

Release Note

AR Series Router Software Release 1.8

Introduction	2
New Release Numbering Scheme	2
New Model Names	3
New Hardware Platforms	4
AR320 and AR330	4
AR720	5
Flow Control	6
Operation	6
HTTP Server	6
Mail Subsystem	7
Integrated Services Digital Network (ISDN)	7
LAPD Signalling	7
Non-Associated Signalling	7
Internet Protocol (IP)	11
Finger Client	11
RIP Timers	11
IP Address Pools	12
Control Over IP Packet Fragmentation	13
NAT	13
Time Division Multiplexing (TDM)	14
E1/T1 TDM	14
Unstructured Mode	14
PBX	15
Layer 2 Tunnelling Protocol (L2TP)	16
IP Security (IPsec)	17
Security Associations	18
Security Policy Database	18
ISAKMP/IKE Key Management	21
Pre-IPsec Security Associations	23
MIBs	24
Other Enhancements	26
Availability	34
Installation	34
Upgrading From an Older Release	34
Upgrading From Software Release 7.2	36
Upgrading From Software Release 7.4	36
Upgrading from Software Release 7.6	37
Upgrading from Software Release 7.7	37

Introduction

Allied Telesyn International announces the release of Software Release 1.8.1 for the AR series of multiprotocol routers. This release note describes the new features and enhancements to the AR Router for Software Release 1.8.1, and should be read in conjunction with the *AR Series Router Reference Manual for Software Release 1.8* (Document Number C613-03014-00 REV B).

WARNING: Information in this release note is subject to change without notice and does not represent a commitment on the part of Allied Telesyn International. While every effort has been made to ensure that the information contained within the document and the features and changes described are accurate, Allied Telesyn International can not accept any type of liability for errors in, or omissions arising from the use of this information.

Software Release 1.8 includes significant changes to the release numbering scheme and model names as a result of Allied Telesyn's acquisition of the Network iQ router product range from Teltrend, Inc. in June 1999.

This release also adds support for three new models to the router product range—the AR320, AR330 and AR720.

The most significant enhancement in Software Release 1.8 is the continued development of the router's advanced Internet security features with the addition of a fully RFC-compliant IPsec (IP Security) implementation.

New Release Numbering Scheme

As a result of Allied Telesyn's acquisition of the Network iQ router range from Teltrend, Inc., the numbering scheme for software releases has been changed. Software Releases are now numbered 1.X.X, rather than 7.X.X. This means that the successor to Software Release 7.7 for Network iQ routers is called Software Release 1.8, not Software Release 7.8. This change brings the release numbering in to line with Allied Telesyn's AR product range.

New Model Names

Products from the Network iQ router range have been incorporated into Allied Telesyn's AR series. The following table shows the relationship between the original Network iQ product names and the new AR series product names.

Table 1: New names of Network iQ products.

Network iQ Product Names	Allied Telesyn AR Series Product Names
NiQ 820	AR300L(S)
NiQ 821	AR350
NiQ 822	AR370(S)
NiQ 822U	AR370(U)
-	AR320
-	AR330
NiQ 840	AR300(S)
NiQ 860	AR310(S)
NiQ 890	AR390
NiQ 891	AR395
-	AR720
PRI E1/T1 ICM	AR020 PIC E1/T1 PRI
BRI1 ICM	AR021(S) PIC BRI (S)
BRI1-U ICM	AR021(U) PIC BRI (U)
ETH1 ICM	AR022 PIC Eth
SYN1 ICM	AR023 PIC Syn
ASYN4 ICM	AR024 PIC Asyn4
PRI1/G.703 ICM	AR025 PIC PRI/G.703
CMAC	AR012 CMAC
EMAC	AR010 EMAC
CEMAC	AR011 ECMAC

In all cases the new AR product has identical functionality to the original Network iQ product.

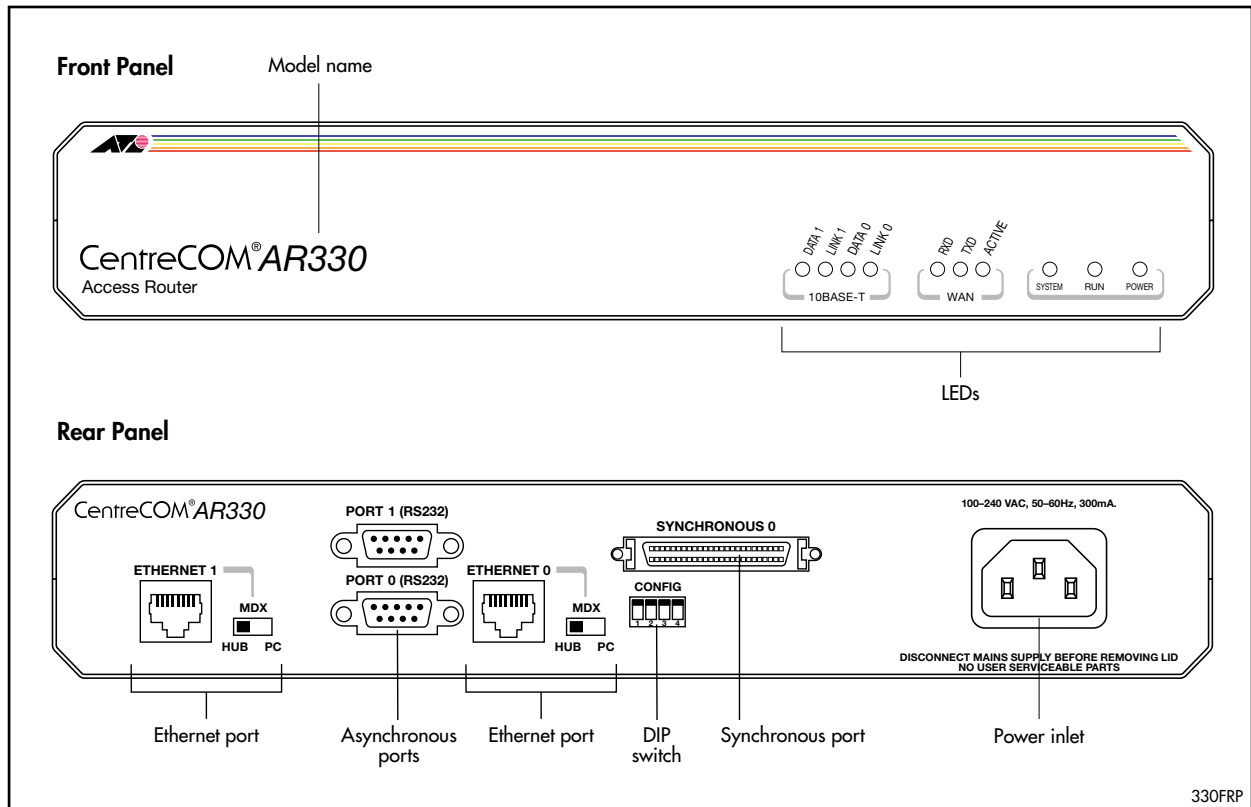
New Hardware Platforms

Software Release 1.8 adds support for three new models—the AR320, AR330 and AR720.

AR320 and AR330

The AR320 and AR330 access routers feature dual Ethernet LAN ports, and are ideally suited to LAN/WAN firewall applications or for organisations wanting secure connections between LAN segments (Figure 1 on page 4).

Figure 1: Front and rear views of the AR320 and AR330 access routers.




The main features of the AR320 and AR330 are:

- 68360 processor.
- 512K of EPROM.
- 8 MBytes of DRAM.
- 2 MBytes of FLASH memory.
- 2 Ethernet LAN 10BASE-T ports.
- 2 RS-232 asynchronous serial ports with DB9 male connectors.
- 1 high speed synchronous serial port (on model AR330).
- A MAC compression/encryption card slot.

AR720

The AR720 router is the first model in a new series of modular network access routers. The AR720 features a high performance RISC-based architecture, 100Mbps Ethernet, two *Port Interface Card (PIC)* bays that accept a range of PIC cards offering extensive network support, and a dedicated slot for MAC cards providing hardware-based encryption and compression (Figure 2 on page 5).

Figure 2: Front and rear views of the AR720 modular network access router.



The main features of the AR720 are:

- 50 MHz RISC processor.
- 1 MByte of EPROM.
- 16 MBytes of synchronous DRAM.
- 4 MBytes of FLASH memory, expandable to 8 MBytes.
- 128 KBytes of battery backed SRAM.
- A high performance 10/100 Fast Ethernet LAN port.
- 2 RS-232 asynchronous serial ports.
- 2 PIC slide-in bays.
- A MAC compression/encryption card slot.

The two PIC bays can accommodate any combination of the following PIC slide-in interface cards:

- AR020 PRI E1/T1 PIC, 1 Primary Rate E1/T1 port.
- AR021(S) BRI-S/T PIC, 1 Basic Rate ISDN S/T port.
- AR021(U) BRI-U PIC, 1 Basic Rate ISDN U port.
- AR022 ETH PIC, 1 Ethernet LAN AUI/10BASET port.
- AR023 SYN PIC, 1 Synchronous port with universal 50-way AMPLIMITE connector.
- AR024 ASYN4 PIC, 4 Asynchronous ports with RJ45 connectors.

- AR025 PRI E1 PIC, 1 Primary Rate/G.703 E1 port.

The dedicated MAC slot can accommodate any of the following MAC cards:

- AR010 EMAC, Encryption MAC card.
- AR011 ECMAC, Compression/Encryption MAC card.
- AR012 CMAC, Compression MAC card

Flow Control

As of Software Release 1.8, asynchronous ports on all AR300 Series routers, the AR720 router and the AR024 ASYNC4 PIC default to hardware flow control. In previous releases, asynchronous ports on some products defaulted to software (Xon/Xoff) flow control.

