
GRF 1.4.10 Release Notes

August 25, 1998

These release notes describe the 1.4.10 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
- 1.4.6 General release January 30, 1998
- 1.4.8 Limited release June 30, 1998

New 1.4.10 Addendums

Two new addendum accompany release 1.4.10. The *GRF 1.4.10 Addendum* contains revised configuration documentation for the new features in the release as well as updated information for bridging, ECMP, Frame Relay and ATMP. The *GRF GateD 1.4.10 Addendum* contains significantly updated sections for the GateD Configuration Statements, the GSM utility, and the `/etc/gated.conf` file. These addendum are available from the ftp site as the files `Addendum.pdf` and `gated.pdf`.

Note: certain media cards not supported

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.

Contents

Upgrading system software	2
New features in 1.4.10	4
Advisories	7
Known problems in 1.4.10	8
ATMP notes	8
ATM OC-3 notes	9
ATM OC-12 notes	11
CFMS notes	13
Ethernet notes	13
FDDI notes	14
Filtering notes	14

Frame Relay notes	15
HIPPI notes	17
HSSI notes	18
SONET notes	18
SNMP notes	18
SPD notes	19
System notes	19
GateD advisories	24
GateD notes	27
GRDIAG	28
Problems fixed in 1.4.10	29
ATM OC-3	29
ATMP	29
ECMP	29
Frame Relay	29
SONET	30
SP Adapter card	30
System	31
GateD	32

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 1.4.10 code, CFMS code, release notes (RN), and GRF 1.4.10 Addendum files:

- 1.4.10.ibm.TAR.gz
- 1.4.10.ibm.root.gz
- CFMS.TAR.gz
- RN1_4_10.pdf
- RN1_4_10.txt
- addendum.pdf
- gated.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_10
```

Set the file format and download the files:

```
ftp> bin
ftp> get 1.4.10.ibm.TAR.gz
ftp> get 1.4.10.ibm.root.gz
```

Change directory to install the remaining files in /flash so they are not removed when the **flashcmd** executes:

```
ftp> lcd ..
ftp> get CFMS.TAR.gz
ftp> get RN1_4_10.pdf
ftp> get RN1_4_10.txt
ftp> get addendum.pdf
ftp> get gated.pdf
ftp> quit
# cd /
```

Use **grwrite** to save any configuration changes made since last boot.

```
# grwrite
```

The next commands unpack zipped files and clean up flash:

```
# grfins --source=/flash/tmp --release=1.4.10.ibm --activate
# flashcmd -w rm -rf /flash/tmp
# umountf
```

Reboot installs the files on RAM, enter:

