

---

# GRF 1.4.6 Release Notes

January 30, 1998

---

This release note describes the 1.4.6 software release of the GRF operating system. Software release 1.4 supports the GRF 400 and GRF 1600 as well as systems configured with RMS nodes.

- 1.4.1 GRF 1600 beta release September 15, 1997
- 1.4.2 Limited release October 23, 1997
- 1.4.3 General release November 17, 1997

Software release 1.4 does not support the first versions (v1) of the ATM OC-3 or FDDI media cards. These cards are informally called ATM Classic or FDDI Classic. Only the hardware-enhanced versions of these cards, ATM/Q and FDDI/Q, will continue to be supported.

Note that SP Switch Router orders have always been filled with ATM/Q and FDDI/Q media cards.

### *GRF 1.4 Addendum*

This release also includes an addendum of configuration information available from the ftp site as the file `Addendum.pdf`. The 1.4 Addendum has three sections:

- ECMP (Equal-cost Multi-path Routing)
- Frame Relay (includes Frame Relay Switching configuration)
- Bridging (this section replaces chapter 15 in the *GRF Configuration Guide 1.4*)

## Contents

Upgrading system software . . . . .	3
New features and advisories in 1.4.6. . . . .	4
Advisories . . . . .	4
New system features . . . . .	4
New GateD features . . . . .	6
Known problems in 1.4.6 . . . . .	7
ATMP notes . . . . .	7
ATM OC-3 notes . . . . .	8
ATM OC-12 notes . . . . .	9
Ethernet notes. . . . .	10
FDDI notes. . . . .	10
Filtering notes . . . . .	11
Frame Relay notes . . . . .	11
HIPPI notes . . . . .	13
HSSI notes . . . . .	13
IS-IS notes . . . . .	14

SONET notes .....	14
System notes .....	14
GateD advisories .....	17
GateD notes .....	20
Problems fixed in 1.4.6 .....	22
ATMP .....	22
ATM OC-3c .....	23
ATM OC-12c .....	23
Frame Relay .....	23
SONET .....	24
System .....	24
GateD .....	25

## Notes about upgrading from 1.2.17 to 1.4.6

### Note 1:

Release 1.4 software does not support early versions (v1) of the ATM OC-3c and FDDI media cards. Remove these cards before you do the installation. These media cards do not boot and come up under 1.4 software. They cause this message to be displayed repeatedly:

```
"Cannot load the card because it cannot find the files."
```

The only ATM OC-3c and FDDI cards that should be installed in the router will have a faceplate label that says "ATM/Q" or "FDDI/Q".

### Note 2:

After you enter the **grfins** command to install the new software, you may see a message similar to the following:

```
Warning; please reboot -i the box at this point. Do not grwrite.
```

This means you need to enter the **reboot -i** command before doing any thing else.

The **-i** option ensures correct file versions are maintained, and prevents troublesome versions of old files from being retained through the upgrade.

### Note 3:

The `/etc/grinchd.conf` configuration file is no longer used. When you upgrade, GRF profiles automatically replace the `/etc/grinchd.conf` file. Variables from `grinchd.conf` are now fields in the Card, User, Dump, Load, and System profiles.

Refer to the *GRF Configuration Guide 1.4* for a description of the CLI and the GRF profiles.

The GRF user interface now has two aspects:

- One is a command-line interface (CLI) interface in which system parameters are configured in profiles. You use the CLI to access the profiles and assign values to the fields they contain.
- The second is a UNIX interface in which system parameters are configured by editing configuration files. This shell is started with the CLI **sh** command. The shell supports standard and GRF versions of UNIX networking and configuration commands. CLI and GRF commands are described in the *GRF Reference Guide 1.4*.

## Upgrading system software

Obtain the release files from the IBM FTP server:

```
service2.boulder.ibm.com
```

You are prompted for the SP Switch Router customer ID and password when you ftp to this server.

This release consists of the 1.4.6 code, release note (RN), and 1.4 Addendum files:

```
- 1.4.6.ibm.TAR.gz
- 1.4.6.ibm.root.gz
- RN1_4_6.txt
- RN1_4_6.pdf
- addendum.pdf
```

To upgrade GRF software, follow this procedure:

Log on as root and start the UNIX shell:

```
super> sh
```

Prepare to load the 1.4 release on the router's internal flash device:

```
# mountf -w
# cd /flash
# mkdir tmp
# cd tmp
```

Then, ftp to this server:

```
# ftp service2.boulder.ibm.com
ftp>
```

Enter the SP Switch Router customer ID and password as requested.

Now change to the /releases directory:

```
ftp> cd releases
ftp> cd A1_4_6
```

Set the file format and download the files:

```
ftp> bin
ftp> get 1.4.6.ibm.TAR.gz
ftp> get 1.4.6.ibm.root.gz
ftp> get RN1_4_6.txt
ftp> get RN1_4_6.pdf
ftp> get addendum.pdf
ftp> quit
# cd /
```

The next commands unpack zipped files and clean up flash:

```
# grfins --source=/flash/tmp --release=1.4.6.ibm --activate
# flashcmd -w rm -rf /flash/tmp
# umountf
```

Reboot installs the files on RAM, enter:

```
# reboot -i
```

or use

```
# shutdown -r now
```

When the router is rebooted, it should display that it is running kernel version 1.4.6.

# New features and advisories in 1.4.6

## Advisories

### - 1MB media cards not supported

Release 1.4 does not support media cards that have less than 4MB of buffer memory per side. These cards have not been supported in releases beyond the NetStar 4.x series.

Before you upgrade from a NetStar release to new software, please consult with Customer Support.

### - Save before you upgrade

When installing a new software version, make sure all changes to files in `/etc` that you intend to be carried forward to the new release have been saved to flash. Execute **grwrite** *-before* – you begin the software upgrade. If you execute the **grwrite** command after doing a **grfins**, the changes will only be saved to the currently running version.

### - Use binary format to ftp release files

You must transfer system images in binary mode when using ftp to get a new release. The system cannot boot if a newly-installed system image is transferred in ASCII format.