Operation

HTTP Server

A built-in HTTP server allows the router to serve HTML pages from FLASH memory in response to requests from a remote web browser. The HTTP server is compatible with any HTTP/1.1-compliant browser and is enabled by default. The HTTP server can be explicitly enabled or disabled using the commands:

```
ENABLE HTTP SERVER
DISABLE HTTP SERVER
```

The HTTP server implements basic authentication. When a user attempts to access the router via a web browser the HTTP server will request authentication from the browser, and the browser will in turn prompt the user for a username and password. The username and password entered by the user must match a user defined in the router's User Authentication Database.

The router's default home page can be set using the command:

```
SET HTTP HOMEPAGE=filename.htm
```

This allows a custom home page to be developed with links to other URLs, or to pages stored in FLASH. The default home page is HOMEPAGE.HTM.

All HTTP requests and authorisation failures are logged to the Logging Facility. Debugging can also be enabled or disabled using the commands:

```
ENABLE HTTP DEBUG
DISABLE HTTP DEBUG
```

Debug messages display TCP state changes, GET and POST requests, status messages sent in response to requests, and authorisation attempts.

The current status of the HTTP server is displayed using the command:

```
SHOW HTTP
```

Mail Subsystem

The router has a built-in email client and SMTP (*Simple Mail Transfer Protocol*) server to enable email messages to be sent from the router to remote mail systems using SMTP. The email client generates messages that comply with RFC 822 (*Standard for the Format of ARPA Internet Text Messages*). The SMTP server implements RFC 821 (*Simple Mail Transfer Protocol*) for the transmission of mail messages. The SMTP server can only transmit email messages; it can not accept email messages from other mail systems.

Software Release 1.8 adds support for the ETRN option:

```
MAIL TO=destination {FILE=filename|MESSAGE=message}  
[SUBJECT=subject] [ETRN=mail-domain]
```

The ETRN option sends an ETRN request (as defined in RFC 1985) to the remote mail server to forward any queued mail messages for the specified mail domain or host name. This can be used to assist mail servers that are connected to the Internet via dial-up rather than permanent connections. A trigger can be created to send an ETRN message to the email service provider each time the router connects to the Internet.

Integrated Services Digital Network (ISDN)

LAPD Signalling

PRI interfaces do not support bus configurations with multiple devices and therefore do not require Terminal Endpoint Identifiers (TEIs). However, for compatibility with some ISDN networks around the world, the logical channel used by LAPD for TEI management is still present on PRI interfaces.

Both Primary Rate and Basic Rate ISDN interfaces use a Service Access Point Identifier (SAPI) of 63 for TEI management and a SAPI of 000 for Q.931 call control.

The RESET Q931 command has been modified to allow a Q.931 interface, an active call, or all active calls on an interface to be reset:

```
RESET Q931=interface [CALL={call-index|ALL}]
```

A RESTART message for the interface or call(s) is sent to the network. The *call-index* value must be the index for the Q.931 call. To display a list of Q.931 calls, use the command:

```
SHOW Q931 CALL
```

Non-Associated Signalling

Normally, an ISDN interface will use its own D channel for signalling for the calls that are made on the interface. However, it is now possible to configure the D channel on one interface to provide the signalling for a number of other ISDN interfaces. The advantage of this is that the D channels which are unused for signalling can then be used as B channels, since on a PRI interface the D channel and B channels have the same bandwidth and underlying signalling structure.

This feature is known as *non-associated signalling* or *common D channel signalling*. The ISDN network must support the feature. Currently the router only supports this feature if the Q.931 profile of the participating interfaces is set to JAPAN. LAPD commands are used to set up the interfaces that are taking part in non-associated signalling, while Q.931 commands are used to give each interface a unique ID.

The SET LAPD command has been modified to support this new feature:

```
SET LAPD=interface [NASMODE={NORMAL|MASTER|SLAVE}]  
[NASMASTER=interface]
```

To set up an ISDN interface to be a master interface for non-associated signalling, use the command:

```
SET LAPD=instance NASMODE=MASTER
```

To set up an ISDN interface to be a slave interface for non-associated signalling, use the command:

```
SET LAPD=instance NASMODE=SLAVE NASMASTER=master-interface
```

where *master-interface* is the instance number or interface name of an ISDN interface whose NASMODE is MASTER.

To identify the ISDN interfaces for non-associated signalling, the SET Q931 command has been modified to allow a Q.931 interface to be assigned an interface identifier:

```
SET Q931=instance INTID=hex-string
```

hex-string is a sequence of hexadecimal digits specifying the interface ID in hexadecimal. The interfaces operating in non-associated signalling mode and their interface IDs will be arranged by subscription to the ISDN service provider. Note that the interface ID is a hexadecimal value; if the interface ID was, for example, the digit "0", the interface ID would have to be entered as INTID=30, since 30 is the hexadecimal value for the digit 0. The format of interface identifiers must be clearly understood and this information should be explicitly requested from the ISDN service provider.

The SHOW LAPD command has been modified to display the corresponding LAPD parameters (Figure 3 on page 9). Two new fields, *NAS Mode* and *NAS Master*, have been added (Table 2 on page 9).

Figure 3: Example output from the SHOW LAPD command for a Primary Rate Interface.

```

Interfaces:
ISDN      Type      TEI Mode      Debug      TEI      NAS mode      NAS master
-----
PRI0      TE        nonAuto      off        000      Normal        -
PRI0      TE        nonAuto      off        000      Normal        -
-----
SAPs:
ISDN      SAPI      T200      T201      T202      T203      N200      N201      N202      k
-----
PRI0      063      000010      -        -        000100      000003      000260      -      007
          000      000010      -        -        000100      000003      000260      -      007
PRI1      063      000010      -        -        000100      000003      000260      -      007
          000      000010      -        -        000100      000003      000260      -      007
-----
DLCs:
ISDN      SAPI      CES      TEI      State      V(S)      V(A)      rxN(S)      V(R)      rxN(R)
-----
PRI0      063      000      127      bcast      -        -        -        -        -
          000      000      000      ALIVE      0021      0021      0076      0077      0021
PRI0      063      000      127      bcast      -        -        -        -        -
          000      000      000      ALIVE      0014      0014      0051      0052      0014
-----
Packet parameters:
-----
PRI0
Packet mode TEIs: -
Packet mode SPIDs: -
PRI1
Packet mode TEIs: -
Packet mode SPIDs: -
-----

```

Table 2: New fields in the output of the SHOW LAPD command.

Parameter	Meaning
NAS mode	Non-associated signalling mode or common D channel mode. One of "Normal" (this interface's D channel does the signalling for this interface and no other interface), "Master" (this interface's D channel does the signalling for this interface and other interfaces) or "Slave" (another interface's D channel does the signalling for this interface).
NAS master	Non-associated signalling or common D channel master interface. If this interface's NAS mode is "Slave", the NAS master gives the interface whose D channel will provide the signalling channel for this interface. If the NAS mode is "Normal" or "Master", this field contains "-".

The SHOW Q931 command has been modified to display the corresponding Q.931 parameters (Figure 4 on page 10). Two new fields, *Common D channel* and *Interface ID*, have been added (Table 3 on page 10).

Figure 4: Example output from the SHOW Q931 command.

```

Q.931 interface ... BRI0
Profile ..... NI1-BR
ASD state ..... Operational
Data rate ..... 64k
Number 1 ..... -
Sub-address 1 ..... -
Number 2 ..... -
Sub-address 2 ..... -
No number ..... Accept
No sub-address .... Accept
DLC1
  State ..... Established
  SPID state ..... OP
  SPID file state ... 3 (Auto SPID successful)
  Current SPID ..... 62155542310101
  USID ..... 0
  Terminal ID ..... 1
DLC2
  State ..... Established
  SPID state ..... OP
  SPID file state ... 3 (Auto SPID successful)
  Current SPID ..... 62155579340101
  USID ..... 0
  Terminal ID ..... 2
Common D channel
  Interface ID .... 00
TSPID ..... 20
T301 ..... -
T302 ..... -
T303 ..... 4
T304 ..... 15
T305 ..... 30
T308 ..... 4
T309 ..... 90
T310 ..... -
T313 ..... 4
T314 ..... -
T316 ..... -
T317 ..... -
T318 ..... -
T319 ..... -
T321 ..... -
T322 ..... 4

```

Table 3: New fields displayed in the output of the SHOW Q931 command.

Parameter	Meaning
Common D channel	Parameters concerning non-associated signalling, or common D channel.
Interface ID	The non-associated signalling, or common D channel, interface identifier.

Internet Protocol (IP)

Finger Client

The finger protocol provides a mechanism for exchanging user information between a finger client and a finger server. The router's finger client can be used to send finger requests to a finger server for information about a specific user, or to request a list of all logged in users on the server. The information returned depends on the finger server, but typically includes the user's login name, real name, home directory, shell type, login details, and mail status. Other information may also be returned, and signature files may also be appended to the reply.

A finger query is sent to the finger server on a host or hosts using the command:

```
FINGER [username]@host[@host]... [DETAIL={HIGH|LOW}]
```

The response from the finger server is sent to the terminal or Telnet session from which the command was entered.

A typical use of the finger client is to trigger the download of new email from an Internet Service Provider (ISP) whenever a link to that ISP is brought up. In this scenario, the ISP's mail host runs a finger server which responds to finger queries from a subscriber by downloading any new email for the subscriber's account. On the router, a trigger is created which is activated whenever the ISDN call comes up and becomes active. The trigger runs a script which sends a finger query to the ISP with the subscriber's username. Whenever the link to the ISP is brought up, the mail host is automatically polled to see if there is any new mail, without any user intervention.