```
# shutdown -r now
```

or use

```
# reboot -i
```

When the router is rebooted, it should display that it is running kernel version 1.4.10.

New features in 1.4.10

More information about each of these features is available in the *GRF 1.4.10 Addendum* that accompanies this release. The *Addendum* is available in PDF format in the ftp release directory.

- Field diagnostic capability - **grdiag**

The **grdiag** command enables site and field personnel to verify media card hardware. **grdiag** performs internal Built-in Self Test (BIST) level diagnostics. A media card that fails this **grdiag** must be replaced. You can run **grdiag** on one or more cards without interrupting other media cards or system operation.

The diagnostics verify the following media card and slot functions:

- all memory
- all media hardware logic (media card)
- all serial hardware logic (serial daughter card)
- the connection between the slot and the switch
- the connection between the slot and the communications bus
- the connection between the slot and power delivery

The diagnostics do not exercise the physical interfaces or transceivers.

- Data collection utility - **grdinfo**

The **grdinfo** utility enables the site to use a single command to collect a comprehensive set of debug and configuration information for the GRF. **grdinfo** options specify the type of information collected, including logs, dumps, media card statistics, protocol statistics, and control board data. Target data can be obtained at the system level or at the card level.

You can execute **grdinfo** while the GRF is running although there will be an impact on performance while the information is collected. Data is saved in compressed TAR files for ease in transfers to Technical Support.

- Support for larger external devices

Prior to this release, the only fully supported external PCMCIA options were the modem and 85MB flash devices. The **iflash** command now determines the geometry of a 85/175/520MB PCMCIA ATA Type-II or Type-III disk device, including spinning disks, and formats the device for use in a GRF.

Ascend supports the following ATA-compliant drives for GRF operation:

- Kingston Datapak 520MB, P/N CT520RM
- Sandisk 175MB Flash, P/N SDP3B
- Sandisk 85MB Flash, P/N SDP3B-85-101
- Aved 85MB Flash, P/N AVEF385MB25ATA501

Ascend offers only the Sandisk 85MB Flash directly (part name GRF-AC-FLASH). Customers may purchase the PCMCIA devices listed above through an external source. PCMCIA devices other than those listed above are not supported by Ascend.

- Remote configuration management via CFMS

The Configuration File Management System (CFMS) tool enables network administrators to store and track changes to host and router configuration files.

CFMS runs on a UNIX server attached to client routers via a secure administrative LAN. CFMS maintains a Concurrent Versions System (CVS) repository on the CFMS server of all configuration files for all hosts under CFMS management. An administrator can add and remove hosts from CFMS management and change the list of files tracked for all hosts. CFMS code is available in a separate gnu zipped tar file (CFMS.TAR.gz) in the 1.4.10 release directory.

- Interface monitoring utilities: Threshpoll and Pinglog

Two new utilities support interface-level monitoring and management.

The Threshpoll utility enables the network administrator to specify various statistics and/or states to be monitored and then compared to a user-specified threshold. If the threshold is met, exceeded, or not reached, a unique trap signaling the problem is sent to a target SNMP manager/gateway. Results of the monitoring are available in Threshpoll logs. These are the some of the statistics and states that can be monitored:

- length of time system has been up
- octets received or transmitted
- unicast packets received or transmitted
- packets dropped on receive side
- packets dropped on transmit side
- errors occurring on receive side
- errors occurring on transmit side

The Pinglog utility enables the network administrator to ping GRF interfaces and attached network devices to determine if the interfaces are alive and responding. If a GRF interface does not respond, a trap is sent to a target SNMP manager/gateway. Detected events are available in Pinglog logs.

- Frame Relay Switching and on-the-fly configuration

Frame Relay switching, NNI, is now implemented on the GRF. The Frame Relay switching feature enables a GRF to function as a switch. When a GRF router functions as a Frame Relay switch, it performs layer-2 switching and forwards incoming data from incoming circuits to the appropriate out-going circuits without touching the payload of the data packets.

Frame Relay switching is supported on the HSSI and SONET media cards. The incoming and outgoing endpoints of a switched circuit can be configured on different media types (HSSI and SONET). A complete Frame Relay network can be built by connecting GRF routers using high-speed links.

Also, Frame Relay configuration has changed significantly from 1.3.

You can now add and delete Frame Relay links, or modify a link parameter, without resetting the media card. You edit the `/etc/grifconfig.conf` and `/etc/grfr.conf` files as needed, and then issue `grfr` commands to activate the changes.

You can add or delete PVCs without resetting the media card by editing the `/etc/grfr.conf` file and then using **grfr -c ccp** to create and install the PVC or **grfr -c crp** to remove the PVC.

- GRF ATMP provides pre- and post-fragmentation options

The GRF home agent encapsulates packets received from the home network by adding the GRE and a second IP header. If the resulting packet is too large to transmit as a single datagram over the connection media, the GRF must fragment these packets. Users can specify that fragmentation be done either before encapsulation (pre-fragmentation) or after (post-fragmentation).

- GRF ATMP supports RIPv2 route transmission

GRF home agents can use RIPv2 to advertise routes for registered mobile nodes to the home networks (the home network routers must also run RIPv2).

The ATMP daemon, **aitmd**, can be configured to multicast RIPv2 packets to the home networks attached to a GRF home agent. These packets advertise routes to each mobile node tunnel registered on the home agent. RIPv2 transmission enables a home network to learn the paths to its mobile nodes. This capability supports the configuration of a GRF as a secondary home agent. RIPv2 transmission is enabled in the home network record of the `/etc/aitmd.conf` file.

An ATMP tunnel “connects” a foreign agent with a specific home agent which, in turn, has a dedicated link to the target home network. If the home agent fails, the tunnel also fails. The mobile node is disconnected and cannot connect to the home network until the home agent is back or the foreign agent negotiates another tunnel. A secondary home agent can be a backup for the primary home agent. If the primary home agent fails, a new tunnel is created to the specified secondary home agent. The home network uses the RIPv2 routes to determine a path to the mobile node.

New GateD features

No new GateD features for 1.4.10.

Advisories

- New Frame Relay configuration tasks in 1.4

Frame Relay configuration has significantly changed since release 1.3.

In release 1.4, the configuration of Frame Relay links and PVCs require **grfr** commands. Display and on-the-fly configuration capabilities of **grfr** commands are also different. New configuration commands and examples are described in the Frame Relay chapter of the *GRF 1.4.10 Addendum*.

- 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 Technical 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.

Known problems in 1.4.10

ATMP notes

Advisory: major changes in ATMP configuration parameters

New parameters are added to support new ATMP features. The “circuit” section is replaced by an “interface name” entry. However, backward compatibility is maintained and `circuit` entries will be parsed. BUT — if you leave `circuit` entries in the configuration file, DO NOT include the new `interface name` entry, the results are unpredictable. The `netmask_size` parameter is no longer used and is now ignored.

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, which is larger than the Ethernet MTU of 1500 bytes. The ATM card must fragment the packet before sending.

This slows down 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

Advisory: SVCs are not supported in this release

ATM SVCs are not supported in release 1.4.10.

- PVC with `proto=vc` loses its ARP entry

A PVC specified with `proto=vc` loses its entire ARP entry, and no ARP entry is observed for that PVC in the `maint 8` output. The ARP entry is restored only after performing `gratm` on the interface, as in `gratm ga0yx`. The loss of the ARP entry for PVC=`vc` is observed to happen intermittently.

- Card with full configuration of PVCs is constantly reset

An ATM OC-3 media card that has been configured with 512 PVCs resets under various conditions including 1) periods of high activity; 2) when it receives an invalid VPI/VCI pair; and 3) occasionally during periods of no activity. Users can work around this by configuring no more than 256 PVCs on a ATM OC-3 card.

- Inverse ARP does not resolve deleted entries

A manually-deleted ARP entry (**grarp -i gaxxxx -d xx.xx.xx.xx**) on an ATM-OC3 card then fails to be re-resolved without resetting the card. Manually configuring the ARP entry in `/etc/grarp.conf` works around the problem.

- Inverse ARP fails to resolve alias IP address

Inverse ARP on an ATM OC-3c card fails to resolve the alias IP address of its partner. Inverse ARP can only resolve the peer IP address for its primary address.

This can be worked around by manually configuring the ARP entry for the alias in the `/etc/grarp.conf` file.

- gratm incorrectly lists dest_vc as optional

Instructions in `/etc/gratm.conf` state that the `dest_vc` is an optional parameter for raw PVC. However, **gratm** does not correctly process a `proto=raw` PVC configured without the `dest_vc` parameter. The raw PVC is not set up correctly, and the following error message appears in `grconslog`:

