
GRF 1.4.20 Release Notes

August 31, 1999

These release notes describe the 1.4.20 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 |
| – 1.4.10 | General release | August 25, 1998 |
| – 1.4.12 | General release | October 9, 1998 |
| – 1.4.14 | Directed Partial Release | December 4, 1998 |
| – 1.4.16 | General release | March 15, 1999 |
| – 1.4.18 | General release | May 29, 1999 |

New GRF manuals: 1.4 Update 2

The GRF manual set has been updated to include new features added since software release 1.4.6. GateD information is provided in the *GRF GateD Manual*. You can obtain PDF versions of the manuals from the Lucent Online Support site at: <http://aos.ascend.com>. After you register as a customer and log in, select "AOS Online Manuals" and scroll to GRF.

The GRF 1.4 Update II documentation set consists of the following manuals:

- *GRF 400/1600 Getting Started - 1.4 Update 2*
- *GRF Configuration and Management - 1.4 Update 2*
- *GRF Reference Guide - 1.4 Update 2*
- *GRF GateD Manual - 1.4 Update 2*

Contents

Upgrading system software	3
New features in 1.4.20	4
Reset for Frame Relay PVC statistics	4
Switch router adapter card netmask for alias address	4
ATMP configuration file syntax verification	4
ATM OC-3c now supports local ARP server function	4
Version 2 of ATM OC-12c has updated SAR component	5
BGP communities options are enhanced 5	
Advisories	7
Known problems in 1.4.20	8
ATMP notes	8
ATM OC-3 notes	8
ATM OC-12 v1 notes	10
ATM OC-12v2 notes	11
Ethernet notes	11
FDDI notes	11
Frame Relay notes	12
grdiag note	13
HIPPI notes	13
HSSI notes	14
SONET notes	14
SNMP notes	14
SPD notes	14
System notes	15
GateD advisories	15
GateD notes	18
Problems fixed in 1.4.20	20
ATM OC-12 v2 notes	20
Ethernet notes	20
Filtering notes	21
GateD notes	21
SNMP notes	21
System notes	21
Switch Router Adapter card notes	22

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.

Upgrading system software

Obtain A1_4_20 release files from the FTP server at `service2.boulder.ibm.com`. This release consists of the 1.4.20 code, release note, `RN1_4_20.pdf`, and the GRF 1.4 Update 2 manuals (`config.pdf`, `getting.pdf`, `refer.pdf`, `gated.pdf`).

- `1.4.20.ibm.TAR.gz`
- `1.4.20.ibm.root.gz`
- `RN1_4_20.pdf`
- `config.pdf`
- `getting.pdf`
- `refer.pdf`
- `gated.pdf`

To upgrade GRF software, follow this procedure:

- 1 Log on as root and start a UNIX shell: `super> sh`
 - 2 Prepare to load the 1.4 release on the router's internal flash device:

```
# mountf -w
# cd /flash
# mkdir tmp
# cd tmp
# ftp ftp.min.ascend.com
ftp> cd releases
ftp> cd A1_4_20
```
 - 3 Set the file format and download the files:

```
ftp> bin
ftp> get 1.4.20.ibm.TAR.gz
ftp> get 1.4.20.ibm.root.gz
```
 - 4 Change directory to install the remaining files in `/flash` so they are not removed when the `flashcmd` executes:

```
ftp> lcd ..
ftp> get RN1_4_20.pdf
ftp> config.pdf
ftp> getting.pdf
ftp> refer.pdf
ftp> gated.pdf
ftp> quit
# cd /
```
 - 5 Use `grwrite` to save any configuration changes made since last boot.

```
# grwrite
```
 - 6 The next commands unpack zipped files and clean up flash:

```
# grfins --source=/flash/tmp --release=1.4.20.ibm --activate
# flashcmd -w rm -rf /flash/tmp
# umountf
```
- A reboot installs the files on RAM:
- ```
shutdown -r now - or use
reboot -i
```

## New features in 1.4.20

Additional information about the new features is available in the *GRF 1.4 Update 2* manuals.

