Oracle® Communications Border Gateway

Essentials Guide

Release S-CX6.4.0

Formerly Net-Net Border Gateway

December 2014

ORACLE

Contents

About This Guide

Overview

The Acme Packet 4500 is offered as a high performance, high capacity border gateway that optimally delivers interactive communications—voice, video, and multimedia sessions—across wireline, wireless, and cable IP network borders. With its compact, single unit, 1U, design, the system provides exceptional functionality in a tightly integrated system.

Audience

This guide is written for network administrators and architects, and provides information about the Border Gateway (BG) application. Supporting configurations are available in the ACLI Configuration Guide, Release version S-CX6.4.0. Please refer to that document as noted.

Supported Platforms

Release Version S-CX6.4.0 is supported on the Acme Packet 4500 and Acme Packet 3800 series platforms.

Related Documentation

The following table lists the members that comprise the documentation set for this release:

Document Name	Document Description
Acme Packet 4500 System Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 4500 system.
Acme Packet 3800 Hardware Installation Guide	Contains information about the components and installation of the Acme Packet 3800 system.
Release Notes	Contains information about the current documentation set release, including new features and management changes.
ACLI Configuration Guide	Contains information about the administration and software configuration SBC.
ACLI Reference Guide	Contains explanations of how to use the ACLI, as an alphabetical listings and descriptions of all ACLI commands and configuration parameters.
Maintenance and Troubleshooting Guide	Contains information about Net-Net SBC logs, performance announcements, system management, inventory management, upgrades, working with configurations, and managing backups and archives.

Document Name	Document Description
MIB Reference Guide	Contains information about Management Information Base (MIBs), Enterprise MIBs, general trap information, including specific details about standard traps and enterprise traps, Simple Network Management Protocol (SNMP) GET query information (including standard and enterprise SNMP GET query names, object identifier names and numbers, and descriptions), examples of scalar and table objects.
Accounting Guide	Contains information about the SBC's accounting support, including details about RADIUS accounting.
HDR Resource Guide	Contains information about the SBC's Historical Data Recording (HDR) feature. This guide includes HDR configuration and system-wide statistical information.
Administrative Security Essentials	Contains information about the SBC's support for its Administrative Security license.

Document Revision History

This section contains a revision history for this document.

Date	Revision Number	Description
March 10, 2008	Rev. 1.0	• Initial release
August 31, 2007	Rev. 1.1	• Add information about historical data recording (HDR) feature

December 17, 2007	Rev. 2.0	• Adds *New Features* section which contains information about the following incrementally available features and other information:
		  • Net-Net BG Licensing
		  • SNMP Additions
		  • System ACLS
		  • Interim QoS Statistics Filename
		  • Hanging Termination
		  • Media Hairpinning across VLANs and Physical Interfaces
		  • Management via Front Interfaces
		  • V3 Spoofing
		  • Backwards and forwards octet strings for ds/dscp parameters
		  • H.248 subtraction enhancements
		  • Failover ServiceChange Command
		• Adds IP address information about vBG being in the HIP
		• Adds information about the **show bgfd vbgs** command

Date	Revision Number	Description

November 1, 2012	Rev. 3.0	• Adds support for Message Session Relay Protocol (MSRP)
		  • *message* media type
		  • *TCP/MSRP* transport protocol
		  • *path* SDP attribute
		  • H.248 *MODIFY* message type
		• Changes previous behavior to allow the establishment of RTCP traffic flows regardless of the endpoint termination mode.
		• Enhances support for relatch signal, defined in H.248.37, *Gateway control protocol: IP NAPT traversal package*

February 27, 2013	Rev.4.0	• Documents the following functionality
		  • *IPv6 support*
		  • Media mirroring support
		  • Hide-Media-Update support
		  • Block SRTP support

December 29, 2014	Rev 4.01	• Removes erroneously added section on QoS VQ statistics

1 Border Gateway Application

Decomposed SBC — Net-Net BG

The Net-Net 4000 can be configured as a BGF logical device as used in the ETSI/TISPAN IMS architecture. A single Net-Net 4000 is configured as a border gateway (BG). This fills a single logical role, whereas the integrated SBC model spans several logical roles. When the Net-Net 4000 is configured to act in the BGF role, it is responsible for controlling media streams as they enter and exit the network. A session controller controls a BG's media operations using H.248 v.2 ETSI/TISPAN Ia profile with long text over a UDP interface.