RIP Timers

The router fully implements the Routing Information Protocol (RIP) as described in RFC 1058, as well as extensions for RIP version 2 (RFC 1723) and RIP on demand (RFC 1582). Software Release 1.8 enhances the router's support for RIP by making the timers that control the operation of RIP user configurable. The operation of RIP is controlled by four timers whose values are set globally:

- The UPDATE timer sets the time interval between RIP updates for all interfaces not using RIP on demand.
- The INVALID timer sets the time interval after which the router will deem a route to be invalid if no update has been received for the route.
- The HOLDDOWN timer sets the time interval, after a route has become invalid, during which the router will ignore updates for the route which would normally make the route valid again.
- The FLUSH timer sets the time interval from the last update of a route until the route is flushed from the route table.

After a valid update, the FLUSH and INVALID timers are restarted. When the INVALID timer expires the route is invalidated and the HOLDDOWN timer started. The FLUSH timer continues to run. When the HOLDDOWN timer expires valid updates for the route will result in the router being reinstated. When the FLUSH timer expires, the route is deleted from the route table.

The timers can be set using the command:

```
SET IP RIPTIMER UPDATE=time INVALID=time HOLDDOWN=time
FLUSH=time
```

The current values of the RIP timers can be displayed using the command:

```
SHOW IP RIPTIMER
```

WARNING: All routers in the network participating in RIP exchanges should use the same set of RIP timer values.

IP Address Pools

A new mechanism for managing the assignment of IP addresses to dynamic connections has been added. An IP address pool is a named collection of IP addresses that ACC, PPP and other modules can use when assigning dynamic IP addresses. The advantage of an address pool is that a finite number of IP addresses can be re-used by many clients. When a client is finished with the IP address (for example, when a dial-in SLIP connection terminates) the IP address is returned to the pool and is available for another client to use.

The router supports multiple methods for assigning IP addresses to dynamic dial-in calls. The router uses the following procedure to select the IP address assigned to a dial-in call:

1. If the user is authenticated via RADIUS, and the RADIUS response supplies an IP address, then that IP address is used.
2. If the user is authenticated via TACACS, and a domain string is defined in ACC or ISDN, and the DNS lookup is successful, then that IP address is used.
3. If the user is authenticated via the router's internal User Authentication Database, and an IP address is set in the User Authentication Database for that user, then that IP address is used.
4. If the call is an ACC call on an asynchronous port with an IP address set, then that IP address is used.
5. If the ACC or PPP call has an IP pool set, and the request to the IP pool is successful, then that IP address is used.

IP address pools are created and destroyed using the commands:

```
CREATE IP POOL=pool-name IP=ipadd[-ipadd]
DESTROY IP POOL=pool-name
```

The currently configured IP address pools, and the status of the IP addresses in the pools, can be displayed using the command:

```
SHOW IP POOL[=pool-name] [IP=ipadd[-ipadd]] [SUMMARY]
```

Once an IP address pool has been created, it must be assigned to a call so that dial-in connections using that call will use IP addresses from the IP address pool. An IP address pool can be assigned to an ACC call, a PPP interface, or a PPP template using the commands:

```
CREATE ACC CALL=call-name IPPOOL=pool-name
[other-acc-options...]
```

```

SET ACC CALL=call-name IPPOOL=pool-name
    [other-acc-options...]

CREATE PPP=ppp-interface OVER=physical-interface
    IPPOOL=pool-name [other-ppp-options...]

SET PPP=ppp-interface IPPOOL=pool-name [other-ppp-options...]

CREATE PPP TEMPLATE=template IPPOOL=pool-name
    [other-template-options...]

SET PPP TEMPLATE=template IPPOOL=pool-name
    [other-template-options...]

```

To remove an IP address pool from an ACC call, PPP interface or PPP template, specify NONE as the pool name, for example:

```
SET ACC CALL=call-name IPPOOL=NONE
```

Control Over IP Packet Fragmentation

Control over the interpretation of the “Do not fragment” bit in IP packets has been added. The default behaviour, and the normal behaviour for IP, is to obey the “Do not fragment” bit in outgoing IP packets. If an outgoing IP packet is larger than the MTU (*Maximum Transmission Unit*) of the interface, and the “Do not fragment” bit is set, the packet will be discarded.

IP interfaces that are configured with Generic Routing Encapsulations (GRE), Security Associations (SA) and/or IPsec encapsulations can potentially increase packet sizes beyond the MTU of the interface. In this situation it is necessary to ignore the “Do not fragment” bit for outgoing packets. The FRAGMENT parameter has been added to control whether or not the “Do not fragment” bit will be obeyed for outgoing IP packets that are larger than the MTU of the interface:

```

ADD IP INTERFACE=interface [FRAGMENT={YES|NO}]
    [other-ip-options...]

SET IP INTERFACE=interface [FRAGMENT={YES|NO}]
    [other-ip-options...]

```

NAT

NAT now supports a wider range of protocols. A NAT translation is added or removed using the commands:

```

ADD IP NAT IP=ipadd [MASK=ipadd] [GBLIP=ipadd]
    [GBLMASK=ipadd] [GBLPORT=port] [GBLINTERFACE=interface]
    [PORT=port]
    [PROTOCOL={protocol|ALL|EGP|GRE|ICMP|OSPF|SA|TCP|UDP}]

DELETE IP NAT IP=ipadd MASK=ipadd GBLINTERFACE=interface
    [GBLMASK=ipadd] [GBLPORT=port] [PORT=port]
    [PROTOCOL={protocol|ALL|EGP|GRE|ICMP|OSPF|SA|TCP|UDP}]

DELETE IP NAT IP=ipadd GLBIP=ipadd [MASK=ipadd]
    [GBLMASK=ipadd] [GBLPORT=port] [PORT=port]
    [PROTOCOL={protocol|ALL|EGP|GRE|ICMP|OSPF|SA|TCP|UDP}]

```

The PROTOCOL parameter specifies either the name of a known IP protocol or an IP protocol number.

Time Division Multiplexing (TDM)

Time Division Multiplexing (TDM) is a mechanism for dividing the bandwidth of a link into separate channels or time slots. The router previously supported TDM in two forms—G.703 TDM and BRI TDM. Two significant changes have been made to the router's support for TDM:

- Enhancement of the G.703 TDM functionality to provide full E1/T1 TDM.
- Support for an unstructured mode on E1 TDM links.

E1/T1 TDM

E1 TDM provides a 2.048 Mbps communications link divided into 32 slots of 64 kbps each. T1 TDM provides a 1.544 Mbps communication link divided into 24 slots of 64 kbps each and an 8 kbps channel for synchronisation and maintenance. E1 and T1 TDM were first used by telephone companies for the transport of digitised voice, but since there is no difference between digitised voice and other kinds of data E1 and T1 TDM are now also used for wide area network links.

Of the 32 time slots in an E1 TDM link, time slot 0 is usually reserved for framing. For Primary Rate ISDN running over E1 where calls need to be set up and cleared dynamically, time slot 16 is reserved for signalling. This leaves up to 30 time slots, each of 64 kbps, for information transfer over ISDN calls and up to 31 time slots for static TDM links. A mixture of static links and dynamic ISDN calls is also possible. An E1 TDM link may also be used in an unstructured mode where all of the 2.048 Mbps of bandwidth is available for data transfer. In this mode only a single link of the full bandwidth is possible and slot 0 is not used to demarcate the slot structure.

For T1 Primary Rate ISDN, time slot 24 is reserved for signalling and the other 23 time slots are available for ISDN calls. When the T1 interface is not used for ISDN, slot 24 may be used for a TDM link. Unstructured mode is not supported for T1 TDM.

E1/T1 TDM support is provided on the router by the AR020 E1/T1 PIC, and can be used as the data link layer transport mechanism for Primary Rate ISDN or as the data channel over which one or more static PPP links can be configured.

Unstructured Mode

E1 TDM links now support an unstructured mode in which the TDM group uses all the available bandwidth on the E1 link. Unstructured mode is enabled when the TDM group is created, using the command:

```
CREATE TDM GROUP=groupname INTERFACE=interface
        {SLOTS=slotlist | UNSTRUCTURED}
```

The SLOTS and UNSTRUCTURED parameters are mutually exclusive and may not be specified together in the same command. The E1 PRI interface must be set to TDM mode, not MIXED mode. Unstructured mode is not supported on a T1 or BRI link. An unstructured group may be the only group on an interface as it effectively uses all the time slots.

For example, to create a TDM group named "bigpipe" that uses all the bandwidth of PRI interface 0, use the command:

```
CREATE TDM GROUP=bigpipe INTERFACE=PRI0 UNSTRUCTURED
```

PBX

Software Release 1.8 adds Centrex™ services to the range of call handling functions provided by the PBX module.

A Centrex™ system allows a business to have offices spread throughout a city or country, yet be able to access any phone connected to that system using simple 3 to 5 digit extension numbers. Standard PABX functions such as call hold, call divert, and call transfer are all possible with the Centrex™ system, as well as other more sophisticated features. These features are in addition to the PABX features on the router, and are duplications in some instances.

A Centrex™ group is a collection of phone lines connected to the Centrex™ system. They each have assigned Centrex™ extensions as well as unique Direct-Dial-In (DDI) numbers. Only extensions in the same Centrex™ group can call one another using the assigned Centrex™ extension numbers. All other calls must be made using the DDI number for that line.

On a single Basic Rate ISDN line with two B channels, both B channels will be tied up if an external call is made to a local user and then transferred to another external number. In this situation, the incoming call comes in on one channel and is put on hold at the router. When the user transfers the call, another call is made on the other B channel to the second external number. When the call is transferred both B channels are then in use.

Centrex™ call transfer enables the incoming external call to be put on hold at the Centrex™ switch, rather than at the local router. The same B channel can then be used to make the call to the second external number. When the calls are transferred, they are actually connected at the switch, freeing up both B channels at the local router.

NOTE: The Centrex™ feature is currently only available in New Zealand. The router territory must be set to New Zealand using the SET SYSTEM TERRITORY command.