### Reset for Frame Relay PVC statistics

A new command is available to reset the PVC statistics reported by `grfr -c dps` for all or selected frame relay links in the GRF. The following is the syntax for the new command:

```
grfr -c crs [-s slot] [-l port] [-i dlci]
```

### Switch router adapter card netmask for alias address

This feature allows the user to have an alias address with a netmask different from the primary address's netmask on a switch router adapter card.

### ATMP configuration file syntax verification

The `aitmd -n` command is used to check the syntax of the ATMP configuration file. By default, the configuration file is `/etc/aitmd.conf`. If the `-f` option is used in conjunction with the `-n` option, then the file specified with the `-f` option is checked.

### ATM OC-3c now supports local ARP server function

A logical interface(s) on a GRF ATM OC-3c media card can now be configured to support ATMARP service for its corresponding IP subnet.

The server function provides ATMARP for IP subnets operating over ATM and supports IP over ATM SVCs. A site can configure and support multiple logical IP subnetworks (LIS) without requiring additional equipment for ATMARP server support.

The GRF supports up to 16 ATMARP servers per media card, with one local server per subnet. A GRF ATM OC-3c card configured with 1–16 server interfaces supports up to 512 ATMARP client entries, a pool shared among the servers configured on that card. The ATM nodes on the IP subnet are assigned an ARP server that is configured on a GRF ATM OC-3c logical interface. The GRF client function is RFC 1577-based, the server function is compatible with both RFC 1577 and RFC 2225 clients.

This feature is supported by additional `maint`, `grstat`, and `grarp` commands. Use `maint 107` to display the local server ARP table, `maint 108` to display the client ARP table, and `maint 109` to display local server statistics. The `grstat arpserver` command returns statistics similar to those from `maint 109`. The `grarp -i interface -r` command displays the contents of the ARP server table (as does `maint 107`). Users can add, delete, and flush entries from the server table using a set of `grarp` commands.

To delete a local server entry from the ARP server table:

```
grarp -i interface -d IP_addr server
```

To add a local server entry to the ARP server table:

```
grarp -i interface -s IP_addr NSAP_addr server
```

To flush all entries from the ARP server table:

```
grarp -i interface -z
```

Refer to chapter 5, ATM OC-3c Configuration, in the *GRF Configuration and Management - 1.4 Update 2* manual for more information.

## Version 2 of ATM OC-12c has updated SAR component

The GRF supports two ATM OC-12c media cards, version 1 and version 2. Version 2 has an updated SAR component. Both versions support the same features and provide a single physical ATM interface that supports 70 logical interfaces. The two cards can operate and co-exist in the same GRF chassis.

Refer to chapter 11, ATM OC-12c Configuration, in the *GRF Configuration and Management - 1.4 Update 2* manual for more information.

## BGP communities options are enhanced

The **send-no-communities** / **send-communities** option is implemented for BGP group statements. It is not supported on peer statements (per-peer behavior is achieved by using subgroups). The default behavior is to send communities. The **send-communities** option enables the user to make that behavior explicit in the configuration file. This option must appear at the end of the **group type** line:

```
group type linktype peeras autonomous_system options
[send-no-community | send-community] { bgp_peer_stmts } ;
```

The functionality of the **mod-aspath** statement for routes being exported to BGP is expanded and an enhanced **aspath-opt** statement can now be used to select BGP routes for export. Note that the **aspath-opt** enhancements only apply to BGP use with exports. The enhancements are not added to **aspath-opt** for imports.