```
Invalid raw request to port 133: VC= 0/0
```

Users should include both `dest_if` and `dest_vc` as mandatory parameters for raw PVCs.

- IP point-to-point function fails

After a user creates a point-to-point ATM OC-3c interface by entering the peer address under the `broad_dest` field of `/etc/grifconfig.conf`, the **ifconfig interface** command shows a different peer address for that interface.

- maint 3 command failed to list all interfaces

The ATM OC-3c `maint 3 1 0` command did not list the entire set of interfaces configured on the card. A **ping** to these interfaces and their directly-attached interfaces succeeds.

- Multicast fails on ATM when PVC proto=llc

Attempts to send multicast packets over an ATM PVC specified with `proto=llc` fail. However, if the the PVC is changed to `proto=ip`, multicast sends work.

- 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 no more than 75 logical interfaces per ATM/Q card.

- 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
```

- 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.

Advisory: 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

- IP point-to-point function fails

After a user creates a point-to-point ATM OC-12c interface by entering the peer address under the `broad_dest` field of `/etc/grifconfig.conf`, the **ifconfig interface** command shows a different peer address for that interface.

- Inverse ARP fails to resolve alias IP address

Inverse ARP on an ATM OC-12c card fails to resolve the alias IP address of its partner. Inverse ARP can only resolve the peer IP address for its primary address.

This can be worked around by manually configuring the ARP entry for the alias in the `/etc/grarp.conf` file.

- Data problem with ping packets > 9788

Pinging between GRF systems across an ATM OC-12c link produces data errors on **ping** sizes greater than 9788 bytes.

- PVC with proto=raw fails

Packets received from the input side of a raw PVC were not transmitted to the output side of raw PVC.

- Card panics while coming up after reboot

The ATM OC-12c cards PANIC while coming up after the GRF was rebooted under no load. This occurs most of the time.

- Card may issue "BOOT ME after system reboot

While the system was coming up, an ATM OC-12c 1 issued a BOOT ME request and started dumping. The console and **grconslog** reported no associated errors. This occurs about half the time.

- Certains pings to attached interfaces fail

Pings to interfaces directly attached to ATM OC-12c PVCs specified with `proto=ipnllc` have failed.

- 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 up to 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.)

- gratm erroneously allows ATM-12 PVC proto=vc

ATM OC-12 does not support `proto=vc` encapsulation. Only protocols that use an LLC/SNAP header are supported. 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
```

CFMS notes

- CFMS runs only from installation directory

CFMS commands can only be issued when the user is in the CFMS installation directory. Change to that directory to work with CFMS. Add ./ to the front of CFMS commands given that you may not have the current directory in your path. In this example, CFMS is installed in directory /opt/cfms:

```
$ cd /opt/cfms
$ ./cfms list
$ ./cfms add_host <host>
```

- CFMS commands ignore extraneous arguments

CFMS commands frequently accept extraneous arguments without generating error messages to user.

Ethernet notes

- Ethernet card not auto-negotiating speed correctly

GRF Ethernet ports may not correctly read the negotiation information coming from some other vendor's Ethernet chip set. As a result, it does not correctly negotiate and sense the speed or duplex of the partner unless both are set to autonegotiate.

- Ethernet ICMP directed broadcast not working

Ethernet ICMP directed broadcast is not working in this release. Note that UDP directed broadcast does work.

- Ethernet 100 full/half duplex exhibit high packet loss

A test on Ethernet media cards running at 100 full and 100 half duplex modes resulted in excessive instances of 100% packet loss. They occur more frequently as packet size increases.

- 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 1.4
```

Using the **grreset** 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

Advisory: Packet logging documentation needs clarification

A statement about packet header logging in the filtering chapter of the *GRF Configuration Guide 1.4* 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.