The Centrex™ feature is disabled by default. Two new commands have been added to enable or disable the Centrex™ feature:

```
ENABLE PBX CENTREX
DISABLE PBX CENTREX
```

The *Centrex* field has been added to the SHOW PBX command to display the current state of the Centrex feature.

A new parameter, CALLINGNUMBER, has been added to the CREATE PBX EXTENSION and SET PBX EXTENSION commands:

```
CREATE PBX EXTENSION=extension-number
    [CALLINGNUMBER={calling-number|OFF}] [other-options...]
SET PBX EXTENSION=extension-number [CALLINGNUMBER={calling-
number|OFF}] [other-options...]
```

The CALLINGNUMBER parameter specifies the format of the calling party number IE (also known as CLI) in outgoing call SETUP messages. The STD number is not included. If OFF is specified, the calling party number IE is not included in the call SETUP message. The default is OFF.

The *Calling number* field has been added to the SHOW PBX EXTENSION command to display the value of the CALLINGNUMBER parameter.

Layer 2 Tunnelling Protocol (L2TP)

The router's implementation of the Layer Two Tunnelling Protocol (L2TP) has been enhanced to be compatible with the Internet Draft *Layer Two Tunnelling "L2TP"*, revision 16, June 1999.

Draft 16 differs significantly from, and is incompatible with, drafts prior to Draft 13. Two new command parameters, NUMBER and PRE13 have been added to the following commands to provide compatibility with pre-Draft 13 implementations:

```

ADD L2TP CALL=name TYPE={ASYNC|ISDN|VIRTUAL} IP=ipadd
    REMOTE=name [DIAL=number] [NUMBER={ON|OFF|STARTUP}]
    [PASSWORD=password] [PRE13={ON|OFF}] [PRECEDENCE={IN|OUT}]
    [SPEED=speed] [SUBADDRESS=subaddress]

SET L2TP CALL=name [TYPE={ASYNC|ISDN|VIRTUAL}] [IP=ipadd]
    [REMOTE=name] [DIAL=number] [NUMBER={ON|OFF|STARTUP}]
    [PASSWORD=password] [PRE13={ON|OFF}] [PRECEDENCE={IN|OUT}]
    [SPEED=speed] [SUBADDRESS=subaddress]

ADD L2TP IP=ipadd[-ipadd] PPPTEMPLATE=ppp-template
    [NUMBER={ON|OFF|STARTUP}] [PRE13={ON|OFF}]

ADD L2TP USER={mapping|ALL|LOCAL|NONE|REMOTE}
    ACTION={DATABASE|DNSLOOKUP|IGNORE|RADIUS} [IP=ipadd]
    [PORT=port] [NUMBER={ON|OFF}] [PASSWORD=password]
    [PRE13={ON|OFF}] [PREFIX=prefix] [TIMEOUT=timeout]

SET L2TP USER={mapping|ALL|LOCAL|NONE|REMOTE}
    [ACTION={DATABASE|DNSLOOKUP|IGNORE|RADIUS}] [IP=ipadd]
    [PORT=port] [NUMBER={ON|OFF}] [PASSWORD=password]
    [PRE13={ON|OFF}] [PREFIX=prefix] [TIMEOUT=timeout]

```

The NUMBER parameter specifies how L2TP will handle the sequence numbering of L2TP data packets. If NUMBER is set to ON, data packets will always be numbered. If NUMBER is set to OFF, sequence numbering will only be used if the remote end requests sequencing. L2TP data packet sequence numbering is used to ensure that packets arrive in the order in which they were sent from the remote L2TP server. This is required if encryption is active over the L2TP link. It is also useful in the initial PPP negotiation phase if the link between the L2TP servers uses multiple routes.

The PRE13 parameter specifies compatibility with pre-Internet Draft 13 L2TP implementations.

The SHOW L2TP IP, SHOW L2TP CALL and SHOW L2TP USER commands have been modified to include two new fields, *Sequence numbering* and *Pre-draft 13 support*, reflecting the current state of the new NUMBER and PRE13 parameters.

Debugging can now be enabled or disabled on a per-call or per-tunnel basis, using the commands:

```

ENABLE L2TP DEBUG={ALL|PKT|STATE}
    {CALL[=callid]|TUNNEL[=tunnelid]}

DISABLE L2TP DEBUG={ALL|PKT|STATE}
    {CALL[=callid]|TUNNEL[=tunnelid]}

```

If a call or tunnel is specified, debugging is enabled or disabled for that call or tunnel. If a call or tunnel is not specified, debugging is enabled or disabled for all calls or tunnels.

The SET L2TP WINDOW command has been removed.

The SHOW L2TP TUNNEL command now accepts an optional call ID or tunnel ID:

```
SHOW L2TP TUNNEL[=tunnelId] [CALL[=callId]] [COUNTER]
```

If a call ID is specified, detailed information about the call is displayed. If a call ID is not specified, summary information about all current calls is displayed. If a tunnel ID is specified, detailed information about the tunnel is displayed. If a tunnel ID is not specified, summary information about all tunnels is displayed.

IP Security (IPsec)

Software Release 1.8 includes an RFC-compliant implementation of the *Internet Protocol Security Facility* (IPsec), the *Internet Security Association Key Management Protocol* (ISAKMP) and the *Internet Key Exchange* (IKE) protocol. IPsec is a set of security protocols which provide encryption and authentication to IP packets. IKE is a mechanism for negotiating IPsec protection and keys. The basic framework used by IKE is provided by ISAKMP.

IPsec provides the following security services for traffic at the IP layer:

- Connectionless integrity—ensuring the data has not been changed en-route.
- Data origin authentication—identifying who sent the data.
- Confidentiality (encryption)—ensuring that the data has not been read en-route.
- Replay protection—detecting packets received more than once to help protect against denial of service attacks.

These services are provided through the use of two security protocols, the *IP Authentication Header* (AH) and the *IP Encapsulating Security Payload* (ESP). The IP Authentication Header (AH) protocol provides authentication of IP packets. AH provides protection to an IP packet and to any further headers added by AH. The Encapsulating Security Payload (ESP) provides one or both of encryption and authentication. ESP provides protection to an IP packet, but not to any further headers added by ESP. AH and ESP are algorithm-independent and may be applied either alone or together to provide the desired set of security services for selected IP packets.

IPsec uses ENCO services for encryption, authentication and compression algorithms. If the ENCO module is not configured to provide the resources required by IPsec, the IPsec configuration will not work as intended.

IPsec can be enabled or disabled using the commands:

```
ENABLE IPSEC  
DISABLE IPSEC
```

The status of IPsec can be displayed using the command:

```
SHOW IPSEC
```

The router's implementation of IPsec consists of two parts:

- The IPsec configuration, which is stored in the IPsec Security Policy Database (SPD).
- The currently active Security Associations and SA bundles. These are created from the IPsec configuration, either manually or by the ISAKMP/IKE key management protocol.

NOTE: When the router is in Security Mode only users with Security Officer privilege may to change the IPsec configuration.

Security Associations

A Security Association (SA) is a simplex connection that provides security services between IPsec peers for selected IP packets. Typically, two SAs are created for a traffic stream, one to protect inbound traffic and one to protect outbound traffic. Security services are provided by an SA using AH, or ESP, but not both. If both AH and ESP protection is required for a traffic stream, a *bundle* of two or more SAs is created and applied to the traffic stream. An SA is uniquely identified by the combination of a random number called the *Security Parameter Index* (SPI), an IP destination address, and a security protocol (either AH or ESP).

Security Policy Database

Three entities are used to define the IPsec configuration in the IPsec Security Policy Database—SA specifications, bundle specifications, and policies.

SA Specifications

SA Specifications are a template for SAs. They specify the attributes an SA will have when it is created. The router creates SAs in pairs, one for outbound packets and one for inbound packets. When an SA pair is created an SA specification is used to determine the attributes of the SAs, including the key management mechanism to be used to create the SAs, the encryption and/or authentication algorithms to be used by the SAs, and how the SAs will provide an anti-replay service (if any).

SA specifications are managed using the commands:

```
CREATE IPSEC SASPECIFICATION=spec-id
  KEYMANAGEMENT={ ISAKMP | MANUAL } PROTOCOL={ AH | COMP | ESP }
  [ ANTIREPLAYENABLED={ TRUE | FALSE } ] [ COMPALGORITHM=LZS ]
  [ ENCALGORITHM={ 3DES2KEY | 3DESINNER | DES | NULL } ]
  [ ENCKEY=key-id ] [ HASHALGORITHM={ DESMAC | MD5 | NULL | SHA } ]
  [ HASHKEY=key-id ] [ INSPI=spi ] [ MODE={ TRANSPORT | TUNNEL } ]
  [ OUTSPI=spi ] [ REPLAYWINDOWSIZE={ 32 | 64 | 128 | 256 } ]

DESTROY IPSEC SASPECIFICATION=spec-id

SET IPSEC SASPECIFICATION=spec-id
  [ ANTIREPLAYENABLED={ TRUE | FALSE } ] [ COMPALGORITHM=LZS ]
  [ ENCALGORITHM={ 3DES2KEY | 3DESINNER | DES | NULL } ]
  [ ENCKEY=key-id ] [ HASHALGORITHM={ DESMAC | MD5 | NULL | SHA } ]
  [ HASHKEY=key-id ] [ INSPI=spi ] [ MODE={ TRANSPORT | TUNNEL } ]
  [ OUTSPI=spi ] [ REPLAYWINDOWSIZE={ 32 | 64 | 128 | 256 } ]

SHOW IPSEC SASPECIFICATION[=spec-id]
```

Bundle Specifications

The router always creates an SA pair as part of an SA bundle. Often, however, there will be only one SA pair in a bundle. Bundle Specifications are a template for SA bundles. They specify the number and order of SAs an SA bundle will have when it is created. A bundle specification used by ISAKMP/IKE to negotiate an SA bundle can specify more than one bundle of SAs, in order, with the most desired bundle first. ISAKMP/IKE will negotiate with the IPsec peer to determine which bundle will be used.