One enhancement is the addition of pattern-matching using the **any** keyword before or after the colon. This can be used to match any ordinary community. Note that pattern matching cannot be used to match the well-known community values, even with "any : any". Those must be matched explicitly.

```
aspath-opt {
 [community (autonomous_system : community-id) | community-id]
 [community no-export | no-advertise | no-export-subconfed | none]
 [community any : community]
 [community autonomous_system : any]
 [community any : any]
}
```

As a second enhancement, the **mod-aspath** statement is extended to delete communities in addition to adding them. The rules for selecting communities to delete are the same as the rules for matching in **aspath-opt**, and the **any** keyword may be used with **delete**.

```
mod-aspath {
 [community ((autonomous_system : community-id)
|community-id | no-export | no-advertise | no-export-subconfed)]
 [delete community ((autonomous_system : community-id)
|community-id | no-export | no-advertise | no-export-subconfed)]
 [delete community autonomous_system : any]
 [delete community any : community-id]
 [delete community any : any]
}
```

For more information, refer to chapter 3, GateD Configuration Statements, in the *GRF GateD Manual - 1.4 Update 2*.

## Advisories

### - IS-IS not supported

The interior gateway protocol IS-IS is not supported in GRF 1.4 software releases.

### - CFMS not supported

The Configuration File Management System (CFMS) is not supported in GRF 1.4 software releases.

### - 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 new *GRF Configuration and Management* manual.

### - 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.

### - 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 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.

## Known problems in 1.4.20

### ATMP notes

#### - System reboot may cause intermittent encapsulation

When a router configured as a home agent in an ATMP network is re-booted, the home agent may intermittently fail (approximately half the time) to GRE-encapsulate packets from the home server to the mobile node. The home agent will remain in this condition until the ATMP interface to the home network router is re-booted.

#### - Advisory: major changes in ATMP configuration parameters

Several parameters have been added to support new ATMP features and the `/etc/aitmd.conf` file has a number of new configuration options. If you are upgrading from an earlier release, you can transfer your ATMP configuration to the parameters used in the new `/etc/aitmd.conf.template` file installed during the upgrade, but you are not required to do so. Your current `/etc/aitmd.conf` file is not replaced during an upgrade.

In the new template file, 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 will be unpredictable. The `netmask_size` parameter is no longer used and is replaced by `vpn_netmask_size`.

#### - 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.

### ATM OC-3 notes

#### - VCs on ATMP home router not established on reboot

ATM interfaces configured with `proto=vc` or `proto=vc_atmp` directly connected to ATM interfaces. Do not get ARP entries established when the GRF reboots.

#### - Bridging across ATM does not fragment packets

When using ATM in a routed bridge network, the bridge does not fragment packets. Pings greater than 1500 will stop.

#### - 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 `proto=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.

**- 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.

**- 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.

**- Server not fully 1577 compatible**

Certain ATMARP clients, do not advertise their own IP address to the ARP server until they have some IP traffic of their own they would like to resolve. They open the SVC to the ARP server, but do not send ARP or Inverse ARP registrations. A pure 1577 client expects the server to send an Inverse ARP request to which the client should generate an Inverse ARP reply. The GRF ATMARP server will not poll these clients for their IP address. Such clients must initiate IP traffic to the GRF to be registered.

**- Filterd dies during SVC initialization**

Filterd dies when bringing up an ATM OC3-V2 running SVCs.

*ATM OC-12 v1 notes*

**- ATM OC-12c - Limited SNMP support**

This release provides limited SNMP support for the ATM OC-12c card. The GigaRouter Enterprise MIB is not supported. MIB-II (RFC 1213) is partially supported. MIB-II was ported from GRF ATM OC-3c but was not updated to reflect OC-12 card parameters such as bit rates and numbers of interfaces. For example, SNMP polling will incorrectly return bit rates of 155 mbs for the ATM OC-12c card.

**- 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 peer. 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 larger than 9788 bytes**

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 system reboot**

An ATM OC-12c card may panic while coming up after the GRF is rebooted under no load. This occurs most of the time.

**- Card may issue BOOT ME after system reboot**

While the system is coming up, an ATM OC-12c card may issue a BOOT ME request and start dumping. This occurs about half the time. The console and `grconslog` report no associated errors.

**- Certain pings to attached interfaces fail**

Pings to interfaces directly attached to ATM OC-12c PVCs specified with `proto=ipnlc` 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 70 configured logical 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` failed.)