### - File names

Various OS utilities used in managing GRF software releases and configuration files do not support file names with spaces. Use the underbar character to separate words in a file name.

### - AAL advisory:

AAL 3/4 is not supported on the ATM OC-3 media card. The GRF supports only AAL 5 on ATM. Circuits configured with other AAL types will still be AAL 5.

## New system features

### - ATMP support for VPN tunnels

ATMP (Ascend Tunnel Management Protocol) is a layer 3 UDP/IP-based protocol that provides a cross-WAN (Internet or other) tunnel mechanism using standard Generic Routing Encapsulation between two Ascend units. GRF routers support a subset of ATMP that enables the GRF to function as an ATMP home agent in IP gateway mode.

The following media support is provided:

- Ethernet from the foreign agent to the home agent (GRF)
- HSSI (Frame Relay) from foreign agent to home agent (GRF)
- HSSI (Frame Relay) from home agent (GRF) to home network
- ATM OC-3c from foreign agent to home agent (GRF)
- ATM OC-3c from home agent (GRF) to home network

Refer to the *GRF Configuration Guide* for configuration information.

### - Equal Cost Multi-path routing - ECMP

The Equal Cost Multi-path (ECMP) feature provides an ability to efficiently modulate traffic to destination networks. With ECMP enabled, multiple gateways for destination network or host prefixes (addresses) can be legally installed in the GRF route table. ECMP routes packets toward a destination network by splitting up the packet load between different, but similar, paths. This release supports a maximum of eight gateways per ECMP group. Routes can be specified statically or dynamically.

GateD supports equal cost multipath prefixes learned via the OSPF and OSPF\_ASE protocols. Also, I/CBGP will resolve prefixes to multiple gateways if the nexthop resolving protocol is OSPF or OSPF\_ASE.

Refer to the *GRF 1.4 Addendum* for configuration information (addendum.pdf file).

### - TACACS+ authentication support

The GRF can be configured as a client of an authentication program running on a remote server. TACACS+ is now added to the RADIUS and securID options.

Refer to the *GRF Configuration Guide* for configuration information.

### - RFC 1483 support

The GRF implements IEEE 802.1D transparent bridging on GRF Ethernet and FDDI interfaces, and on ATM OC-3c interfaces using RFC 1483 encapsulated bridging over PVCs. Simultaneous bridging and routing over the same interface is supported, a GRF interface participating in a bridge group can still route normally.

Refer to the *GRF 1.4 Addendum* (addendum.pdf file) for configuration information. The *Addendum* updates and corrects chapter 15 in the *GRF Configuration Guide*.

### - Frame Relay NNI support

GRF HSSI and SONET OC-3c media cards now support both routed and switched Frame Relay (UNI and NNI). In the same GRF system, HSSI or SONET media cards can be configured to route Frame Relay packets while different HSSI or SONET cards can be configured to switch Frame Relay packets. Frame Relay links are configurable as UNI-DTE, UNI-DCE or NNI.

Refer to the *GRF 1.4 Addendum* (addendum.pdf file) for configuration information.

### - Minimum system memory now 128MB RAM

GRF 400 and GRF 1600 systems now ship with 128 MB of RAM. Additionally, a new 128MB DRAM package has been released to support up to 512MB of system memory.

### - New ATM/Q statistics commands

In this release, the ATM/Q maintenance command **maint 15** is deleted. **maint 16** now displays the ATM OAM F5 statistics per VC. Here is an example of the receive side:

```
GR 2> maint 16 0 0 100
[RX] OAM F5 Stats for Port 0 VPI/VCI 0/0 [RX] Receive
-----
[RX]           End to End  Segment
[RX] Fault Management
[RX] AIS:             0    0
```

```

[RX] RDI:                0    0
[RX] Loopback:           0    0
[RX] Continuity:        0    0
[RX]
[RX] Performance Management
[RX] BWD Reporting:     0    0
[RX] FWD Monitoring    0    0
[RX] Monitoring & Reporting: 0    0
[RX]
[RX] Activation/Deactivation
[RX] Continuity:       0    0
[RX] Performance Mon: 0    0

```

Additionally, RX DMA Error, TX DMA Error, and RX Descriptor Unavailable are now reported by the **grstat -a switch** command. Here is an example:

```

grstat -a switch

Switch statistics
count description
0 RX packets
0 RX bytes
0 RX errors
0 Switch receiver reset
0 Receive Encoding Errors
0 Receive Running Disparity Errors
0 Receive Receiver Errors
0 Receive Running Checksum Errors
0 Receive Descriptor Unavailable
0 Receive DMA Error
0 TX packets
0 TX bytes
0 Transmit Data Errors
0 Transmit Fifo Parity Errors
0 Transmit Internal Parity Errors
0 Transmit Connection Rejects
0 Transmit FIFO full
0 Transmit no descriptor
0 Transmit RX don't free bit set
0 Transmit adjust header up - not used
0 Transmit adjust header down - not used 0 Transmit DMA Error

```

## New GateD features

### - ECMP support

GateD supports equal cost multipath prefixes learned via the OSPF and OSPF\_ASE protocols. Also, I/CBGP will resolve prefixes to multiple gateways if the nexthop resolving protocol is OSPF or OSPF\_ASE.

---

# Known problems in 1.4.6

## ATMP notes

### **Advisory: ATM large packet performance**

Due to the processing overload for IP fragmentation, the ATM card does not forward full-size encapsulated packets to the Ethernet card. The GRE packet is 1528 bytes, larger than the Ethernet MTU of 1500 bytes. The ATM card must fragment the packet before sending.

This slows the Ethernet transmit rate.

### **Advisory: ATMP memory requirements**

On the media card, tunnel connection for each mobile node consumes memory comparable to space required by two routes. As additional tunnels are negotiated, less media card memory is available for route tables.