The BG performs the following tasks on media traffic (RTP and RTCP):

- VLAN tagging
- DSCP Marking
- Resource allocation and reservation
- Media supervision
- QoS Statistics Collection & Reporting
- DoS protection
- Fault management
- Bandwidth Policing
- Media Latching for HNT

The diagram below shows how endpoints first send their call signaling information to the session controller. Based on signaling commands from an endpoint, the SC directs the BG to perform the appropriate action on the media stream thus acting as a gateway to the network.

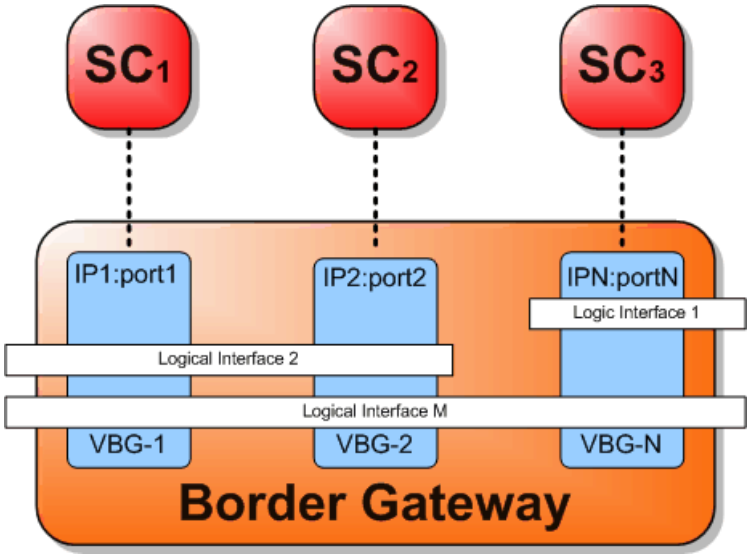

BGF Media Stream Control

The communication between an SC and BG can be decoupled from the signaling events that make up a basic SIP call. The BG gets its directives directly from an SC; there is no call signaling routing logic that occurs on the BG. Different SCs may use different sequences of H.248 control messages to achieve the same results.

The events and their order described in the following call scenario do not necessarily describe the prototype or typical model. They depict only one version of the sequence of events that could happen. One aspect of the BG's functionality that remains constant is that it is essentially controlled by the SC, and does exactly what the SC instructs it to do. The BG does not make any application-based decisions on its own.

The Net-Net BG supports multiple Virtual Border Gateways (VBG), which are logical instances of a BG. A vBG has one H.248 control association identified by a unique IP address and port combination, known as an MID. Control interfaces associate a vBG to an SC and use the Net-Net 4000's front-panel, media interfaces.

When the BG is initialized, each vBG contacts its configured, primary SC, in the order listed in the configuration. The BG attempts to create an H.248 control association between itself and the SC it contacts. Once this control association has been established, the SC can send media control messages to its corresponding vBG.

vBG Network interface Association

Each vBG steers media through realms (and network interfaces implicitly) as listed in the associated-realms parameter. A vBG can be configured to control an exclusive list of realms or all configured realms. We recommend that you configure each vBG with the most restrictive list of associated realms as possible. This is useful if you do not want a particular vBG to have access to certain interfaces, due to topology. When a Net-Net BG's interface fails, or when a next-hop device becomes unreachable, the vBG notifies its controlling SC of the event. You may want to consider this behavior so that SCs are excluded from receiving non-relevant failure messages (i.e., service change commands).