*ATM OC-12v2 notes*

**- Card panics (malloc error) with too many interfaces configured**

The ATMOC12v2 card will panic with a "bus timeout" error due to memory/malloc issues. This is occurring with a total of 78 interfaces configured in the GRF (65 on OC12 and 13 on OC3).

**- 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 peer. 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.

**Advisory - Media cards tipover running ping tests**

Sending large pings between two GRFs (for example: `ping -s 15,000`) can cause cards to reboot in a sending or receiving GRF. This condition is not caused by forwarded IP traffic.

*Ethernet notes*

**- Ethernet card not auto-negotiating speed correctly**

TR3348

A GRF Ethernet port configured as autonegotiate may not correctly read the negotiation information coming from some other vendor's Ethernet chip set. However, if both ports are set to autonegotiate, the GRF port correctly senses the other's speed and duplex.

**- 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.

*FDDI notes*

**- Binary load may fail**

Occasionally a load of FDDI/Q run-time code gets out of sequence. The following is an example of the output messages:

```
...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.

### *Frame Relay notes*

#### **- Frame Relay switch configuration 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 the HSSI card. The interfaces are cleared up by rebooting the system.

#### **- HSSI FR circuit under heavy load can lose LICS data**

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 that 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.

#### **- 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 `/etc/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`.

*grdiag note***- Cards may not return to running state after grdiag operates**

After running **grdiag**, a media card may not return to the running state but may stay in the configuration state and continue to reboot. Please do not use **grdiag** until further notice.

To get the card out of configuration state, exit the UNIX shell and read the Card profile of the media card that has not booted correctly. Change to the `load` section and delete the `hw-table` field of the affected card.

Here is an example:

```
super> grcard
0 ATM_OC12_V1 running
1 HSSI_V1 running
2 ATM_OC3_V2 configuring
4 FDDI_V2 running
super> read card 2
CARD/2 read
super> cd load
config = 0
hw-table = <{atm-oc3-v2 0""0""on 6 1 <{1 atm-oc3-v2 /usr/libexec+
boot-seq-index = 1
boot-seq-state = 3
boot-seq-diagcode = 0
```

The `hw-table` field should be empty and look like this: `hw-table = < >`. To empty this field, you use the **delete** command. The only parameter required is the media card type, in this case it is the ATM OC-3c type, `atm-oc3-v2`:

```
super> delete hw-table atm-oc3-v2
```

*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 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 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 status, 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 may not report a failed card

TR1000109

The `mib2d` process periodically looks at all of the cards to check their state. However, if a card loses power, it no longer communicates. `mib2d` does not detect whether there used to be an active card in a particular slot.

## 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. There should be 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.

**System notes****- CLI profiles reference unsupported media cards**

Some CLI profiles include media cards not supported on the GRF. For example, the Load profile displays E1 and T1 run-time binary path names. Please ignore these references.

**- 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. As shown in this excerpt, it reports “incomplete”:

```
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
```

**- 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.

**- GRF will not reboot and reports 1.4.10 under NCLI**

TR 4009

**GateD advisories****- 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 1.4.20 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. To do this, define the following interface in the `/etc/grifconfig.conf` file:

```
#name address netmask broad_dest arguments
lo0 127.0.0.1 255.0.0.0 #A: Standard loopback interface
```

### - 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.

### - Assign an alias as routerid

Customer Support strongly recommends that you configure an alias address on the loopback interface of your router and use this address as the routerid. 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 0xffffffff0
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

IS-IS

### - No IGP group support

This release of GateD does not support the BGP group type **igp**. Use the BGP group type **routing** 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, Lucent recommends 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, perform the following procedure 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 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*

#### **- static routes with retain lost 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. `route add`, `route delete`, and `route flush` commands should never be used while GateD is active.

#### **- 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 ;
```

#### **- OSPF does not work properly with ATM and HSSI point-to-point links**

TR258712, 36795

Currently, OSPF does not work properly with point-to-point links on HSSI and ATM cards. As a workaround, use NBMA interfaces instead of point-to-point interfaces for ATM and HSSI.