A null permit matches all packets

A implicit deny filter with a "null" permit block allows all traffic to be passed. A permit block can become "null" if the user inadvertently comments out all entries within it, as in:

```
filter denyall {
  implicit deny:
  permit {
    #insert commented out filter here}
  }
}
```

In other words, if all entries for the filter are commented out, that is equivalent to not doing any filtering. Thus, the implicit deny will not be invoked.

- Filtering documentation omits SONET and Ethernet support

The Filtering chapter in the *GRF Configuration Guide 1.4* 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

- Frame Relay configuration process changed from 1.3 to 1.4

Frame Relay configuration has significantly changed since release 1.3. In 1.4, the configuration of Frame Relay links and PVCs require **grfr** commands. Display and on-the-fly configuration capabilities of **grfr** commands are also different. New configuration commands and examples are included in the Frame Relay chapter in the *GRF 1.4.10 Addendum*.

- Frame Relay switch config files not backward compatible

Frame Relay switching configuration files are not backward compatible between 1.3 and 1.4 software releases. Backing out to a previous release requires manual intervention by the user,

the process is not automated. User must make backup copies of all GRF configuration files (normally `/etc/*.conf`) prior to any upgrade.

- pings may fail with 128 HSSI Frame Relay interfaces

When 128 Frame Relay interfaces are configured on one HSSI port, the user may not be able to ping through all of them. Although the ARP table shows the ARP entry for the other side, the message in **grconslog** is that Arp Entry Not Found.

This condition occurs when interfaces are configured on-the-fly and/or after resetting card. The interfaces are cleared up by rebooting the system.

- HSSI port 0 and 1 interfaces required for FR switching

Interface 0 must be configured on port 0 for HSSI Frame Relay NNI interfaces configured on port 1 to operate. Without an interface defined in `/etc/grifconfig.conf`, the Frame Relay switching links and circuits defined on those ports do not become active. The interfaces for **fred** to talk LMI with the cards are created and active, but the actual circuits the user defines stay in the "Inact" state.

The entry can be either an active or dummy interface. For interface `gs0380` to work, you must define the following interfaces in `/etc/grifconfig.conf`:

```
# name address netmask broad_dest argument
gs0380 192.168.0.1 255.255.255.0
gs030 - - - up
```

If the media card has already booted, then interface 0 must be added to the `/etc/grifconfig.conf` file and the **fred** daemon must be restarted. This requirement applies only to HSSI Frame Relay switching.

- HSSI FR circuit under heavy load can lose LICs

When a HSSI Frame Relay circuit is under a heavy load, it can lose LICs data. 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.

- grfr -c ccl does not restore PVCs with restored link

When a link is removed using the **grfr -c ccl** command, the link and all underlying PVCs are removed. A subsequent **grfr -c ccl** command restores the link but not the PVCs. Use the **grfr -c ccp** command to restart the PVCs.

- fred does not support all documented options

Two Frame Relay command options, **grfr -c dss** and **grfr -c dls**, have been replaced by two other command options, **grfr -c dsc** and **grfr -c dlc**. The **grfr -c dss** and **grfr -c dls** options have not as yet been removed from the **grfr** man page.

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: controlling PVCs autoadded by a remote system

In the Link section of `grfr.conf`, the `AutoAddGrif` option can be selected (set to `Auto`) to allow the local GRF to add a new circuit to an interface when it detects through CMI that a remote system has added a circuit. 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 ccl** 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 five 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 five times to configure the five PVCs.

- 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

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***- maint 133 returns incorrect IP statistics**

The HIPPI card is not keeping IP statistics counters correctly. As a result, output from **maint 133** reports incorrect numbers of packets received and forwarded.

- tcpdump does not capture packets on HIPPI

The **tcpdump** utility is not capturing and counting packets on HIPPI interfaces.

- 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

- Interface goes inactive with 230+ DLCIs configured

A HSSI logical interface configured with Frame Relay and more than 230 DLCIs becomes non-operational. Limiting the number of DLCIs on any given HSSI logical interface to less than 230 will avoid the problem.

- 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 count high enough to make it less likely to drop all keepalives and lose the 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**.

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).

SNMP notes

- SNMP changes do not take effect

Changes to the `/etc/snmpd.conf` configuration file do not take effect when **snmpd** is restarted.

A workaround is to reconfigure the SNMP agent using this procedure: and save.

1. Make your changes as needed in `/etc/snmpd.conf` and save.
2. Remove the `/var/run/snmpd.NOV` file.

```
3. Restart snmpd by using grep to obtain the PID and then restart the daemon:  
# ps -ax|grep snmpd  
# kill -9 <PID>
```

SPD notes

- ATM OC-3c can lose high priority packets

On SPD settings between 1 and 90, the ATM OC-3c media card has a high loss of the high priority packets. No packet loss for high priority traffic is detected above an SPD setting of 90.