Logical interfaces assigned to specific vBGs

One vBG may be assigned to one or more logical interface, and one logical interface may be assigned to one or more vBG. The SC is thus responsible for adding a termination on the vBG that controls the required logical interfaces for that call. When an SC tries to add a termination on a vBG on an interface that is not assigned correctly, the BG returns a 430 error code (unknown termination ID) to the SC.

In the following example,

- Logical interface 1 is assigned to VBG-N, any failures on this interface will be reported to SC3 only
- Logical interface 2 is assigned to VBG 1 and VBG2, any failures of this interface will be reported to SC1 and SC2 only
- Logical interface M is assigned to all vBGs, any failures of this interface will be reported to all SCs.

- Each vBG can create a call on the logical interfaces that it exclusively touches.

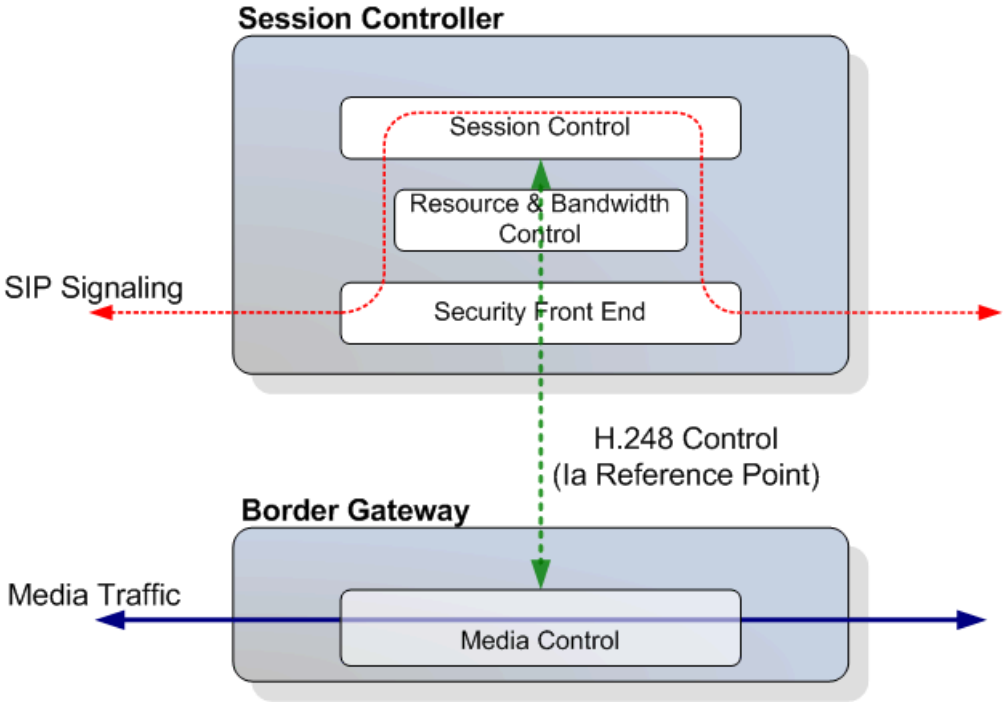

Termination ID Structure

The Termination ID structure is supported as follows:

```
<prefix>/<group>/<interface>/<id>
```

Where:

- `<prefix>`—Is a prefix string always set to IP.
- `<group>`—3 digit number that identifies the group of the interface and id. The group can range from 0-255, per the Ia Profile.
- `<interface>`—Alphanumeric identifier the refers to a configured realm on a Net-Net BG.
- `<id>`—Specific, non-zero, 32 bit termination identifier which is a unique ID for each termination on the Net-Net BG.

A few other H.248 specific terms are:

- Termination—The source or destination of a stream.
- Context—An association between a collection of terminations. A context represents a call on the SC.
- MID— message Identifier (MID) of a message is set to a provisioned name (e.g., domain address/domain name/device name) of the entity transmitting the message (i.e., the SC or vBG).
- Mode—The direction of a flow between two terminations within the Net-Net BG. Valid modes are send only, receive only, send receive, and inactive.