When an SA bundle is created a bundle specification is used to determine the attributes of the SA Bundle. The attributes of a bundle specification are the key management mechanism to be used to create the bundle, the lifetime expiry limits of the SA pairs to be created for the bundle, and the SA pairs to be created for the bundle. The SA pairs are represented by SA specification identification numbers in a *bundle string*, which can also contain the connectors "AND", "OR" and ",". Examples of valid bundle strings are:

```
"1 OR 2 AND 3, 1 OR 2"
"1 AND 3, 1, 2 AND 3, 2"
"1, 2 AND 3"
"1, 2"
```

Bundle specifications are managed using the commands:

```
CREATE IPSEC BUNDLESPECIFICATION=bundle-spec-id
  KEYMANAGEMENT={ISAKMP|MANUAL} STRING="bundle-string"
  [EXPIRYKBYTES=kbytes] [EXPIRYSECONDS=seconds]

DESTROY IPSEC BUNDLESPECIFICATION=bundle-spec-id

SET IPSEC BUNDLESPECIFICATION=bundle-spec-id
  [EXPIRYKBYTES=kbytes] [EXPIRYSECONDS=seconds]

SHOW IPSEC BUNDLESPECIFICATION[=bundle-spec-id]
```

Policies

An IPsec policy binds a packet selection rule to an action. A policy may be attached to only one IP logical interface, but multiple IPsec policies may be attached to the same IP logical interface. When multiple policies are attached to one IP logical interface, the policies are ordered and packets traversing the interface are matched against the policies' selection rules in order. Policies with more specific selection rules must be placed ahead of those with more general rules. IPsec policies will only be attached to IP logical interfaces, and receive packets for processing, when the IPsec module is enabled.

In general, the action used to process each IP packet is determined by matching information from the IP and transport layer headers of the IP packet against IPsec policies in the SPD. The pieces of information used to select the packet are termed *selectors*. IPsec supports the following selectors:

- Remote IP address—a single address, a range of addresses or a masked subnet
- Local IP address—a single address, a range of addresses or a masked subnet
- Remote port
- Local port
- Transport protocol (e.g. UDP, TCP, ICMP)

- Local system name or user name
- Remote system name or user name

If a selection value is not assigned to a selector, the default is to match any value. An IP packet must have fields which match all the selectors of a policy to be selected by that policy.

The actions an IPsec policy can define are:

- PERMIT—allow matching packets to bypass IPsec and be processed normally by the router.
- DENY—discard any matching packets.
- IPSEC—apply IPsec processing to any matching packets.

If the action is IPSEC the policy must also specify:

- The IP address of the IPsec peer, or “DYNAMIC” if the peer’s IP address is dynamically assigned.
- The key management method to be used for creating an SA bundle to process the selected packets (either ISAKMP or MANUAL).
- The bundle specification to be used when creating an SA bundle to process the selected packets.

If the key management is ISAKMP the policy may also specify:

- That ISAKMP/IKE must use Perfect Forward Secrecy when creating keys for the bundle.
- An Oakley group to be used by ISAKMP/IKE when negotiating the bundle.

IPsec policies are managed using the commands:

```
CREATE IPSEC POLICY=name INTERFACE=interface
ACTION={DENY|IPSEC|PERMIT}
[BUNDLESPECIFICATION=bundlespec-id] [GROUP={1|2}]
[KEYMANAGEMENT={ISAKMP|MANUAL}]
[LADDRESS={ANY|ipadd[-ipadd]}] [LMASK=ipadd]
[LNAME={ANY|system-name}] [LPORT={ANY|OPAQUE|port}]
[PEERADDRESS={ipadd|DYNAMIC}] [POSITION=pos]
[RADDRESS={ANY|ipadd[-ipadd]}] [RMASK=ipadd]
[RNAME={ANY|system-name}] [RPORT={ANY|port|OPAQUE}]
[SASELECTORFROMPKT={ALL|LADDRESS|LPORT|NONE|RADDRESS|RPORT|
TRANSPORTPROTOCOL}]
[TRANSPORTPROTOCOL={ANY|EGP|ESP|GRE|ICMP|OPAQUE|OSPF|RSVP|
TCP|UDP|protocol}] [USEPFSKEY={TRUE|FALSE}]

DESTROY IPSEC POLICY=name [SABUNDLE=bundle-id]

SET IPSEC POLICY=name [INTERFACE=interface]
[ACTION={DENY|IPSEC|PERMIT}]
[BUNDLESPECIFICATION=bundlespec-id] [GROUP={1|2}]
[LADDRESS={ANY|ipadd[-ipadd]}] [LMASK=ipadd]
[LNAME={ANY|system-name}] [LPORT={ANY|OPAQUE|port}]
[PEERADDRESS={ipadd|DYNAMIC}] [POSITION=pos]
[RADDRESS={ANY|ipadd[-ipadd]}] [RMASK=ipadd]
[RNAME={ANY|system-name}] [RPORT={ANY|port|OPAQUE}]
[TRANSPORTPROTOCOL={ANY|EGP|ESP|GRE|ICMP|OPAQUE|OSPF|RSVP|
TCP|UDP|protocol}] [USEPFSKEY={TRUE|FALSE}]

SHOW IPSEC POLICY[=name] [SABUNDLE] [COUNTER]
```

ISAKMP/IKE Key Management

The *Internet Security Association and Key Management Protocol* (ISAKMP) is defined in RFC 2408. It prescribes procedures and packet formats to securely establish, negotiate, modify and delete Security Associations (SAs) and their attributes, including secret keys. It provides a common framework for a key management implementation to follow and is independent of key generation, authentication, encryption algorithms and SA definitions. The *Internet Key Exchange* (IKE) protocol, defined in RFC 2409, is based on the ISAKMP framework and provides an authenticated key exchange method using the Diffie-Hellman algorithm.

Key management refers to the creation, distribution, storage and deletion of keys. If an attacker obtains an encryption key they can decrypt sessions they recorded any length of time ago. Therefore a key management protocol must ensure keys are never compromised.

Session keys are used to increase the security of encryption keys. If the key is changed frequently it is near impossible for an attacker to try a brute force attack on the encrypted data. If a key is compromised in the future, only the data encrypted with that particular session key can be decrypted. Session keys need to be changed on a regular basis, so manually keyed IPsec soon becomes very unmanageable. Session keys, therefore, need to be changed automatically, and more importantly, securely.

NOTE: Software Release 1.8 supports both manual key management and ISAKMP/IKE key management.

ISAKMP/IKE ensures that only the two negotiating parties know the value of the exchanged session key. If any key is compromised then it should not cause any other keys to be compromised. This property is called *Perfect Forward Secrecy* (PFS). Other properties such as replay protection ensure that the key management protocol is not susceptible to denial of service and man-in-the-middle attacks.

When an IPsec policy is configured to use ISAKMP/IKE key management, IPsec will request that a new SA bundle be negotiated when the IP module passes it a packet for a policy which either does not have an SA bundle or does not have an SA bundle with selectors matching the packet. Outbound packets will be queued while IPsec is waiting for ISAKMP/IKE to negotiate a bundle. SA bundles will also be created if an IPsec peer requests that ISAKMP/IKE negotiates a bundle.

SA pairs negotiated by ISAKMP/IKE have lifetime expiry limits, specified in seconds and/or kilobytes of data processed. When an SA pair in an SA bundle nears its lifetime, IPsec will request that ISAKMP/IKE negotiate a replacement bundle. The replacement bundle will be used in preference to the old bundle as soon as it is negotiated.

ISAKMP negotiates security associations on behalf of IPsec. If an IPsec policy specifies ISAKMP for key management, then ISAKMP must be enabled and an ISAKMP policy must exist for the ISAKMP peer.

NOTE: An ISAKMP feature licence must be installed before ISAKMP can be enabled. Modification of the router's ISAKMP configuration requires Security Officer privilege. ISAKMP also requires DES encryption and the Diffie-Hellman key exchange algorithm.

ISAKMP can be enabled or disabled using the commands:

```
ENABLE ISAKMP [LOCALRSAKEY=key-id]
DISABLE ISAKMP
```

The status of ISAKMP can be displayed using the command:

```
SHOW ISAKMP
```

An ISAKMP policy specifies how to communicate with, and how to authenticate, an ISAKMP peer. The information in the ISAKMP policy is used during ISAKMP exchanges and in the creation of an ISAKMP SA. An ISAKMP policy specifies:

- An encryption algorithm and hash algorithm. The encryption algorithm is used to encrypt ISAKMP messages to protect them against eavesdropping. The hash algorithm is used to authenticate ISAKMP messages to prevent man-in-the-middle attacks.
- The address of the remote ISAKMP peer. This is used to match an ISAKMP policy to an IPsec peer.
- The method used to authenticate an ISAKMP peer—a shared secret key or an RSA public key.
- Optionally, a different Diffie-Hellman group to the default “OAKLEY MODP Group 1”. The more secure group “OAKLEY MODP Group 2” may be used at the expense of negotiation speed.