- SPD settings cause FDDI to lose high priority packets

FDDI loses all high priority traffic with SPD settings of 73 and 100. We should see no traffic losses at these settings.

- Error not reported when SPD assigned to ATM OC-12c

The CLI should return an error if the user tries to set an SPD value on the ATM-OC12 card using a CLI profile setting. On this card, SPD is under hardware control.

- Configuration Guide incorrectly indicates SPD on ATM OC-12c

The ATM OC-12c chapter in the *GRF Configuration Guide 1.4* incorrectly indicates that SPD applies to the ATM OC-12c media card. You cannot configure SPD on that card.

- SONET can lose high priority packets

SONET shows a small loss on high priority traffic on any SPD setting. When SPD was set to 100 to allow only high priority traffic, the SONET also passed low priority traffic.

- SPD causes SP Adapter card to lose packets

The SP Switch Router Adapter card shows 14 to 15 percent packet loss on high priority traffic at any SPD setting. Due to its high traffic rate, this card should not operate with SPD.

System notes

Advisory: grifconfig can reconfigure addresses on-the-fly

A 1.4.3 release note stated that the **grifconfig** command could add/delete interfaces on-the-fly. This is incorrect. The note should read as follows:

Users can add/change interface *addresses* on-the-fly by using the **grifconfig** `<interface_name>` command. This command takes the place of the **ifconfig** command for assigning, changing, or removing IP addresses. (The use of **ifconfig** for these tasks has been deprecated. Refer to the **grifconfig** man page for more information."

- grifconfig incorrectly sets address of point-to-point link

After a user deletes an interface using **ifconfig interface delete**, using **grifconfig** to recreate the interface does not correctly set the IP address of the other end of a point-to-point link.

- Spaces cannot follow a grinch command

Do not follow a **grinch** command with spaces. This command fails because of the spaces before and after the equal (=) sign:

```
# grinch -p0 2.12.2.1.15.3.1 = 10000
grinch: No value to set.
Ignoring ''.
2.12.2.1.15.3.1: 10
10000: no such item
```

Without the space, the command executes properly:

```
# grinch -p0 2.12.2.1.15.3.1=100000
2.12.2.1.15.3.1: OK
```

- Route table memory size can be exceeded

By injecting more active routes than the media cards are able to install, it is possible to exceed the hard-coded limit on route table memory size.

- adduser does not create files to new home directory

The **adduser** program does not correctly create proper home directories for new users. Administrators can either manually create the home directory for the user or modify the **perl** script `/usr/sbin/adduser` so that it will add all file types. Modification to the **perl** script is required to copy all file types until this problem is corrected in software.

1. Make two changes to `/usr/sbin/adduser`:

- modify line 476 to add a space between the " *.", so change " *.* " to " * .* "
- change

```
while ($file = <*. *>) {
```

 to

```
while ($file = <* .* >) {
```

2. Copy designated files (`.cshrc`, `.profile`, `.login`, etc) to directory to `/usr/share/skel`

This copies files from the `skel` directory to the user's home directory when executing "adduser"

3. Then execute this series of commands:

```
# cd /usr/share
# grsite --perm skel          ## to save directory and contents outside
                             /etc directory over upgrades
# grsite skel                ## to save directory and contents outside
                             /etc directory over a system boot
# cd /usr/sbin               ## this part will not be needed when problem is fixed
# grsite --perm adduser      ## to save files outside /etc directory over upgrades
# grsite adduser             ## to save files outside /etc directory over a system boot
```

4. Run **adduser**

- grarp reports incorrect link-level address for HSSI Frame Relay

The **grarp** command does not report the link-level addresses (DLCI) for each IP address configured on HSSI Frame Relay, it reports "incomplete" as shown in this excerpt:

```
# grarp -i gs010 202.5.1.152
gs010 (26): 202.5.1.152 at (incomplete) permanent
gs010 (26): 222.5.1.152 at (incomplete) permanent
# grarp -i gs020 202.5.2.1
gs020 (31): 202.5.2.1 at (incomplete) permanent
```

- finger causes GRF to dump core if not on remote host

Running the **finger** command to find a user on a remote box will cause **finger** to abort if the **finger** daemon is not running on the remote box. If the user does **finger @<ip_address>** command, and the host responding to **ip_address** refuses the request, the **finger** program on the GRF dumps core.

- 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.

- 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. These options support large external PCMCIA devices.

- 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.

- 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 umount 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.

- 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

- Using OSPF to advertise interface address(es)

By default, OSPF does not advertise the addresses of any atmp interfaces. To get OSPF to advertise these addresses you must add a stubhosts entry to the appropriate OSPF area definition, as shown in the following example:

```
ospf yes {
    backbone {
        stubhosts { 10.1.19.45 cost 40; };
    }
}
```

The IP address in the stubhosts entry is the address assigned to the atmp interface you want to advertise.