### **Advisory: starting the ATMP daemon**

By default, **aitmd** is not running. To enable ATMP, the administrator creates the file `/etc/aitmd.run`. This is done in the UNIX shell with the command:

```
# touch /etc/aitmd.run
```

Save the configuration change with:

```
# grwrite
```

The daemon starts within 15 seconds and restarts automatically on future reboots. The file is removed normally with **rm** and the removal again saved with **grwrite**.

Removing `/etc/aitmd.run` does not kill the daemon, it only prevents **aitmd** from being restarted. The daemon loads the ATMP configuration from the file `/etc/aitmd.conf`. The configuration can be changed on the fly by editing the configuration file and then sending the process a HUP signal.

### **Advisory: IPX not supported**

Although Max and TNT implementations of ATMP tunnels support IPX, the GRF does not support tunneling of IPX packets.

### **- ARP reply overrides configured Frame Relay address**

An inverse ARP reply received on a Frame Relay circuit overrides the operator-configured IP address. As a result, the GRF can no longer use that circuit. This does not affect the `pvcatmp` connection to a home network since ATMP operation does not depend upon any addressing of the `pvcatmp`. The workaround is to reset the HSSI card on the GRF. HSSI interoperation with the TNT is the source of the inverse ARP.

### **- Media support for ATMP tunneling**

Traffic from a Max or TNT ATMP foreign agent can arrive at the GRF home agent via HSSI, ATM, or Ethernet media cards. Tunneled traffic is not properly decapsulated if it arrives on

other types of media cards such as FDDI. Also, tunnel traffic is not supported over the control board's administrative Ethernet interface.

#### **- Reassembly limitations**

If an ATMP tunneled packet from a foreign agent is fragmented, the GRF is unable to reassemble it, and the packet is not decapsulated or sent to the home network.

Fragmentation can occur in the most typical of networks. For example, when a 1500-byte IP packet is encapsulated and tunneled over Ethernet to the GRF home agent, the encapsulation has made the packet larger, and so it is fragmented down to meet Ethernet's 1500 byte MTU.

Two strategies can deal with the reassembly limitation, one places the control of the MTU in the tunnel link, the second places control at the MAX TNT foreign agent.

One workaround for this problem is to use a frame relay link between the foreign agent and the GRF home agent. Frame relay has a large MTU, and fragmentation will never be required if the mobile node and home networks are using Ethernet or common dial-up connections. Additionally, this workaround eliminates the performance overhead of fragmentation.

The MAX TNT can fragment traffic prior to encapsulation. This avoids the reassembly problem and is the preferred solution when Ethernet is used to connect the GRF to a TNT.

In the TNT ATMP profile, set the `mtu-limit = parameter` to 1440. It is also probably necessary to set the `force-fragmentation = parameter` to "yes" because common endstations unnecessarily set the "don't fragment" bit in every packet. Check with Customer Support to find out if this setting is appropriate.

#### **- Logging information to syslog may be truncated**

If `aitmd` generates large bursts of logging output to `syslog`, such as when dumping state information in response to a SIGINFO signal, the output may be truncated. This problem is not present if `aitmd` is logging to a file. See the `aitmd` man page for a complete description of `aitmd`'s logging options.

## **ATM OC-3 notes**

#### **- Up to 75 logical interfaces/card recommended**

An ATM/Q media card configured with up to 128 logical interfaces may panic and reset itself. In this release, we recommend a site configure up to 75 logical interfaces per ATM/Q card.

#### **- ATM card may not boot under heavy switch load**

If the ATM media card is running under a heavy load from the backplane switch and requires a reboot, it may not be able to complete a reboot. The card may instead enter a loop in which it repeatedly but unsuccessfully attempts to come back up. Remove the load and reset the ATM card.

#### **- SVCs fail if card is reset in traffic**

It has been observed that if an ATM/Q card is under load and an SVC is torn down and attempted to be brought up again, it may not succeed. Remove the load and retry.



**- Traffic shape must follow service on gratm.conf entry**

The following statement will produce a syntax error from the **gratm** parser:

```
Interface ga000 traffic_shape=hshq service=arp0
```

The following statement parses correctly:

```
Interface ga000 service=arp0 traffic_shape=hshq
```

**- Invalid PCR or SCR settings allowed by gratm**

Users are not notified when they configure an invalid setting for PCR or SCR. Invalid settings are not caught by **gratm** and cause the ATM card to set the default to 0 (zero) Kb per second. Setting PCR and SCR to below the minimum range (below 130Kb for PCR, for example) can cause the ATM card to crash as it attempts to move data at less than 16 bytes per second.

**- Old SVC settings prevent connection retry attempts**

If an ATM media card which has an SVC in place is rebooted, it retains the SVC configuration information. When the external node attempts to set up the SVC again, it is refused by the ATM card because one already seems to exist.

**- SPANS signalling not tolerated by ATM media cards**

When a GRF ATM card is connected to a switch with SPANS signalling enabled, ATM cells are dropped at the ATM card. In **maint 4** output, all packets are shown as errors. To work around, disable SPANS signalling on the switch, and reset the ATM media card.

**- Local ping to broadcast address not handled properly**

A local ping to the all 0s broadcast address is not actually sent, and no ARP server or ARP entry is made.

## ATM OC-12 notes

**- Advisory: maximum value of MBS traffic shaping parameter**

The largest allowable maximum burst size (MBS) for the ATM OC-12 card is 255 cells, the range is any value between 0 and 255. Note that MBS values for ATM OC-3c and OC-12c cards are different, this is due to vendor SAR chip design.

**- Ping advisory:**

It has been observed that a ping from a GRF over an OC-12 link to the next hop on OC-12 failed when traffic shape values were set to peak=7500 (no sustain value). The problem disappeared when traffic shape was changed to peak=622000, sustain=622000, burst=32.

**- Configured interfaces advisory**

Currently, 1.4 software supports about 100 configured interfaces on ATM OC-12.

**- Large pings are dropped**

Some ATM OC-12c cards can corrupt packet data on transfers from the RMS where the packet size is greater than the MTU. This condition does not affect forwarded IP traffic. (Example: **ping -s 9788** was seen to fail.)