In the Interfaces section of the `/etc/gated.conf` file, do not use the "define ...pointtopoint..." statement for ATM and HSSI interfaces. Instead, use the OSPF "interface ...nonbroadcast" statement. This is described at the end of the OSPF section in the *GateD 1.4 Update 2 Manual*.

#### **- Communities: no-export-subconfed is mishandled**

When using Communities in a BGP Confederation, add the following statements at export policy as it applies to group type confed peers to ensure proper handling of no-export-subconfed:

```
proto bgp aspath .* origin any
aspath-opt { community no-export-subconfed }
{all restrict};
```

## Problems fixed in 1.4.20

### *ATM OC-12 v2 notes*

The following problems have been corrected on the verion 2 card only.

**- ATMOC12 card locks up**

TR 4193

The problem has been corrected so that the ATMOC12 card does not lock up when booting up.

**- ATMOC12 card slow buffer leak**

TR 2990

Implemented software and hardware changes needed to ensure that there is no ATMOC12 card slow buffer leak.

**- Under RX load, only 25% of packets were properly converted to IP or ARP**

**packets**

TR 2702

The problem has been corrected by implementing software changes needed for new TX and RX SAR.

**- ATMOC12 cards are showing packet loss**

TR 2775

The problem has been corrected by implementing software changes for a new SAR.

**- ADSM transfer problems**

TR4961

The problem has been corrected by implementing software changes for a new SAR.

**- Packets not forwarding when MTU set to 9,180**

TR 3371

Ping packets greater than 10,000 bytes now are correctly received and responded to.

### *Ethernet notes*

**- Port set for 10 BaseT may mis-report good link to connecting port**

TR3334

Additional checks were added to detect mismatches when hardware indicates that a link is up.

### *Filtering notes*

#### **- Logical interface number 255 (0xff) is reserved as GRF broadcast address**

TR 1000199

On the GRF, logical interface 255 (0xff) is reserved as the system broadcast address. Configuring this address as a data path interferes with the internal communications bus. Do not assign filters to this interface. It cannot be configured as a regular IP address on HSSI and ATM media cards.

HSSI media cards are assigned interface addresses in the range of 0–7f (upper physical port) and 80–fe (lower physical port).

ATM OC-3c media cards are assigned up to 70 logical interfaces on each physical port, in the range of 0–7f (upper physical port) and 80–fe (lower physical port).

ATM OC-12c media cards are assigned up to 70 logical interfaces in the range of 0–fe.

### *GateD notes*

#### **- BGP: ASPath Filtering with community tags**

TR 3324

The problem has been corrected so that customers can now filter and set communities.

#### **- GateD crashes with various errors**

TR 3513

This problem has been corrected in a previous release.

#### **- gdc reconfig crashes with define in export statement**

TR 3681

This problem has been corrected in release 1.4.18.

#### **- When running BGP, the HPPI crashes**

TR 5284

HPPI card will now work properly with larger route table configurations.

### *SNMP notes*

#### **-SNMPwalk enterprise mib causes mib2d to restart and it resets SNMP uptime**

TR 4705

Corrected this problem in the `clntnext.c` file without modifying `gratm mib` entries.

### *System notes*

#### **- Grrt man page does not match grrt usage line**

TR 3240

The man page for the **grrt** command has been corrected. Several options and some examples were removed from the man page in order to match the functionality of this command. The following options were removed:

-a|-A    -r|-R    -Z    -q

**- grarp command times out**

TR 1000163

This problem has been corrected so the **grarp** command does not time out. Timeout value has been changed.

**- Cannot access GRF when disk is receiving errors**

TR 4559

The problem occurs because the logins attempt to look up the file `/var/log/wtmp` which is on the failing disk. The disk error recovery code has been modified so that this is less likely. This problem can be prevented by keeping `/var/log/wtmp` on the "RAM" disk.

*Switch Router Adapter card notes*

**Configuration displays wrong switch number in large configurations (greater than 16)**

TR 4681

Configuration of the switch number is wrong in large configurations (greater than 16).