Please refer to ITU-T H.248 for additional details.

Basic Call

The following example walks through the communication between the SC and BG for setting up and tearing down a basic SIP call. The Net-Net BG supports multiple media streams per context/termination. The diagram below shows the relationship between BG, SC, and the types of traffic in a typical scenario.

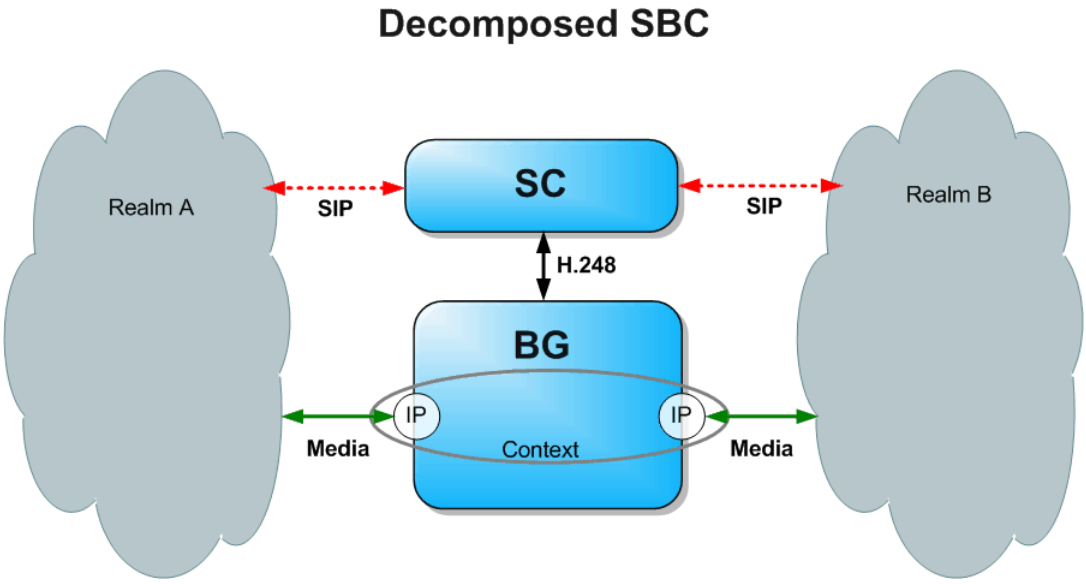

The following diagram illustrates the call flow that is described next.