**- ATM OC-12 may boot repeatedly**

An OC-12 card may boot repeatedly when receiving traffic. Eventually it boots successfully.

**- gratm erroneously allows ATM-12 PVC proto=vc**

ATM OC-12 does not support `proto=vc` encapsulation. The **gratm** parser does not flag the following statement as an error when configuring an OC-12 media card:

```
PVC ga0180 0/32 proto=vc traffic_shape=high_speed_high_quality
```

## Ethernet notes

**- maint 4 command reports bad statistic**

The RX Collision statistic reported in **maint 4** output should be ignored. Erroneous counts are often displayed. No packets are lost or corrupted because of the reporting problem.

**- ICMP UNREACHABLE messages not generated**

The Ethernet card does not generate and send back ICMP UNREACHABLE messages when attempting to forward an IP packet to a destination not found by ARP on the Ethernet.

**- No promiscuous mode**

Ethernet interfaces do not operate in promiscuous mode. If a program, e.g., **tcpdump**, tries to make an Ethernet interface promiscuous, the Router will ignore the request.

## FDDI notes

**- Binary load may fail**

Occasionally a load of FDDI/Q run-time code gets out of sequence. Resulting output messages look like this:

```
...Bad Sequence Number: 1, expecting 2
...Bad Sequence Number: 2, expecting 3
...Bad Sequence Number: 3, expecting 4
```

This is caused by an out-of-date boot loader on the FDDI/Q card. A card's boot loader can be identified by a line as follows in the output from **grconslog** when the card is rebooted:

```
FDDI/Q boot loader - version 1.2
```

Using the **greset** command to force the FDDI/Q card to reboot will clear the internal communications bus and boot the card successfully. Under controlled conditions, it is possible to update a card's flash memory with an up-to-date boot loader that would eliminate the possibility of this problem occurring. Contact customer support for assistance if you experience this problem and would like to update your card's boot loader. Note that this is only a problem when a card is reset; it does not cause any problem during normal operation.

**- LEDs not consistent**

When the lower port pair on a FDDI/Q card is configured as a single DAS interface and it goes to WRAP-A state, then the wrong faceplate LED (port B TRX) comes on for transmit.

**- No FDDI/Q promiscuous mode**

FDDI/Q interfaces do not operate in promiscuous mode. If a program, e.g., **tcpdump**, tries to make a FDDI/Q interface promiscuous, the Router will ignore the request.

**Filtering notes****- Packet logging documentation needs clarification**

A statement about packet header logging in the filtering chapter of the *GRF Configuration Guide* is incomplete. The documentation says that “logging is a recording activity in which a copy of a filtered packet is sent to a logging host.” The following sentence is needed to make the description clear:

When the action assigned to a filter is logging, as it is in packet header logging, the “filtered” packet is detected, its header (or specified portion) is logged, but the filtered packet is not dropped.

**- Filtering documentation omits SONET and Ethernet support**

The Filtering chapter in the *GRF Configuration Guide* does not accurately list the media cards that support filtering (p. 11-7). Supporting cards are: HIPPI, FDDI, ATM OC-3c, HSSI, Ethernet, and SONET. Filtering is not currently supported on ATM OC-12c. The **filterd** man page does not include Ethernet and SONET support.

**- Packet header logging misconfiguration can cause loops**

If a packet header logging filter is configured on the receive side of a media card to log ICMP messages, AND the filter is configured to log packet headers to a system out the same media card, a loop can occur if there is no application listening on the UDP port configured to receive the logged packet header.

This loop occurs because the logging system must send back an ICMP port unreachable message to the GRF when there is no application running, and this ICMP message will be selected by the filter and cause another attempt to log a packet header, which will in turn cause another ICMP port unreachable message to be sent to the filter on the GRF.

Avoid this situation by configuring the logging filter to ignore ICMP messages from the logging host to the media card's configured IP addresses.

**- Packet header logging to the RMS not supported**

Packet header logging will not forward packets to the IP address of the administrative ethernet (de0 or ef0), or to an IP address out that interface.

To log packet headers to the RMS itself, the workaround is to use the IP address of the interface on which the filter is configured.

**Frame Relay notes****Advisory: stopping and restarting Frame Relay**

The Frame Relay daemon, **fred**, starts automatically at system boot. However, if you enter the command “**fred**” when Frame Relay is already running, you will create a second instance of the daemon. The results of this are unpredictable.

Under normal operations, **fred** will not need to be restarted. However, if you get no response from a **grfr** display command (**grfr -c lss**, for example) **fred** may have hung. You will need to determine the **fred** process ID and then use the **kill process\_id** command to restart **fred**.

#### Advisory: Configure HSSI port 0 first

If only one port on a HSSI card is to be used, configure and cable port 0 first. Port 1 will not operate if port 0 is not active.

#### Advisory: grfr displays LMI and other information

The **grfr** command supports the GRF Frame Relay daemon, **fred**, and provides a large number of link and PVC information displays. Several media card **maint** commands are replaced by the expanded **grfr** capability. For example, the HSSI media card **maint 3** command no longer displays link LMI type. The information is available via the **grfr -c dlc** display command. Please refer to the **grfr** man page for a complete list of **grfr** display options.

#### Advisory: controlling PVCs autoadded by a remote system

In the Link section of **grfr.conf**, the **AutoAddGrif** option can be selected (set to **Auto**) for a GRF interface to allow remote systems to add a PVC to that GRF interface. These autoadded PVCs are configured when **fred** detects a PVC added by the remote system. To delete an autoadded PVC:

1. Reconfigure the link parameter so **fred** does not add the PVC again. You must remove the **AutoAddGrif** parameter from the link statement and run the **grfr -c clp** command.
2. Then you must use the **grfr -c crp** command to delete the autoadded PVC, it will not be deleted automatically.

#### - Cannot add multiple PVCs at one time

In this release, the Frame Relay daemon, **fred**, cannot accept multiple configuration changes. For example, if you add entries for 5 PVCs in the **grfr.conf** file, only one PVC is configured when you run **grfr**. As a workaround, please add one PVC at a time. For this example, run the **grfr** command 5 times to configure the 5 PVCs.