ISAKMP policies are managed using the commands:

```
CREATE ISAKMP POLICY=name PEER=ipadd
[AUTHTYPE={PRESHARED|RSAENCR}]
[ENCALG={3DES2KEY|3DESINNER|DES}] [EXPIRYKBYTES=kbytes]
[EXPIRYSECONDS=seconds] [GROUP={1|2}] [HASHALG={MD5|SHA}]
[KEY=key-id] [PRENEGOTIATE={TRUE|FALSE}]

DESTROY ISAKMP POLICY=name

SHOW ISAKMP POLICY[=name]
```

Communication between two ISAKMP peers takes the form of an ISAKMP exchange. An ISAKMP exchange consists of a finite number of messages and is dynamically created and destroyed as required. The details of any currently active ISAKMP exchanges can be displayed using the command:

```
SHOW ISAKMP EXCHANGE[=exchange-id]
```

The ISAKMP SA is used to protect ISAKMP traffic between the local router and the remote ISAKMP peer. A single ISAKMP SA is dynamically created for each ISAKMP peer. The algorithms used to encrypt and authenticate messages from the remote ISAKMP peer, along with all keys required for the algorithms are stored in the ISAKMP SA. ISAKMP uses information stored in the ISAKMP SA and ENCO encryption services to protect traffic between ISAKMP peers. ISAKMP also uses the ENCO Diffie-Hellman key exchange algorithm to create keys for the ISAKMP SA. The details of existing ISAKMP SAs can be displayed using the command:

```
SHOW ISAKMP SA[=sa-id]
```

Pre-IPsec Security Associations

For backward compatibility, the router supports the Security Association implementation based on 1998 Internet Drafts, as implemented in Software Release 7.4.

The software allows for mixed networks of old Security Association implementations and IPsec. An IP logical interface may have both IPsec policies and old SAs attached to it. This allows a router to communicate using old SAs with one set of routers and to communicate using IPsec with another set of routers. To enable backward compatibility, use the command:

```
ENABLE IPSEC OLDSA INTERFACE=interface
```

To upgrade a router from the old SA implementation to IPsec, use the command:

```
ACTIVATE IPSEC CONVERTOLDSA [SA=sa-id]
```

to convert the old SA configuration in memory into a functionally equivalent IPsec configuration in memory. The old SA configuration will be removed from memory. After the conversion process the IPsec policies created will not have valid IPsec peer addresses. These need to be entered manually before the new IPsec configuration can be used. This conversion process does not change the router configuration file in any way. To retain the new IPsec configuration over a router restart, the configuration file must be updated using the command:

```
CREATE CONFIG=filename
```

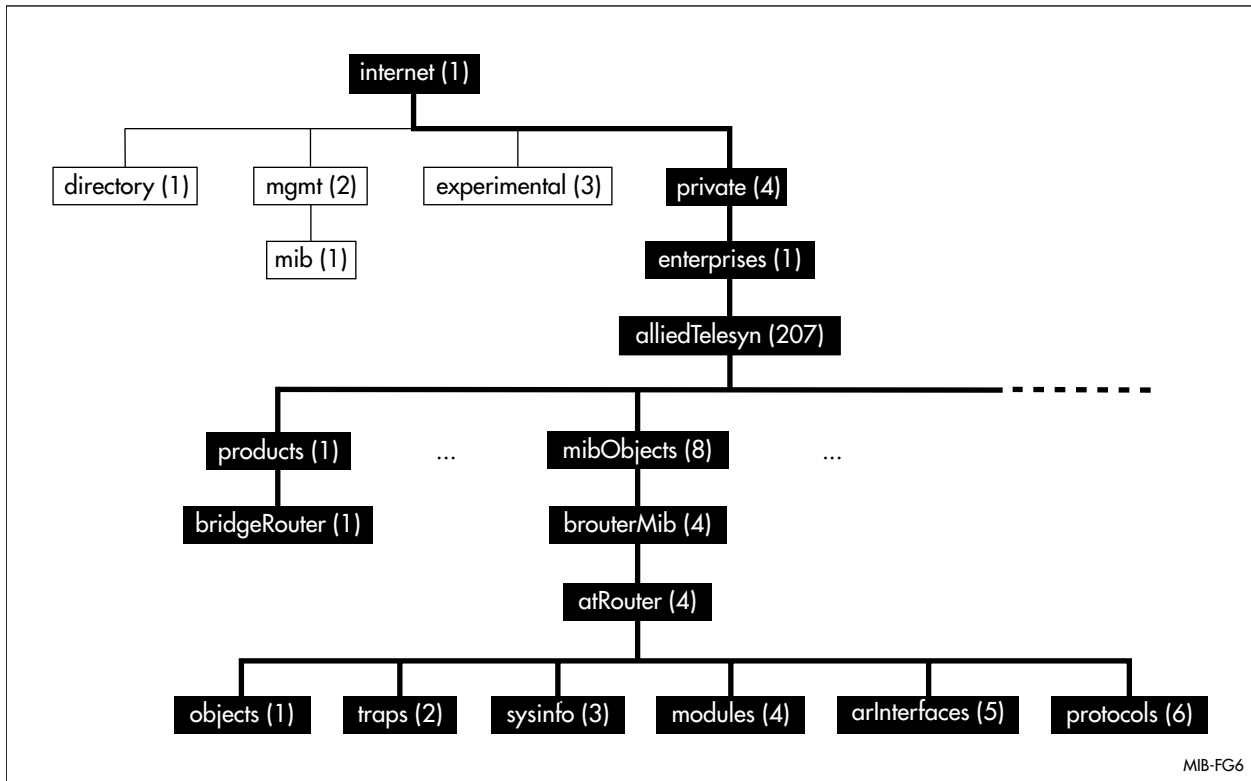
or by editing the current boot configuration script, using the command:

```
EDIT filename
```

MIBs

The router's Enterprise MIB is now incorporated in the Allied Telesyn Enterprise MIB (Figure 5 on page 24). The Products sub-tree (`{ enterprises (1) alliedTelesyn(207) products(1) }`) contains a set of object identifiers for Allied Telesyn products. Within this sub-tree, objects with the object identifier prefix `bridgeRouter` (`{ products 1 }`) are specific to the AR Router family.

Figure 5: The Allied Telesyn Enterprise MIB sub-tree of the Internet-standard Management Information Base (MIB).



The AT Router sub-tree contains a set of objects for managing the AR Router family of multiprotocol routers. Objects are arranged into six groups:

- The Objects Group contains four sets of object identifiers for boards, releases, interface types and chip sets. These object identifiers are for use with the `hrDeviceID` object in the Host Resources MIB.
- The arInterfaces Group contains objects that describe the boards, slots and physical interfaces in the router. A router consists of a number of “boards”. Each board may have a number of “positions”, each of which contains a single physical interface. Each board may also have a number of “slots”, which are places which can take other boards. Thus the physical construction of a router may be seen as a tree whose nodes are boards and interfaces, and whose links are positions and slots.
- The Modules Group contains objects that describe particular software modules in the router that are not covered by standard MIBs. The Modules Group includes the following object groups:
 - The Ethernet Group has the object identifier prefix `ethernet` (`{ modules 23 }`), and contains the following objects that describe Ethernet interface(s) on the router:

- The ISDN Call Control Group contains objects that describe ISDN call definitions, active call details and call history on the router, and has the object identifier prefix *cc* ({ modules 37 }).
- The BRI Group contains objects that describe BRI interfaces on the router, and has the object identifier prefix *bri* ({ modules 41 }).
- The PRI Group contains objects that describe PRI interfaces on the router, and has the object identifier prefix *pri* ({ modules 42 }).
- The Loader Group contains objects for managing the LOAD module which uses TFTP to download releases, patches, configuration scripts and other files from a TFTP server to NVS or FLASH storage in the router. Objects in this group have the object identifier prefix *loader* ({ modules 48 }).
- The Install Group contains objects for managing the INSTALL module which controls the software release and patch running on the router. Objects in this group have the object identifier prefix *install* ({ modules 49 }).
- The File Group contains objects for managing the file system in the router. Objects in this group have the object identifier prefix *file* ({ modules 56 }).
- The Firewall Group contains objects describing traps generated by the firewall in the router. Objects in this group have the object identifier prefix *firewall* ({ modules 77 }).

Other Enhancements

- The RADIUS attribute *Session-Timeout* (27) is now supported by the router. The *Session-Timeout* attribute is only applied to PPP sessions and not directly to logged in users.

- The RESTART ROUTER command now accepts the keyword NONE as a configuration file name:

```
RESTART ROUTER [CONFIG={filename|NONE}]
```

The value NONE forces the router to reboot without executing a configuration file. This overrides any default configuration file set with the SET CONFIG command.

- The SHOW BUFFER SCAN command now displays the number of buffers for each source using a wider output field to avoid overflow errors on routers with 16MB of memory Figure 6 on page 26.

Figure 6: Example output from the SHOW BUFFER SCAN command.

```
Scan of buffers in use

00093d62    2 001338a2    1 0013d27c    1 000cd26a    1 000ccfc2    7
000cd326    5 000cd542    1 0006d1f0    1 000a03e4    1 000a4256    1
001f544e    1 001f5484    1 001f54c0    1 000a50da    1 00082e52    1
0013fe40    2 0008c8b0    1 0008c8f0    1 0008c92c    1 0008f7f6    1
000ebd32    1 000ec0a2    2 000ec364    3 00080048    8 00081352    1
0016ef96    1 0012fd76    1 0012f64a    1 00086e3c    1 0008871a    1
000b6866    1 001f5338   10 001526e0    1 0011e892    2 00099486    1
001194d4    1 0011deb0   17 0011fd6a    2 0011d278    1 001139a4    1
0011b354    1 0011d7e8    1 001fe0ca    1 001fb446    1 001fb48c    2
001fb4e8    2 001fb52a    1 0005e95c    1 0005e9f8    1 000d3976    1
00161596    1 00153b60    1 000994ae    1 000d133e    1 000bbc3a    1
00163154    1 001069fc    1 000a4916    1 000a5298    1 00141e26    1
00157156    1 000f4028    1 00169bd8    1 000a9654    1 001352a4   16
000892ae    1 001524fa    1 00087014    1 00089666    1 0008625c    1
0012f6d2    1 00141e30    1 00141e3a    1 0014190e    1 00141940    1
000c512a   15 00087624    1

Total buffers in use - 333

Memory ( DRAM ) ..... 8192 kB
Free Memory ..... 61 %
Free buffers ..... 2860
Total buffers ..... 3193
Buffer level 3 ..... 638 (don't process input frames)
Buffer level 2 ..... 1277 (don't do monitor or command output)
Buffer level 1 ..... 2235 (don't buffer up log messages)
```

- The SHOW SYSTEM command has been modified to reflect the new model naming scheme and other product enhancements (Figure 7 on page 27). The *Nick Name* and *Part Name* fields have been replaced by a single *Board Name* field, and the *sysDescription* field now contains a complete description of the router model and software release.