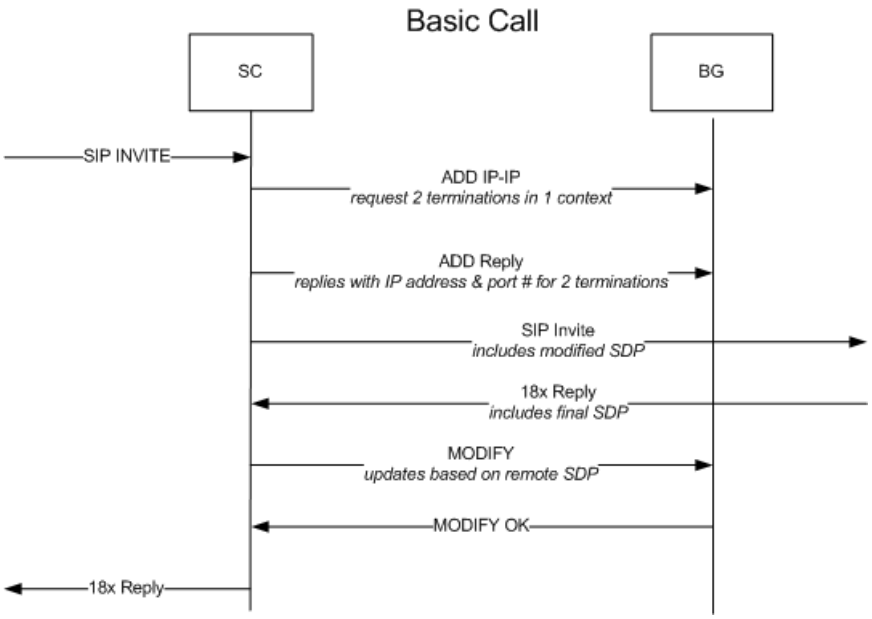

ADD Request from SC to BG

The SC receives a SIP Invite with an SDP offer from the caller and determines the egress realm (from the vBG's point of view) for this call. The SC then sends an H.248 message to a vBG telling it to create a new context with two terminations. One termination should be on the ingress network (realm) of the call, and one termination should be on the egress network (realm) of the call. The SC specifies the realms where the two contexts are created and the type of flow in each direction that

should be created. The SC also includes the caller's remote IP address and port, used as the source address of the flow on the ingress side of the call.

In this H.248 action (which include 2 H.248 ADD commands), the SC requests a Context ID and Termination ID for each of the two terminations. The SC also requests the IP addresses and ports on both sides of the BG where RTP/RTCP traffic will flow into and out of.

ADD Reply from BG to SC

The BG replies to the SC with two termination IDs and one context ID. One of the two termination IDs describes the ingress termination, and the other termination ID describes the egress termination. The BG also replies with the local IP addresses and ports on both the egress and ingress networks (realms) where media will flow in to and out of.

After the SC receives the successful ADD REPLYs from the BG, it can forward the SIP INVITE to its destination in the egress realm. The SIP INVITE can now contain the BG's egress realm IP address and port to successfully receive the media stream.

At this point, the BG has been provisioned with an IP address and port on the caller and network sides. Media can freely flow between two realms.

MODIFY Request From SC to BG

The Modify message sent from the SC to the BG is often used to modify the remote descriptor of the egress termination, as specified by the called party, in response to the SIP Invite. The Modify message can be used for several purposes, including:

- Provide the SDP media details for the flow through the BG, as noted in the endpoint's SDP answer
- Set the Mode of the flows (to open/close gates)
- Set up media policing
- Set ToS
- Update any source/destination IP based filtering
- Send a codec change

MODIFY Reply from BG to SC

The BG sends a reply to the modify message acknowledging that it received and implemented the modify message.

SUBTRACT Request from SC to BG

When signaling indicates that the call has completed and should be torn down, the SC instructs the BG to remove all terminations in the context. It can also request a statistics returned from the BG. QoS monitoring must be enabled on the Net-Net BG if you want to collect statistics.

SUBTRACT Reply from the BG to the SC

The BG replies with end-of-call statistics to the SC after removing this call's terminations and context. The subtract reply functions as the subtract acknowledgement.

ACLI Instructions and Examples

Configuring the Net-Net BG requires a baseline configuration for creating media flows, and BG-specific configuration. The BG license and all other supporting licenses must be installed on this system.

A baseline configuration requires that you have the following configuration elements defined and logically configured:

- system-config
- phy-interface
- network-interface
- realm-config
- steering-pool
- media-manager

In addition, you must create and configure the bgf-config, and one or more vbg-config elements.

The bgf-config essentially turns on the Net-Net BG and sets the debugging level of this task.

Enabling the Net-Net BG and Setting the Log Level

To enable the Net-Net BG using the ACLI:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ```

3. Type **bgf-config** and press <Enter>. The prompt changes to indicate you can configure individual parameters.

   ```
   ACMEPACKET(media-manager)# bgf-config
   ACMEPACKET(bgf-config)#
   ```

   From this point, you can configure the BG global parameters. To view all Net-Net BG configuration parameters, enter a **?** at the system prompt.

