

Following is a comparable RADIUS profile in which the TOS filter is specified by the Filter-ID attribute:

```
jfan-pc Password = "johnfan"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.168.6.120,
    Framed-IP-Netmask = 255.255.255.0,
    Filter-ID = "jfans-tos-filter"
```

Following is a RADIUS profile in which the TOS filter is specified within the profile:

```
jfan-pc Password = "johnfan"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.168.6.120,
    Framed-IP-Netmask = 255.255.255.0,
Ascend-Filter = "iptos in dstip 10.1.1.1/32 dstport = 23 precedence
010 type-of-service latency"
```

Note Filter definitions cannot contain newline indicators. The preceding example shows the specification on two lines for printing purposes only.

Examples of applying a route filter to a WAN or LAN IP interface

Route filters specify which routes in RIP update packets will be allowed to affect the routing table. They can also be used to increase the metric that the system assigns to a route before adding it to the routing table.

When a route filter is applied to an IP interface, the system monitors RIP packets on that interface and takes a specified action if a route matches the filter specifications. Depending on how the filter is defined, it can apply to inbound RIP packets, outbound RIP packets, or both. Route filters are supported only in `filter` profiles defined locally in the command-line interface, not in filters defined in RADIUS.

Route filters do not stop RIP update packets from being forwarded. Rather, their action determines whether the system adds matching routes to its routing table.

Following is an example of applying a filter in a `connection` profile:

```
admin> read conn bdv
CONNECTION/bdv read

admin> set ip-options route-filter = route-test

admin> write
CONNECTION/bdv written
```

Following is an example of applying a route filter to a local IP interface:

```
admin> read ip-interface { { 1 8 1 } 0 }
IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 } read

admin> set route-filter = route-test

admin> write
IP-INTERFACE/{ { shelf-1 first-control-module 1 } 0 } written
```

Example of applying a filter to a LAN interface

Ethernet interfaces are connected routes, so call filters are not applicable. However, you can apply a data filter that affects which packets are allowed to enter or leave the Ethernet network. A filter applied to an Ethernet interface takes effect immediately. If you change any settings in a `filter` profile, the changes apply as soon as you save the `filter` profile.

Note Use caution when applying a filter to the Ethernet interface. You could inadvertently render the Stinger unit inaccessible from the local LAN.

The following set of commands applies a filter to an Ethernet interface of a T1000 module in slot 3:

```
admin> read ether { 1 3 1 }
ETHERNET/{ shelf-1 slot-3 1 } read

admin> set filter-name = dstip

admin> write
ETHERNET/{ shelf-1 slot-3 1 } written
```

Index

A

B

C

D

E

F

G

H

I

L

M

N

O

P

Q

R

S

T

U

V

W