#### - Frame Relay interface 0 configuration requirement

When used in Frame Relay mode, HSSI and SONET card configuration requires that logical interface number 0 on each port **must** be configured. This enables the Frame Relay daemon (**fred**) to properly communicate LICS traffic via the card.

Each interface 0 **must** be identified in **/etc/grifconfig.conf**.

For example, given that a HSSI card in slot 3 is configured for Frame Relay, then interfaces **gs030** (slot 3, physical interface 0, logical interface 0) and **gs0380** (slot 3, physical interface 1, logical interface 0) need to exist in the card configuration. Only one entry is required for the SONET card's single active interface.

These are the **required** entries in **/etc/grifconfig.conf** (SONET in slot 1 included):

#	name	address	netmask	broad_dest	arguments
	gs030	-	-	-	up
	gs0380	-	-	-	up
	go010	-	-	-	up

**- Frame Relay misinterprets link status**

When a GRF is cabled in a point-to-point (PPP) Frame Relay HSSI configuration, the Frame Relay daemon (**fred**) may decide a link is inactive when it really should be active. If you are convinced that the link is really up and correctly configured, then try resetting the HSSI card on each side (see the **grreset** command man page).

**- Frame Relay (HSSI) link goes up and down**

It has been observed that when a load is put on the HSSI interface which is greater than its packet rate, the link goes up and down because the LMI information is lost. Decrease the load to clear this problem.

**- ARP reply overrides configured Frame Relay address****12825**

When an inverse ARP comes in, Frame Relay on HSSI extracts the peer address, places it in its ARP cache, and overwrites any previous entry for that address even if that address was set in `/etc/grarp.conf`.

## HIPPI notes

**- HIPPI IBM H0 mode loses pings at about 956 bytes**

This problem is caused by HIPPI not stripping IBM H0 padding from the end of the ECHO REPLY packet when sending it back to kernel.

## HSSI notes

**- HSSI grarp entries dropped for no apparent reason**

The HSSI media card will sometimes not respond to **grarp** request for list of ARP entries.

**- HSSI can lose PPP connection during high packet rates**

The HSSI media card can drop LCP keepalives and lose connectivity during high packet rates. Set the keepalive count high enough to make it less likely to drop all of them and lose connection.

**- HSSI pings sent across link that should get ICMP redirected**

When a ping that should get ICMP redirected gets sent across a HSSI link, the redirect comes back, but, in addition, the ICMP request itself comes back.

**- HSSI MTU not set correctly when protocol changes**

When changing a HSSI framing protocol from Frame Relay to PPP, the MTU does not get correctly set to 1500. This causes fragmentation problems when routing through this interface from higher MTU media. The MTU can be set correctly by setting the correct value in `/etc/grifconfig.conf` and configuring the interface with **grifconfig**.

## IS-IS notes

### - IS-IS supports 126 circuits, not 128

Note that the IS-IS protocol can support 126 circuits, documentation specifies 128.

### - IS-IS does not maintain adjacencies over HSSI

IS-IS will form an initial adjacency over HSSI circuits, but after a short period will lose the adjacency.

## SONET notes

### - LQR quality remains at zero

It has been seen that according to LQR, quality is always reported as 0 (zero).

### - PPP LQR status reports incorrect remote period

It has been seen that according to LQR, remote period is always reported as 0 (zero).

## System notes

### - Directed broadcast packets are not forwarded

GRF media cards do not correctly respond to and do not forward directed broadcast packets destined for networks to which the GRF is directly connected. This problem only applies to broadcasts to directly-connected networks. If there is another hop to the broadcast target, then the GRF forwards directed broadcast packets correctly.

### - IP aliases are assigned wrong netmask

Although an IP alias is assigned its own netmask in the `grifconfig.conf` file, **grifconfig** assigns the alias the netmask of the primary interface instead.

### - Reset the router when changing media cards

When replacing the media card in a slot and changing from media type A to different media type B, you must reboot the GRF. Replacing media card type A with another type A media card does not require a reboot.

### - MTU default setting problem

Administrators may set a non-default MTU value for an interface on any GRF media card by specifying this on the interface line in the `/etc/grifconfig.conf` configuration file. This value can only be changed by modifying the associated line in this file and running **grifconfig** for the corresponding interface.

Default MTU settings can be reset on an interface by either explicitly setting the default value for the interface in the configuration file, or by clearing the value in the configuration file and resetting the card.

**- Finger to a station not running daemon makes GRF coredump**

Running the "finger" program to find a user on a remote box will cause the program "finger" to abort if finger daemon is not running on the remote box.

**- iflash command options have been updated**

Please refer to the **iflash** man page for new options, they do not appear in the printed 1.4 manuals.

**- TACACS+ does not support encryption**

In this release of 1.4, the TACACS+ implementation does not support encryption.

**- grarp shows wrong interface for FDDI ARPs**

The **grarp** command does not correctly display interface names for FDDI interfaces.

**- grarp command shows incorrect information**

The **grarp -p<card> -a** command should return a list of the ARP table entries on a given card. The erroneous behaviors of this command are:

1. If there is no interface zero defined on the card (e.g. `ga030`), the command does not list any interfaces.
2. If other interfaces besides interface zero exist, the command will list all the ARP table entries as belonging to interface zero, regardless of which interface actually owns them.

The workaround is to execute **grarp -i <ifname> -a** for each interface on the card.

**- telnet sessions time out**

If a telnet session is opened to a GRF and left running for six hours or more with no traffic through the telnet connection, the telnet session will be dropped. You must re-establish the telnet session, it will come back up with no problem.

**- Route table mismatch can occur**

When this message appears in the `messages.log`, it usually means there are route mismatches between the media card and the kernel.

```
Messages: kernel: Error rsp 65525 to rt upd for dst[0]
```

The media card will keep forwarding packets based on the route information it has. To fix the route problem, use the **greset** command to reset the card and cause a new snapshot of the correct route table to be downloaded to the card.