4. **state**—Set this parameter to enabled to globally enable the Net-Net BG.
5. **log-level**—*Optional*. Set this parameter to override the system log level for this task.
6. Save your work using the ACLI **done** command.

Enabling High Availability (HA)

This section shows you how to enable HA on your Net-Net BG; you enable this feature in the bgf-config.

To enable HA for the Net-Net BG:

1. **red-bgf-port**—Enter the port number to use for HA synchronization for the two systems operating as Net-Net BGs. The Net-Net BGs will listen for HA messages they exchange on this port. Valid values between `1025` and `65535`, and the default is `1994`.
2. **red-max-trans**—Enter the maximum number of HA synchronized transactions to maintain on the active system in the HA node. The valid range is `0` to `999999999`, and the default is `10000`.

3. **red-sync-start-time**—Enter the amount of time in milliseconds that the active Net-Net SBC checks to confirm that it is still the active system in the HA node. If the active system is still adequately healthy, this timer will simply reset itself. If for any reason the active has become the standby, it will start to checkpoint with the newly active system when this timer expires. The valid range is `0` to `999999999`, and the default is `5000`.
4. **red-sync-comp-time**—Enter amount of time in milliseconds that determines how frequently after synchronization the standby Net-Net SBC checkpoints with the active Net-Net SBC. The first interval occurs after initial synchronizations of the systems; this is the timeout for subsequent synchronization requests. The valid range is `0` to `999999999`, and the default is `1000`.