- Configurable Export Best BGP (CEBB) feature is documented

The documentation that describes the CEBB feature is available in the *GRF 1.4.10 Addendum* in the GateD Configuration Statements Chapter. The **disable export best** keyword (for disabling this functionality, which is on by default) in the `/etc/gated.conf` file is also documented.

- Equal Cost Multipath (ECMP) statement documented

The documentation for configuring the EMCP statement in the `/etc/gated.conf` file is available in the *GRF 1.4.10 Addendum* in the GateD Configuration Statements Chapter (in the Definitions statement subsection).

- Assign an alias as routerid

We strongly recommend that you configure an alias address on the loopback interface of your router and use this address as the routerid in . Also, the alias should be used as the local address (**lcladdr**) in peer statements for IBGP. This enables proper route calculations across all interfaces.

To configure an alias of the loopback interface, add an interface definition to `/etc/grifconfig.conf`, as shown in the following example. 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 `/etc/grifconfig.conf`:

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

Using this configuration, executing the **netstat -in** command 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
```

Using this configuration, executing the **ifconfig lo0** command produces the following output:

```
# ifconfig lo0
lo0: loop flags=8009<UP,LOOPBACK,MULTICAST> mtu 1536
```



```
inet 127.0.0.1 netmask 0xffffffff00
grit 0:0x48:0
```

- GateD trace option

Tracing significantly decreases GateD speed and efficiency. Unless dynamic routing protocols are being debugged, GateD trace options should be turned off by using one of the following mechanisms:

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

- GateD multiexit discriminators (MEDs)

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

The following `/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

localas

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

EGP

BGP 2

SNMP

- Redirect statement not supported

This version does not support the **redirect** statement. in the `/etc/gated.conf` file.

- No IGP group support

This release of GateD does not support the BGP group type **igp**. Use of the BGP group type **routing** is recommended for internal peers.

- 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.

- Avoid enabling warning messages with BGP

If BGP is running with warning messages enabled, excessive scrolling to the console can be encountered. In general, it is recommended that the **show-warnings** keyword in 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)

- BGP Destination Preference Attributes not supported

Beginning with the 1.4.6 release, BGP destination preference attributes are no longer supported.

- 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 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.

GateD notes

- Gated cores if interface lo0 not defined

The default configuration of a GRF does not have the lo0 interface defined. To start GateD on a system running the 1.4.10 release, it is necessary to define the lo0 interface. Defining this interface must be done whether you are upgrading an existing GRF or installing a new GRF system. Define the following interface in `/etc/grifconfig.conf`:

```
#name      address      netmask      broad_dest      arguments
lo0        127.0.0.1    255.0.0.0
```

- metricout does not work on external BGP group peer statement

GateD allows you to specify the **metricout** keyword on a peer statement in an **group type external** definition, as shown in the following example:

```
group type external peeras 5 {
    peer 10.1.1.1 metricout 45;
}
```

The **metricout** keyword is ignored. All peers within the same peer group must receive the same **metricout** value for any particular route. This value must be specified with the **group type external** definition, as shown in the following example:

```
group type external peeras 5 metricout 45 {
    peer 10.1.1.1;
}
```

- Multicast RIPv2 advertisements on multiple interfaces not multicast

When GateD is configured so that RIPv2 is active for multicast route advertisements on multiple interfaces, RIPv2 advertisements only get sent on one of the configured interfaces.

- ECMP multipath change requires GateD restart

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

- 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.

- 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 ;
```

- GateD local synchronization not supported

Under certain heavy-load conditions, GateD can, 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`, goes down and then comes back up, GateD does not recover the routing policy associated with that interface.

The policy can be recovered by manually executing the **gdc reconfig** command, 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 can influence which routes GateD selects. This will be reviewed and changed, or documented, in a future release.

- 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 GRF appears to have “gone away” and is not considered to be receiving incoming traffic.

GRDIAG

- Cards did not return to running state after running GRDIAG

After running GRDIAG, the cards did not return to the running state but stayed in the configuration state and continued to reboot. To solve this problem, start `ncli`, read `CARD X` (where X is the card that does not boot correctly), `cd load`, and delete `hw-table/x` (where x is the card configuration type, for example, ethernet).

Problems fixed in 1.4.10

ATM OC-3

- PVC on-the-fly locks up ATM/Q port with card under load

A problem has been fixed in which an ATM/Q port would lock up while receiving data via a PVC configured on-the-fly. This situation would occur when the media card was under load.

- Adding on-the-fly PVC would hang card

A problem has been fixed in which adding a PVC on-the-fly could hang the ATM OC-3c media card. If the card did not recover, it would have to be reset.

- Invalid PCR or SCR settings allowed by gratm

A problem has been fixed in which users would not be notified when they configured an invalid setting for PCR or SCR. Invalid settings were not caught by **gratm** and caused 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.