**- telnetd rejects connections not resolved via dns**

If DNS is configured, the telnet daemon does not allow any connections from a given host when it cannot resolve the source IP address of the host via DNS.

**- Cannot amount without reboot**

While using NFS on a GRF, if the NFS server reboots, the mount point is no longer accessible even when the server is back up and running. The only way to clear this situation is to reboot the GRF.

**- Using ntpdate/date to set the date/time freezes system**

Date changes to the GRF system MUST be made by interrupting the boot process and setting the date while in single user mode.

**- Use grstat, not netstat, to return layer 2, 3 data**

The `netstat -I <interface> -w 1` command fails to report GRF layer 2 and 3 data. Use the `grstat` command to get layer 2 and layer 3 information per interface. See the `grstat` man page for specific command information.

**- Do not use suspend command in CLI UNIX shell**

Issuing the `suspend` command from a shell spawned by the CLI `sh` command causes the connection to hang. Recover from the hang by exiting this session and logging in again.

**- Leading spaces cause errors in maint commands**

When a `maint` command has one or more leading white spaces, the `grmb` program does not recognize the command. This mis-typed example does not work:

```
GRX> maint 4 7
```

This command is entered correctly:

```
GRX>maint 4 7
```

**- ping: caution against flood pinging an interface locally**

Do not perform a flood ping from the operating system to a media card interface. This causes excess traffic on the internal communications bus that can degrade the reliability of the system.

**- ping: responding to a deleted route**

If you delete an attached network route for a logical interface on a media card, but leave the interface otherwise configured, the media card will still respond to pings (ICMP ECHO packets) of the interface's IP address.

**- tcpdump: causes problems during heavy traffic**

Using `tcpdump` against any HSSI, or Ethernet interface undergoing heavy traffic is not recommended. The flood of combus traffic generated by `tcpdump` can cause other cards to reset.

**- GRF 400 control boards crash with RAM parity error**

A problem has been observed on some GRF control boards shipped with Revision 2 (dated 03/06/97) BIOS. This caused systems configured with 256M of memory to crash with a message "RAM parity error". The BIOS should be upgraded to Revision 3 (dated 04/17/97).

NOTE: Earlier boards with Revision 1 (dated 11/26/96) BIOS do not have this problem. However, the DRAM timing must be set to 70ns instead of 60ns for proper operation.



**- GR-II resets ATM/Q after download to Ethernet card**

GR-II routers may reset another card after downloading a large route table, one in excess of 50K routes.

**- ICMP throttle settings in CLI do not take immediate effect**

The CLI settings that control how many ICMP messages of each type a card will generate each second do not take immediate effect on the media card. To have them take immediate effect without resetting the card, execute any of the following **grinch** commands:

1. to set echo-reply, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.15.3.1=<value>
```

2. to set unreachable, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.15.3.2=<value>
```

3. to set redirect, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.15.3.4=<value>
```

4. to set TTL timeout, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.15.3.9=<value>
```

5. to set param-problem, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.15.3.9=<value>
```

6. to set time-stamp-reply, use

```
grinch -p<slot> 2.12.2.card+1.4.port+1.15.3.11=<value>
```

**- HDLC settings not set in card in CLI**

You must use **grinch** commands to change HDLC settings.

1. to set debug level, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.2.3.1=<value>
```

2. to set keepalive-enabled, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.2.3.2=on
```

3. to set keepalive-interval, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.2.3.3=<value>
```

4. to set keepalive-error-threshold, use:

```
grinch -p<slot> 2.12.2.card+1.4.port+1.2.3.4=<value>
```

Note: if you now reboot the box, you must rerun the command(s).

## GateD advisories

**Advisory: ECMP multipath parameter revised**

The printed *GRF Reference Guide 1.4* has incorrect options for the ECMP multipath parameter. The ECMP parameter is available in a Definition Statement:

```
multipath { on | yes | off | no } ;
```

This parameter enables/disables the installation of multiple gateways for network or host prefixes into the kernel route table. The default is *off*.

The `on` option is the same as `yes`, and enables GateD to install multiple routes for a single source with the same destination but different nexthops.

The `off` option is the same as `no`, and means that each route GateD installs in the kernel will have a unique destination and nexthop.

### Advisory: internal features require additional resources

This release of GateD incorporates several internal features that use more memory and processing resources than do earlier versions. When running a GRF with 60,000 prefixes, you will see that GateD uses up to 10MB more memory. CPU usage will fluctuate 15-25%, at minimum. Route flap or larger routing tables will increase memory and CPU usage.

### Advisory: assign an alias as routerid

It is strongly recommended that you configure an alias address on the loopback interface of your router and use this address as the `routerid` in `gated.conf`. Also, the alias should be used as the local address (`lcladdr`) in peer statements for IBGP. This will enable proper route calculations across all interfaces.

To configure an alias of the loopback interface, add an interface definition to `grifconfig.conf` as shown in the example here. The example sets the IP address of the loopback interface to 127.0.0.1 with an alias of 210.210.210.1. It is important to maintain the 127.0.0.1 address on the loopback interface. Enter the 127.0.0.1 line so it appears before the alias in `grifconfig.conf`.

```
lo0 127.0.0.1 255.255.255.0
lo0 210.210.210.1 255.255.255.0
```

In this configuration, `netstat -in` produces the following output:

```
# netstat -in | grep lo0
lo0 1536 <link3> 909 0 909 0 0
lo0 1536 127/24 127.0.0.1 909 0 909 0 0
lo0 1536 210.210.210 210.210.210.1 909 0 909 0 0
lo0 1536 <GRIT> 0:0x48:0 909 0 909 0 0
```

In this configuration, `ifconfig lo0` command produces:

```
# ifconfig lo0
lo0: loop flags=8009<UP,LOOPBACK,MULTICAST> mtu 1536
inet 127.0.0.1 netmask 0xffffffff
grit 0:0x48:0
```

### GateD older versions advisory:

Versions of GateD older than GateD Consortium distribution 3.5B3 and other BGP implementations based on the code in these older GateD versions have a bug that makes them incorrectly process any BGP optional, transitive attributes with non-zero attribute value lengths.