Configuring the vBG

To configure a vBG using the ACLI:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter>.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Type **vbg-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# vbg-config
   ACMEPACKET(vbg-config)#
   ```

   From this point, you can configure a vBG. To view all Net-Net BG configuration parameters, enter a **?** at the system prompt.

4. **ip-address**—Set the IP address of this vBG. The IP address and port number combination must be unique to specify a vBG. In addition, this IP address must match that of an existing network interface configuration, or at least be added as a HIP entry.
5. **port**—Set the port of this vBG. The well-known H.248 port number is 2944. The IP address and port number combination must be unique to specify a vBG.
6. **mid**—Set the MID that this vBG uses when sending H.248 messages to an SC. The convention for this string is "[vGB's IP address]:vBG's port number". For example: `[192.168.0.15]:2944.` The following syntax is valid for an MID:

   ```
   [1.1.1.1]:2944
   [1.1.1.1]
   <a.b.c>:2944
   <a.b.c>
   ```

   We recommend that you set the mid to the corresponding values of the real IP address and port, i.e., [*1.1.1.1*]:*2944*
7. **state**—Set this parameter to enabled to activate this vBG.
8. **realm-id**—Set the realm/interface where the control association between this vBG and the SC exists.
9. **transport-protocol**—Set this parameter to **UDP**.
10. **encoding**—Enter **text** to set the H.248 message encoding type to uncompressed (i.e. long) format.

11. **associated-realms**—Enter a list of realms through which this vBG can set flows up in. You can also set this parameter to **none** or **all** to indicate no associated realms or that all configured realms are considered as associated realms.
12. Save your work using the ACLI **done** command.

To configure SCs that this vBG may be controlled by:

You must manually provision all SCs that this vBG can communicate with. The order in which you configure them is relevant. The first-configured SC should be the primary SC that this vBG is meant to communicate with. Each additional SC represents the next-in-place SC that this vBG attempts to reach during an SC failure situation.

Continuing from the previous procedure, you will configure a vBG's SC list.

To configure a vBG's SC list using the ACLI:

1. Confirm that you are at the vbg-config configuration sub element level.

   ```
   ACMEPACKET(vbg-config)#
   ```

2. Type **session-controller** and press <Enter>. The prompt changes to indicate you can configure individual parameters.

   ```
   ACMEPACKET(vbg-config)# session-controller
   ACMEPACKET(vbg-sc-config)#
   ```

   From this point, you can configure a vBG's session controller. To view all session controller configuration parameters, enter a **?** at the system prompt.
3. **ip-address**—Set the IP address of this SC.
4. **port**—Set the port of this SC.
5. Save your work using the ACLI **done** command. Repeat this process to configure additional SCs.

H.248 Package Support

In addition to pure media control, the Net-Net BG also supports several other features.

Bandwidth Policing

The BG can provide constant bit rate (CBR) bandwidth policing for call admission control into a network. CBR support is based on the "Sustainable Data Rate" parameter from the H.248 Traffic Management package.

The SC informs the BG of bandwidth allowed for a call, and the BG will impose bandwidth limiting for that call's media traffic as it exits the Net-Net BG.

Differentiated Services

The BG can add DiffServ markings (RFC 2474) onto IP packets exiting into an egress network. The SC informs the BG of a value to insert into IP packets' TOS octets as the packets exits the BG. Diffserv markings are added on a call-by-call basis. This support is based upon the H.248 Differentiated Services package.

If the SC does not specify a value for the ds/dscp property, then the BG will use the TOS bits from the ingress interface or the egress interface. This behavior deviates from the H.248 Differentiated Services package standard which implies a default value of 0x00.

DiffServ markings are configured in the media policy configuration element. For information about configuring DiffServ markings, refer to the "Realm-Based Packet

Marking" section in the *Realms and Nested Realms* chapter in the *Net-Net 4000 ACLI Configuration Guide, Release version 4.1*.

Media Flow Timers

The BG supports three media flow timers which are used to guard against unused or timed-out flows through the BG. The three flow timers are:

1. Initial flow—This flow timer controls how long the BG can wait between the time the flow is created and the first packet is received. The default value is 300 seconds.
2. Subsequent flow—This flow timer controls the maximum length of time that can elapse between two subsequent IP packets before the NAT flow times out and is removed. The default value is 300 seconds.
3. Maximum duration—This flow timer controls the maximum time a NAT flow can exist. Once this time limit has been exceeded, the flow is removed from the BG. The default value is 86400 seconds.

When one of the timers expires, the BG sends a service change command to the SC for the appropriate ContextID and TerminationID. The service change reason accompanying this message is a "910 Media Capability Failure" with method=forced. The SC should reply to the BG to remove the flow. If none of the media flow timers are configured, then this service change command is never sent.

The flow timers are configured globally in the media-manager configuration element. These timers appear as `flow-time-limit`, `initial-guard-timer`, `subsq-guard-timer` parameters.

To configure media flow timers on the Net-Net BG:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ```

3. Type **media-manager** and press <Enter>. The prompt changes to indicate you can configure individual parameters.

   ```
   ACMEPACKET(media-manager-config)#
   ```

4. **flow-time-limit**—Set this parameter maximum total session time for any flow in seconds.