ATMP

- ATMP accepts tunnels regardless of circuit state

A problem has been fixed in which ATMP would continue to accept tunnels when the circuit to the home network was down. No error message would be returned to the user.

ECMP

- ECMP continues to forward down links without link0

A problem has been fixed in which ECMP incorrectly continued to forward packets. Even though the cable had been pulled from a downstream card in ECMP, and `link0` was not there any longer, ECMP would continue to forward packets to an interface on that card.

- ECMP could corrupt kernel

A problem has been fixed in the way ECMP handles deletions on ECMP routes which would cause ECMP to corrupt the kernel.

Frame Relay

- Frame Relay not aware when media card is removed

A timing problem with the Frame Relay daemon has been fixed that caused **fred** to be unaware that a HSSI media card or cable had been removed. Any configured PVCs remained active, and would resume activity if the card was replaced.

- Power cycle required for PVC to be recognized

A Frame Relay PVC is configured and an IP address is added to the interface on which the PVC is assigned. After the card is reset, the **grfr -c dpc** command does not show the PVC. You must reset the GRF to have the PVC recognized.

This is fixed by use of **grfr -c ccp** command to restart the PVC.

SONET

- Cannot set MTU on SONET HDLC card

A problem has been fixed in which the MTU could not be set on a SONET card configured as HDLC. Using the **ifconfig** command and modifying the `/etc/grifconfig.conf` file both fail. as a result, the MTU would default to 1500.

SP Switch Router Adapter

- SP switch router adapter hangs when OSPF is enabled

A problem has been fixed where enabling OSPF between an SP Switch Router Adapter and an IBM SP node over an IBM SP switch hangs the SP Switch Router Adapter.

- ARP disabled configuration does not work

ARP disable appears to be disabled using IBM **Enadmin** tool, but ARP remains enabled. A workaround is to manually edit the `/etc/grdev1.conf` file to change the switch ARP entry (1 or 2). To activate the change, reset the card or run the **grinch** command.

This problem has been fixed.

- Traffic over card hangs when using OSPF

When OSPF is running on the GRF, **ping** and all other IP traffic over the SP Switch Router Adapter card gradually slows until the card hangs and becomes inoperable. Output from **gsm** shows the OSPF peering state to be either INIT, IDLE, or DOWN, but never FULL.

This problem has been fixed.

- pings fail when ARP is disabled

When ARP is disabled on an SP Adapter card, you cannot ping its local interface. If ARP is enabled, there is no problem.

This problem has been fixed.

- SP card does not handle loss of switch node

If the switch node to which an SP Switch Router Adapter card is attached becomes administratively fenced, the card does not successfully transition. The card's switch port begins to issue parity errors. The SP Switch Router Adapter card may panic or also be fenced.

The SP Switch Router Adapter card did not correctly process the new SP switch routing table coming from the primary node, this problem has been fixed.

System

- Date command works correctly

A problem with the UNIX **date** command is fixed. Unless you execute these commands in single user mode, **ntpdate** and **date** freeze the GRF. Date changes to the GRF system **MUST** be made by rebooting and interrupting the boot process, entering **boot -s**, and setting the date while in single user mode. Please use caution when changing date and time settings.

All GRFs ship to customers with date and time settings in place. Change these settings only when the GRF is to operate in a different time zone. During initial configuration you can specify a different time zone in the **netstart** script. To reconfigure after initial configuration and specify a new host name or a different time zone, use the **config_netstart** command to initiate this script.

- Directed broadcast could cause cards to panic

A problem has been fixed in which directed broadcast packets could cause a media card to panic or dump. Directed broadcast packets coming into HSSI, SONET, and Ethernet media cards had the potential of causing the ingress card to generate a bad packet on the switch backplane, possibly causing an unrelated card to panic or report an error.

- Directed broadcast packets are not forwarded

A problem has been fixed in which GRF media cards would not correctly respond to and did not forward directed broadcast packets destined for networks to which the GRF is directly connected. This problem only applied to broadcasts to directly-connected networks. If there was another hop to the broadcast target, then the GRF would forward directed broadcast packets correctly.

- TACACS+ now supports encryption

In this release of 1.4, the TACACS+ implementation does support encryption.

- Route table mismatch can occur

A problem in which route mismatches could occur and the media card would continue to forward packets has been fixed.

- grifconfig will delete interface addresses

With the new and improved **grifconfig**, it is possible to remove an IP address from an interface without rebooting the system.

GateD

- Cannot restrict routes based on AS path

When using the **restrict** parameter in an export policy statement for **proto bgp aspath**, the parameter is ignored and the routes are still exported.

- Route filters on import policy statement are ignored

When specifying routes filters in an import policy statement, the route filters are not being used, and all routes are being imported.

- Cannot restrict routes for exporting to BGP based on protocol

When attempting to use the **restrict** keyword in the source portion of a BGP export policy statement (that is, **proto static restrict**;, **proto direct restrict**;, and so on), the desired routes are not restricted and are exported.