This bug has been fixed in GigaRouter versions of GateD, but other systems that still have the bug might not interoperate with the GigaRouter's version. Versions of GateD shipped after July 8, 1996, with GigaRouter software versions have had this fix applied. Consult the mailing list archives or contact Customer Support for a copy of the fix for earlier versions.

### GateD trace option advisory:

Tracing significantly decreases GateD speed and efficiency. Unless dynamic routing protocols are being debugged, GateD trace options should be turned off. Two mechanisms are available:

1. remove/add traces in `/etc/gated.conf`
2. use `gdc toggletrace`

**GateD MEDs advisory:**

The BGP implementation treats a route with no declared MED as preferable to a route with any MED. This conforms to the de facto standard implemented by other router vendors, and prevents routing loops when talking to some routers. This implementation differs from the new Internet draft for BGP 4 (draft-ietf-idr-bgp4-03.txt) which specifies that a route with any MED shall be preferred to a route without a MED. This treatment of MEDs may be revised in a future release.

**GateD keywords not supported:**

Certain `/etc/gated.conf` keywords are not supported in the version of GateD distributed with this release.

On the **kernel** statement:

`remnantholdtime`

On the **bgp** statement:

`noauthcheck`

`holdtime 0`

`indelay`

These keywords may be supported in a future release.

**GateD protocols not supported:**

The following protocols are documented in the GateD Consortium's distribution of this version of GateD, but are not officially supported in this release:

HELLO

Router Discovery (client side)

BGP 2

SNMP

**GateD advisory: redirect statement not supported**

This version does not support the `/etc/gated.conf` "redirect" statement.

**GateD: no IGP group support**

This version of GateD does not support the BGP group type "IGP". The BGP group type "routing" is recommended for internal peers.

**GateD: effects of peering upon Route Reflection server**

Configuring one or more Route Reflection clients to peer with a neighboring AS - when the Route Reflection server is also peering with a neighboring AS - can cause a significant degradation in the convergence time for the Route Reflection server.

**GateD: avoid enabling warning messages with BGP**

If BGP is run with warning messages enabled, excessive scrolling to the console may be encountered. In general, it is recommended that the keyword 'show-warnings' within the BGP clause not be used.

If 'show-warnings' must be used, a workaround for excessive scrolling is to direct the GateD trace to a file and to display the file with `more`. If background mode is used, this allows the trace file to be accessible for display at the console, but also allows commands to be entered at the console.

After starting GateD, do the following to set up foreground/ background processing:

1. Enter: `more /var/tmp/gated_trace`  
(or whatever trace file is being used, `less` can be used instead of `more`)
2. Do a Control Z to suspend display by `more`
3. Enter: `bg` (to move `more` to background mode)
4. Enter console commands as desired
5. Enter: `fg` (to return to `more`'s display of GateD trace file)
6. Enter sub-commands to `more` to display the trace file
7. Do a Control Z followed by `bg` to return `more` to the background  
(at any time)

**GateD notes****- ECMP multipath change requires GateD restart**

ECMP (multipath) cannot be changed to on or off with a `gdc reconfig` command. You must restart GateD to change this parameter.

**- BGP Destination Preference Attributes not supported**

Beginning with the 1.4.6 release, BGP Destination Preference Attributes are no longer supported.

**- GateD static routes with retain not overridden in route flush**

When static routes having the retain flag set are configured in GateD, these routes are not kept if the `route flush` command is run.

**- GateD not started by monitor subshell with parser error**

If the `/etc/gated.conf` file does not parse cleanly and there are warnings or errors generated by the parser, GateD is not started by the GateD monitor sub-shell. The `/etc/gated.conf` file MUST parse cleanly for the monitor program to start GateD. Use `/usr/contrib/bin/gdc checkconf` to confirm a clean parse of the GateD configuration file.

**- GSM configuration fails to return to default**

If you configure GSM access using the `gsm` keyword in the `/etc/gated.conf` file and then you later comment out the `gsm` configuration, the command `gdc reconfig` does not return GateD to the default GSM access rules.

To return to the default GSM access rules, provide a minimal GSM configuration with the following command:

```
gsm on ;
```

#### **- Certain IS-IS routes not withdrawn if GateD is stopped**

It has been observed that if the IS-IS protocol is enabled and a route is learned via the protocol from more than one adjacency, GateD may not remove the second, more expensive path to the route when GateD is stopped. The incorrect action does not occur every time one stops GateD (`gdc stop` command), it is timing dependent.

#### **- OSPF terminates when GateD reconfigured**

If the OSPF protocol is running and the user executes the following command:

```
# gdc reconfig
```

the user will observe the termination of the OSPF protocol, including the flush of the LSDB (link state data base) and the termination of adjacencies. Other protocols dependent upon the routes learned or transmitted via OSPF may be affected. This termination occurs every time one is running OSPF and “reconfigures” GateD.

#### **- GateD local synchronization not supported**

Under certain heavy-load conditions, GateD may, by default, send newly-available routes to other peers before installing them in the local kernel routing table and media cards.

This can cause momentary routing black holes. Sometimes the behavior is referred to as routing without local synchronization. Routing with local synchronization is not supported in this release.

#### **- GateD loss of policy configuration**

If an interface is defined as “active” in `/etc/gated.conf` and goes down and then comes back up, GateD does not recover the routing policy associated with that interface.

The policy may be recovered by manually issuing the command

```
# gdc reconfig
```

which rereads the configuration file and does not stop or restart GateD.

#### **- GateD order of policy statements**

The order of the policy statements in the `/etc/gated.conf` file may influence which routes GateD will select. This will be reviewed and changed, or documented, in a future release.

#### **- GateD import filters disregarded (IBGP)**

The Import statement is not required for importation of routes from external or internal peers. If a peer is defined in the group/peer statement, by default, all routes will be imported.