Figure 7: Example output from the SHOW SYSTEM command.

```

Router System Status                               Time 17:10:06 Date 25-Sep-1999.
Board      ID  Bay Board Name                               Rev    Serial number
-----
Base       62      AR720                               M1-0   6845218
IC Module  40    0  AR022 PIC Eth                               M2-0   6844595
IC Module  38    1  AR023 PIC Sync                             M1-1   6844715
MAC        67      AR012 CMAC                             M2-0   33636409
-----
Memory -   DRAM : 16384 kB   FLASH : 4096 kB
-----
SysDescription
CentreCOM AR720 version 1.8.1-00 08-Sep-1999
SysContact
David Johns, ext 8331
SysLocation
Laboratory, First Floor, Head Office Building
SysName
LAB
SysUpTime
250074 ( 00:41:40 )
Software Version: 1.8.1-00 08-Sep-1999
Release Version : 1.8.1-00 08-Sep-1999
Patch Installed : NONE
Territory      : europe
Help File      : help.hlp

Boot configuration file: load.cfg (exists)
Current configuration: load.cfg
Security Mode   : Disabled

Patch files
Name           Device      Size      Version
-----
52772-02.paz   flash      94856     7.7.2-2
-----

```

- The SHOW PORT command has been modified to display additional debugging information (Figure 8 on page 28). Three fields—*PPP Index*, *TX ACCM* and *Transmit frame*—have been added (Table 4 on page 29).

Figure 8: Example output from the SHOW PORT command.

```

PORT 2 : 0000070953 seconds   Last change at: 0000009023 seconds

PORT information
Name ..... Port 2
Status ..... enabled
Mode ..... PPP
PPP Index ..... 1
TX ACCM ..... 00000000
Data rate ..... 38400
Parity ..... none
Data bits ..... 8
Stop bits ..... 1
Test mode ..... no
In flow state (mode) ..... on (Hardware)
Out flow state (mode) ..... off (Hardware)
Autobaud mode ..... disabled
Max tx queue length ..... 100
TX queue length ..... 0
Transmit frame ..... none
RX queue length ..... 0
IP address ..... none
Max transmission unit ..... 1500
IPX Network ..... none

Control signals
  DTR (out) ..... on on      1
  RTS (out) ..... on -      1
  CD (in) ..... off connect 0
  CTS (in) ..... off -      0
  RNG (in) ..... off -      0

TTY information
Instance ..... 18
Login Name .....
Description ..... Port 2
Secure ..... yes
Connections to .....
Current connection ..... none
In flow state ..... on
Out flow state ..... on
Attached module ..... ASYN Call Control
Attached module instance .. 2
Type ..... VT100
Service ..... none
Prompt ..... login
Echo ..... yes
Attention ..... break
Manager ..... no
Edit mode ..... insert
History length ..... 20
Page size ..... 22

```

Table 4: New fields in the output of the SHOW PORT command.

Parameter	Meaning
PPP Index	The index for the current PPP session. This field is only displayed if the port is being used by ACC for a PPP session.
TX ACCM	The current ACCM used by PPP for transmitted control characters. This field is only displayed if the port is being used by ACC for a PPP session.
Transmit frame	The address of the current frame being transmitted by the port, or "none" if no frame is currently being transmitted.

- The SHOW FRAMERELAY command has been modified to include the status of the interface (Figure 9 on page 29). One field, *Active*, has been added (Table 5 on page 29).

Figure 9: Example output from the SHOW FRAMERELAY command.

Interface	Enabled	IfIndex	Over	Active	Logical Interfaces
fr0	YES	05	syn0	YES	0,1,2,3,7
fr1	YES	06	syn1	YES	0,2,3
fr2	YES	07	syn2	YES	0

Table 5: New fields in the output of the SHOW FRAMERELAY command.

Parameter	Meaning
Active	Whether or not the Frame Relay interface is active; one of "YES" or "NO".

- The SHOW FRAMERELAY LI command now displays *IfIndex* in a wider field to support values over 100 (Figure 10 on page 29).

Figure 10: Example output from the SHOW FRAMERELAY LI command.

Interface	LI	IfIndex	Type	Number of DLCs	User Modules
fr0	All				IPX
	0	06	NBMA	4	IP
	1	07	PTP	1	ARP
	2	08	PTP	1	IP
	3	10	PTP	1	ARP
	7	11	NBMA	3	IP
					ARP
fr1	0	09	NBMA	3	IP
	2	12	PTP	1	ARP
	3	13	NBMA	1	IP
				ARP	

- The SHOW PRI COUNTERS=DIAGNOSTIC command now displays a different output for a T1-ESF interface operating in message-oriented mode (Figure 11 on page 30, Table 6 on page 30).

Figure 11: Example output from the SHOW PRI COUNTERS=DIAGNOSTIC command for a T1-ESF interface operating in message-oriented mode.

```

PRI instance 0:          1544 seconds          Last change at:          1499 seconds

Interface-global Diagnostic Counters

  Device Independent Diagnostic Counters

EventQueueFulls                0

Data Link counters

  Receive:                      Transmit:
Bytes                          10439          Bytes                    260
Messages                       19            Messages                 20
FIFOOverflows                  0            FIFOUnderruns           0
UnderlengthMessages            0            DiscardedMessages       0
OverlengthMessages             0
UnrecognisedMessages           0
Aborts                         0
Errors                         0

```

Table 6: Parameters displayed in the output of SHOW PRI COUNTERS=DIAGNOSTIC command for a T1-ESF interface operating in message-oriented mode.

Parameter	Meaning
PRI instance	The instance number of the PRI interface.
seconds	The current value of sysUpTime.
Last change at	The value of sysUpTime at the time the interface entered its current operational state.
EventQueueFulls	The number of times the queue of events for the layer 1 state machine has become full. This could occur if very many events are occurring in a short time period due to an unstable link.
Bytes	The number of bytes received/transmitted over the T1-ESF Data Link.
Messages	The number of messages received/transmitted over the T1-ESF Data Link.
FIFOOverflows	The number of messages lost due to a receiver overrun, possibly due to router overload.
UnderlengthMessages	The number of received messages discarded because they were less than 6 bytes in length, possibly an artifact of changing from bit-patterned to message oriented mode.
OverlengthMessages	The number of received messages discarded because they were longer than a receive buffer, possibly an artifact of changing from bit-patterned to message oriented mode.
UnrecognisedMessages	The number of received messages discarded because they were neither a T1.403 Performance Report Message nor a AT&T 54016 request.

Table 6: Parameters displayed in the output of SHOW PRI COUNTERS=DIAGNOSTIC command for a T1-ESF interface operating in message-oriented mode. (Continued)

Parameter	Meaning
Aborts	The number of received messages terminated in an abort, possibly due to noise on the line.
Errors	The number of received messages discarded due to a FCS error or because they were not an integral number of octets, possibly due to noise on the line.
FIFOUnderruns	The number of times a message had to be retransmitted due to a transmitter underrun, possibly due to router overload.
DiscardedMessages	The number of messages waiting to be transmitted that were discarded when the interface was reset.

- The SHOW IP INTERFACE command now displays information about IPsec policies attached to IP (Figure 12 on page 31, Table 7 on page 31).

Figure 12: Example output from the SHOW IP INTERFACE command.

Interface Pri. Filt	Type Pol. Filt	IP Address Network Mask	Bc Fr MTU	PArp VJC	Filt GRE	RIP Met. OSPF Met.	SAMode DBcast	IPSc Mul.
LOCAL	-	Not Set	- n	-	---	-	-	--
---	----	-	-	-	---	-	-	---
eth0	Static	192.168.163.39	1 y	On	---	01	Pass	--
---	---	255.255.255.0	1500	-	---	0000000001	No	On
ppp1	Dynamic	0.0.0.0	1 y	-	---	01	Pass	--
---	---	255.255.255.255	1500	Off	---	0000000001	No	On
ppp2	Inactive	192.168.23.3	1 n	-	---	01	Pass	--
---	---	255.255.255.0	1500	Off	---	0000000001	Yes	Off

Table 7: New fields in the output of the SHOW IP INTERFACE command.

Parameter	Meaning
IPSc	Whether or not IPsec policies are attached to the interface; one of "Yes" or "No".

- The SHOW SERVICE command syntax has been modified. If a service name is specified it must be the complete name of a service and not just the first few characters of the name:

```
SHOW SERVICE [=service-name ]
      [TYPE={TELNET|INTERACTIVE|TELBIN} ]
```

- The SHOW BRIDGE command now displays the value of the aging timer (Figure 13 on page 32, Table 8 on page 32).

Figure 13: Example output from the SHOW BRIDGE command.

```

Remote Bridge
-----
Bridge Address      : 00-00-cd-00-0d-4d
Bridge Name        : CentreCOM AR720 version 1.8.1-00 08-Sep-1999
Spanning Tree Protocol : ON
Filter Learning    : ON
Number LAN Ports   : 2
  Port Number      : 1
  Port Address     : 00-00-cd-00-0d-4d
  CAM              : Enabled

  Port Number      : 2
  Port Address     : 00-00-cd-00-0d-82
  CAM              : Enabled
Number Virtual Ports : 1
  Port Number      : 3
Number of Groups   : 1
Ageingtime         : 300
Uptime             : 12133
-----

```

Table 8: New fields in the output of the SHOW BRIDGE command.

Parameter	Meaning
Ageingtime	The value in seconds of the ageing timer, after which a dynamic entry is removed from the filtering database.