- Specifying 'export proto ospf tag XX' does not set tag in route

When attempting to set tag values in OSPF AS external routes by using **export proto ospf tag XX** (where **XX** is the desired tag value), the tag is not properly set for the desired route(s).

- Configurable Export Best BGP option not working properly

The ability to control whether the CEBB behavior is on or off within GateD is not working properly. CEBB is on by default and cannot be disabled. Also, the configuration statement for turning CEBB off is not documented.

- Not able to reference BGP subgroups in "exports" to other AS's

When attempting to configure BGP subgroup values in the source portion of an export policy statement, parsing errors would occur for `/etc/gated.conf` when they should not.

- Unwanted routes announced for BGP subgroups specified in export policy

When a BGP subgroup is configured in the source portion of an export policy statement, all routes corresponding to all BGP subgroups are announced instead of just the configured BGP subgroup(s).

- Router Discovery Server requires 'address all ignore' for broadcast interfaces

When Router Discovery Server is configured so that Router Discovery advertisements are broadcast (instead of multicast, which is the default) on an interface, no advertisements are actually sent on the interface unless the **routerdiscovery server** statement contains the sub-statement **address all ignore** before any other address sub-statements.

- Configuration of OSPF over atmp0 interface causes GateD to abort

When attempting to configure OSPF over an atmp0 interface, which is not defined as a loopback interface, GateD aborts.

- Specifying AS path regular expression in export policy causes GateD to abort

When specifying AS path regular expressions in export policy that contain multiple parameters, GateD aborts.

- /etc/gated.conf parse error occurs when specifying null in AS path

When specifying AS path regular expressions in export policy that contain multiple AS numbers, parsing errors occur and GateD does not restart.

- /etc/gated.conf parse error occurs when specifying multiple AS path numbers

When specifying AS path regular expressions in export policy that contain an or (|) operator, parsing errors occur and GateD does not restart.

- /etc/gated.conf parse error occurs when specifying | (or operator) in AS path

When specifying AS path regular expressions in export policy that contain an or operator (|), parsing errors occur and GateD does not restart.

- nexthopself configuration on BGP peer not working correctly

The **nexthopself** keyword when configured for a BGP peer should cause the next hop for routes learned from this BGP peer to be set to the local router's interface address. GateD was not performing this correctly.

- If telnet to GRF is terminated without first exiting GSM, GateD aborts

If a telnet session, which is actively running GSM, is terminated without first exiting from GSM, GateD aborts.

- Selective community tagging of BGP export routes no longer works

When using the **mod-aspath** parameter on export policy for BGP routes to modify community tags, the BGP routes are not being modified as intended.

- /etc/gated.conf parse error when using "0/0 restrict"

Route filter **0/0 restrict** in **define**, **export**, **import**, and other statements where route filters can be defined, a parse error occurs.

- OSPF ASE routes for "overlapping" addresses are not installed

When two OSPF ASE routes are received that have "overlapping" addresses (that is, IP address numbers that are the same but have different subnet masks), only one of the OSPF routes is installed in the kernel routing table.

- GSM commands do not display all routes they should

Various GSM commands that show lists of routes do not completely show all routes that should be shown. This includes GSM commands such as **show ip** and **show bgp detail [ptx | prx]**.

- localpref parameter cannot be set on import proto bgp aspath

When specifying the **localpref** parameter after a route filter in the **import proto bgp aspath** policy statement, a parser error is reported. This syntax is documented as valid.

- Lines greater than 80 characters in gated.conf cause errors

If the `/etc/gated.conf` configuration file contains lines greater than 80 characters, parsing errors occur.

- GateD crashes at startup if Router Discovery server is enabled

If the **routerdiscovery server on** statement is configured in `/etc/gated.conf`, GateD sometimes aborts when GRF is rebooted.

- Non-existent interface for Router Discovery causes GateD to abort

When an interface name is configured in the **interface** sub-statement for the **routerdiscovery server** statement, parsing errors do not occur, but GateD aborts.

- GSM command show ip aggregate can cause GateD to abort

In some instances, executing the GSM **show ip aggregate** command can cause GateD to abort. More specifically, if one of the contributing routes for the aggregate is the same as another route learned from another router, this can happen.

- GateD aborts after interface is brought down/up

Specifically, when a new route is added to the kernel route table using **route add**, and then an interface is brought down and then up by using **ifconfig**, GateD aborts.

- IS-IS does not maintain adjacencies over HSSI

IS-IS forms an initial adjacency over HSSI circuits, but after a short period loses the adjacency.

- 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 **gdc reconfig** command, the user observes the termination of the OSPF protocol, including the flush of the link state database (LSDB) and the termination of adjacencies. Other protocols dependent upon the routes learned or transmitted via OSPF can be affected. This termination occurs every time one is running OSPF and reconfigures GateD.