#### **- GateD under stress, may not send/reply to keepalives to/from peers**

Under extremely heavy traffic conditions, BGP time-outs can occur. And, although GateD may have received routes, it may not have replied to or sent keepalive messages to its peers. If this happens, the Router will appear to have “gone away” and will not be considered to receive incoming traffic.

---

# Problems fixed in 1.4.6

## ATMP

### - Configuration changes load correctly

A problem has been fixed in which configuration changes to `/etc/aitmd.conf` were not loaded correctly in certain cases, such as when reconfiguring an existing home network record to use a different circuit.

### - No longer requires PVC to be set as inactive interface

ATMP no longer requires that HSSI and ATM interfaces configured as dedicated circuits to home networks must also be configured in `/etc/grifconfig.conf` as blank interfaces.

### - Pings normally through GRF to other interface

A problem has been fixed in which an end station could ping its locally-attached GRF interface, but not the address of another interface on the same GRF. The interface was not affected in any way and still passed traffic.

### - aitmd command does not affect configuration

A problem has been fixed in which mobile node addresses and other tunnel information would be lost if the user entered the `aitmd` command when `aitmd` was running.

To check that `aitmd` is running, use the `ps` command as shown here:

```
# ps ax | grep ai
```

If `aitmd` is running, it is reported in the second line of output. If there is no second line, `aitmd` is not running:

```
10199 p2 S+    0:00.02 grep ai
      390 00- I    0:00.04 /usr/sbin/aitmd -F
```

### - Configuration changes loaded normally

A problem has been fixed in which some configuration changes to `/etc/aitmd.conf` might not process correctly the first time the HUP signal is sent to reconfigure the `aitmd` daemon.

### - aitmd logs across NFS normally

Normally the ATMP tunnel management daemon, `aitmd`, logs via `syslog`. However, it can log locally with the `-l filename` option to log to a file. If the `-l` option is used to log and the path name is to an NFS-mounted file system, `aitmd` will fail to open the file and exit. The `-l` option only works to log to local files.

### - aitmd reports authentication errors correctly

A reporting problem has been fixed so that `aitmd` reports the foreign agent's primary IP address.

**- ATMP error message at system start-up**

A transient ATMP error message that would be displayed when the GRF booted has been removed.

**ATM OC-3c****- Supports 512 VCs per physical interface**

A problem has been fixed in which ATM/Q media cards did not support the usual 512 VCs per physical port.

**- ARP entries set correctly during on-the-fly config**

A problem has been fixed so that ARP information associated with a newly-configured or reconfigured interface gets set correctly.

**- Statistics displays corrected**

The **maint 15** command is replaced by **maint 16** and output from **grstat**.

**- Logical interface configuration not restricted**

A problem has been fixed in which interfaces other than `ga0s0` and `ga0s1` on an ATM OC-3c card could not be configured and utilized.

**- Invalid ARP packets no longer cause panic**

Invalid ARP packets no longer cause the ATM OC-3 media card to panic.

**ATM OC-12c****- Advisory: MTU size**

A problem has been fixed which had prevented the ATM OC-12c MTU from being set to the default (9180).

**- Cards operate in any GRF 1600 chassis slot**

A problem has been fixed in which an ATM OC-12c card could not reside in slot 0 in the GRF 1600.

**Frame Relay****- Frame Relay comes up right after card reboot**

A condition has been fixed in which Frame Relay interfaces would be held in the down state for up to 40 seconds after the media card was reset.

**- On-the-fly config of physical links supported**

The current Frame Relay implementation supports on-the-fly configuration of physical links using the **grfr -c ccp** command.

**- Must restart fred to see physical link changes**

A condition which required users to restart the Frame Relay daemon, **fred**, for configuration changes to take effect has been corrected.

**- lo0 setting not required**

A Frame Relay requirement that the interface lo0 in `/etc/grifconfig.conf` must be set to 127.0.0.1 has been removed.

## SONET

**- SONET: maint 82 displays correct data**

The **maint 82** command now only returns valid transmittal and other debugging data.

## System

**- ARP entries sent by default**

To ensure that ARP entries are sent during on-the-fly configuration of ATM virtual circuits, a change was made to **gratm** to invoke **grarp** by default unless the **-a** option is specified.

**- grflash command loads images onto media cards**

The **grflash** command loads new code into media card PROM and is now supported on all media cards. DO NOT use the **grburn** command.

**- GR-II media cards reboot normally**

A problem has been fixed in which one or more media cards would remain in the Configuring state after a system boot and would need to be reset to get them to the running state.

**- route delete command works correctly**

A problem has been fixed in which manually deleting a route not in the route table could cause a wrong route to be deleted

**-Spurious grinchd.conf request deleted**

The remnant of an earlier configuration file, `/etc/grinchd.conf`, no longer appears after a 1.3 GRF system is upgraded to a 1.4 release.

**- grifconfig.conf will parse a second hyphen**

The **grifconfig** parser no longer misinterprets a second hyphen used on the statement line in place of optional arguments.

**- If the NFS connection fails, syslog does not back up**

A problem has been fixed in which the loss of the NFS connection caused **syslog** to back up while waiting for the connection to re-establish. Now, if the connection breaks, **syslog** drops packets as needed and does not fill memory.



**- Spurious GRF 1600 state message deleted**

The following superfluous card state message is no longer sent to the console:

```
Sep  2 11:16:44 grf1 kernel: atmp_ifaddr_cardq_load: Loaded 0
```

**GateD****- netstat -rn only shows active interfaces**

The route table no longer contains routes for down interfaces.

**- GateD supports IS-IS adjacencies over FDDI**

A problem has been fixed that caused incomplete IS-IS adjacencies to be formed over FDDI networks and that resulted in the network device not installing routes received in the LSPs (Link State Packets).

**- OSPF functions through aliased interface**

A problem has been fixed in which the OSPF protocol would not function properly through an interface aliased with secondary IP addresses.

**- GateD starts correctly**

An intermittent problem that would cause GateD to not start correctly or not function properly has been fixed. In this condition, GateD required a directly-attached interface to be active or, during boot, an active interface with an IP in the network defined on the gateway statement.