- The ENABLE ACC CALL DEBUG command syntax has been modified. The PORT parameter has been added to specify the port to which debugging information is to be sent. This allows debugging to be enabled from a script:

```

ENABLE ACC CALL=name DEBUG={UTILISATION|DEMAND|PACKET|
  PKT|PORT|SCRIPTS|DIALIN|ALL} [PORT=port-number]

```

- The SHOW ACC CALL command now displays the currently assigned IP address pool and the port state of any ports used by the call (Figure 13 on page 32). Two new fields, *IP Pool* and *Port State*, have been added (Table 8 on page 32).

Figure 14: Example output from the SHOW ACC CALL command.

```

ACC call details

Name: dialin
  State ..... Enabled
  Direction ..... Both
  Line ..... Modem
  PPP Template ..... 2
  IP Pool ..... Not set
  Remote Call ..... Not set
  Encapsulation ..... Enquire
  Authentication .... Auto
  Dial Number ..... Not set
  Reset Script ..... nvs:reset.mds
  Dial Script ..... flash:dial.mds
  Connect Script .... nvs:connect.mds
  Active ..... No
  Accounting ..... Disabled
  Debug ..... Disabled

Port(s):
  Port3:
    Port State ..... idle
    Number activations ..... 2
    Start time last activation ... 06-Mar-1996 13:36:49
    End time last activation ..... 06-Mar-1996 13:40:18
    Last user ..... joeb

```

Table 9: New fields in the output of the SHOW ACC CALL command.

Parameter	Meaning
IP Pool	The name of the IP address pool used to assign IP addresses when this call is used for dial-in SLIP sessions, or "Not set" if an IP address pool has not been assigned.
Port State	The state of the port; one of the values listed in Table 10 on page 33

Table 10: ACC asynchronous port states.

State	Description
Idle	The port is not active
Dialing	The port is running a dial script
Connecting	The port is running a connect script
Open	The port is connected to a remote system
Waiting	The port is waiting in a connect script to receive the specified input
Sending	The port is sending a text line from a connect script to the modem
Connect wait	The port has dialled the remote end and is waiting for the remote end to answer
Deactivate	The port is currently hanging up the current call
DTR drop	The port has dropped the DTR line to the port to hang-up the call
Dial sending	The port is sending a text line from a dial script to the modem

Table 10: ACC asynchronous port states. (Continued)

State	Description
Dial waiting	The port is waiting in a dial script to receive the specified input
Reset	The port is currently being reset
Reset start	The port is starting the reset process
Reset waiting	The port is waiting in a reset script to receive the specified input
Reset sending	The port is sending a text line from a reset script to the modem
Reset wait done	The port is waiting for a short period after sending the reset script, for the modem to settle
Reset done	The port reset is complete
Dial backoff	The port has backed off dialing the remote end as it detected the remote end was busy

- ACC call names and RSVP proxy names can now include the hyphen character "-". This brings them in to line with ISDN call names.

Availability

Software Release 1.8.1 is available immediately as a FLASH release for upgrading existing routers. The release file can be downloaded directly from the Software Updates area of the Allied Telesyn web site (<http://www.alliedtelesyn.co.nz/support/updates/patches.html>).

Software releases must be licenced and require a password to activate. To obtain a licence and password, download a Software Upgrade request form from the Software Updates area of the Allied Telesyn web site (<http://www.alliedtelesyn.co.nz/support/updates/patches.html>), complete the form and fax it along with a company purchase order to Allied Telesyn Sales on +64-3-377 8870, or call Teltrend Customer Service on 0800 808 909.

Software Release 1.8.1 will not be available as an EPROM upgrade.

Installation

WARNING: For routers upgraded from Software Release 7.0 or earlier, once the release has been fully installed, retrofitting previous releases can not be safely accomplished without completely re-initialising NVS memory. See "Upgrading From an Older Release" below.

Upgrading From an Older Release

The following points should be noted when upgrading from Software Release 7.0 or earlier. Refer to the release notes provided with Software Release 7.0 when upgrading from earlier releases. These issues do not affect routers running Software Release 7.2 or later.

Enhancements to Asynchronous Call Control

Asynchronous Call Control (ACC) has had major enhancements since Software Release 7.0. As a result of structure changes in the NVS configuration storage, all ACC calls must be deleted and then redefined when upgrading from an earlier software release. It is recommended that users create a text file containing the list of commands used to create all their ACC calls. The new script processor can be used to add the calls once the upgrade has been completed.

More Efficient Use of Battery-Backed RAM

Because of the importance of the following information, it is included in this release note for those upgrading from software releases earlier than 7.0.

Software Release 7.0 allowed a previously unavailable block of 24 KBytes of BBR to be used for storing patches and configuration information. When a pre-Software Release 7.0 router is upgraded to 7.x for the first time, this space will be reclaimed. There is a side effect of this when considering backwards compatibility with previous software releases. The consequence of changing back to the previous 6.8 release subsequent to booting with 7.x, is that the IP configuration will be lost, unless a special patch is run with 6.8.

Two cases arise, and should be considered:

1. The Software Release 6.8 EPROMs are removed and Software Release 7.x EPROMs installed. If Software Release 6.8 is then re-installed, and the router rebooted, the IP configuration will be lost. This is because the configuration of the relevant 6.8 patch with the EPROMs can not be achieved while running 7.x (i.e. a 6.8 patch can not run with 7.x EPROMs).
2. The Software Release 6.8 EPROMs are left in the router and Software Release 7.x is run from FLASH memory. In this case, the IP configuration can be retained by making the Software Release 6.8 EPROM and patch the DEFAULT install, and making the Software Release 7.x FLASH release the preferred install. By deleting the preferred (FLASH) install, the Software Release 6.8 release can be safely run.

ISDN Call Parameters

When upgrading from a release earlier than 7.0, note the following from the 7.0 release notes:

1. Disconnect any ISDN cables before turning on the router with new EPROMs.
2. Enter the following command for each ISDN interface after installing the updated software:

```
SET Q931=n SUB= SPID= SPID2= SUB2= NUM2=
```

where *n* is the ISDN interface number. For example, if the router has one BRI interface and two PRI interfaces, enter the following commands:

```
SET Q931=0 SUB= SPID= SPID2= SUB2= NUM2=
```

```
SET Q931=1 SUB= SPID= SPID2= SUB2= NUM2=
```

```
SET Q931=2 SUB= SPID= SPID2= SUB2= NUM2=
```

3. Reconnect the ISDN cables.

IP RIP Definition

When upgrading from any release of software older than 7.2, it is necessary to note and redefine the RIP parameters. With the advent of enhanced RIP, the way RIP is defined has changed. The following highlights the implementation differences:

- RIP has been enhanced to define the interfaces on which RIP is to be sent and received, and whether to accept RIP V1, V2 or both types. In older versions of software, RIP was accepted on any interface. This is no longer the case.
- When defining a RIP entry, it was previously only necessary to define the address to "RIP" to, with no other option. With Software Release 7.2, the RIP IP address is no longer necessary when defining a RIP entry to a broadcast address on a specific interface. The router automatically works out the broadcast address as a function of the interface address and subnet mask.
- The most important change is that if an IP address is defined on an interface, that address is the only address that RIP will be received from on that specific interface. It follows that it is possible to define multiple RIP receives on a single interface. This offers enhanced security and control over which RIP broadcasts are accepted by the router. It is still possible to RIP to specific addresses by interface, but if both a broadcast and specific RIP are required, then two RIP entries need to be added for the interface.

Examples

Old syntax: `ADD IP RIP=192.168.1.255 Send RIP via broadcast address, and receive from any source address on the interface associated with this subnet.`

New syntax: `ADD IP RIP INTERFACE=ETH0 Send RIP on broadcast address on this interface, and receive from any source address.`

Incorrect: `ADD IP RIP=192.168.1.255 INTERFACE=ETH0` This is incorrect as the only source address from which RIP information will be accepted on this interface is a broadcast address.

Upgrading From Software Release 7.2

Due to major changes in 7.4 encryption capabilities, 7.4 PPP and Frame Relay encryption are not compatible with 7.2 PPP and Frame Relay encryption. It is recommended that when any routers in a network using PPP or Frame Relay encryption are upgraded, all routers in the network are upgraded at the same time.

Upgrading From Software Release 7.4

In Software Release 7.4 the router was given the ability to have an explicitly set local IP address. The purpose of this feature was that any IP packet that was created locally within the router would be sent out with the local IP address as its source address, unless a different source address was specifically set.

Due to an oversight, this feature was not fully implemented. In particular, the GRE module did not follow the rule that local generated packets must use the local IP address as source IP address. As a result, the source address applied to a GRE-generated packet was the IP address on the port through which the packet was sent.

From Software Release 7.6.1, this oversight has been cleaned up, so that GRE-generated packets will use the local IP address if one has been defined. This will mean that some configurations which worked correctly on Software Release 7.4 software will not necessarily work correctly on Software Release 7.6.1 or higher. In particular, if the router is configured to apply both GRE and NAT, and has been configured with a local IP address which is within the range of addresses to which NAT is to be applied, it is possible that packets that are being tunnelled by GRE will no longer be successfully delivered. The solution to this problem is to either remove the local IP address that has been configured on the router, or to set the local IP address to an address that is not within the range of addresses to which NAT is to be applied.

Upgrading from Software Release 7.6

This release makes some significant changes to the command structure of PPP and NAT/Firewall in particular. Configuration and scripts may not work as expected. Please consult the release notes for 1.8.1 to check the impact that changes are likely to have.

In particular, note that (if it is licensed) the Firewall module in 1.8.1 changes the way that NAT was configured. These areas should be given careful consideration. Save all existing script and configuration files prior to upgrade.

Upgrading from Software Release 7.7

There are no issues upgrading from Software Release 7.7 to Software Release 1.8.