5. **initial-guard-timer**—Set this parameter for how long the Net-Net BG can wait between the time the flow is created and the first packet is received in seconds.
6. **subsq-guard-timer**—Set this parameter to the maximum length of time that can elapse between two subsequent IP packets before the flow times out and is removed (in seconds).
7. Save your work using the ACLI **done** command.

Congestion State

A congestion situation occurs when the Net-Net BG's CPU usage exceeds a given amount. When this happens, the BG notifies the corresponding SC of a congestion condition whenever it receives an H.248 context-creating ADD command. By setting the BG's congestion action parameter to none, the SC is responsible for managing the congestion situation and must adaptively throttle the rate at which it sets up sessions on the BG.

The congestion threshold is defined on the Net-Net BG as a percentage of CPU load. A second critical congestion level is also defined as a CPU load even higher than the congestion state where the BG only admits emergency calls. The critical congestion level is ostensibly configured higher than the congestion level.

Net-Net BG Smoothing

The Net-Net BG can manage a congestion state internally with the congestion action parameter set to smoothing. Smoothing means that the Net-Net BG uses an algorithm to determine how to react in such a case.

When the Net-Net BG enters a congestion state, and the SC tries to create a call, the BG will reject the call with a 510 error code and notify the SC of an mg-overload event (if the ocp package is enabled). In addition, the BG will begin to drop calls using the following algorithm.

$$\text{Drop Rate} = \left(\frac{\text{Current CPU load} - \text{Congestion threshold}}{(100 - \text{Congestion threshold})}\right)$$

In a critical congestion state, the Net-Net BG only allows emergency calls and drops all other calls. An emergency call is identified by the inclusion of an emergency token as sent by the SC.

To configure the Net-Net BG's congestion state management using the ACLI:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ```

3. Type **bgf-config** and press <Enter>. The prompt changes to indicate you can configure individual parameters.

   ```
   ACMEPACKET(media-manager)# bgf-config
   ACMEPACKET(bgf-config)#
   ```

   From this point, you can configure the BG global parameters. To view all Net-Net BG configuration parameters, enter a **?** at the system prompt.

4. **con-threshold**—Set this parameter to the CPU load percentage that triggers a congested state. Valid values range from 1 - 100%.
5. **con-interval**—Set this parameter to the interval in seconds at which the Net-Net BG is monitored at in congestion state. Valid values range from 1 - 60 seconds.
6. **crit-con-threshold**—Set this parameter to the CPU load percentage that triggers a critical congested state. Valid values range from 1 - 100%.
7. **crit-con-interval**—Set this parameter to the interval in seconds at which the Net-Net BG is monitored at in critical congestion state. Valid values range from 1 - 60 seconds.
8. **con-action**—Set this parameter to **none** for the Net-Net BG to take no action and expect the SC to take action, when the BG enters a congested state. Set this parameter to **smoothing** to take action by dropping incoming calls when the Net-Net BG enters a congested state.
9. Save your work using the ACLI **done** command.

Service Faults

There are four common fault situations that can occur on the Net-Net BG:

- Loss of H.248 Control Association between the Net-Net BG and SC
- SC goes out of service
- Net-Net BG goes out of service (or operational state is set to disabled)
- Net-Net BG failsover from the active to the standby
- Media interface/network failure

On all of these cases, the Net-Net BG sends a Service Change message to the BG to attempt recovery procedures.

Net-Net BG Failover

When an BG is configured for high availability, a switchover between redundant Net-Net BGs does not affect established sessions. Media state is maintained across the two redundant systems. All contexts are migrated over to the newly in-service Net-Net BG such that the switchover is transparent to the network. The Net-Net BG that assumes the active role and sends Restart / 902 Warm Boot service change command to its associated SC.

SC Failure / Loss of Control Association

In normal operation, the BG exchanges control messages with an SC. If no operations are pending, the SC sends empty audit value commands to the Net-Net BG that act as heartbeat messages. Initially, the SC provisions the Net-Net BG with an inactivity timer. If the Net-Net BG doesn't receive a message from the SC within the inactivity timer period, the BG sends an event notification to the SC and looks for another SC.

When the vBG loses a control association with its SC, the vBG tries to recontact with the SC by sending a Disconnected/900 Service Change command. If the SC fails to respond, then the Net-Net BG attempts to contact the next-configured SC with a Failover / 909 Service Change command. The vBG continues through the list of SCs until an SC replies. After the vBG exhausts the list of configured SCs, it returns to the original SC and sends a Disconnected/900 Service Change command.

Media interface failure and recovery on a BG

When the Net-Net BG detects a failure on an interface, the Net-Net BG sends a ServiceChange Forced message on a wildcarded termination ID to all SCs that can control the failed interface. If the Net-Net BG then receives an ADD command for a termination on the failed interface, the BG returns an Error Code 503 to the SC. The the failed interface goes back in service, the Net-Net BG sends a Restart / Service Restored message for that interface.

Media Hairpinning

There are several cases in which the Net-Net BG takes action when media needs to be hairpinned through it.

Case 1

In the following case, the SC sets up a single context to interconnect two endpoints. The media will flow from the endpoints' access realm(s) through the BG.

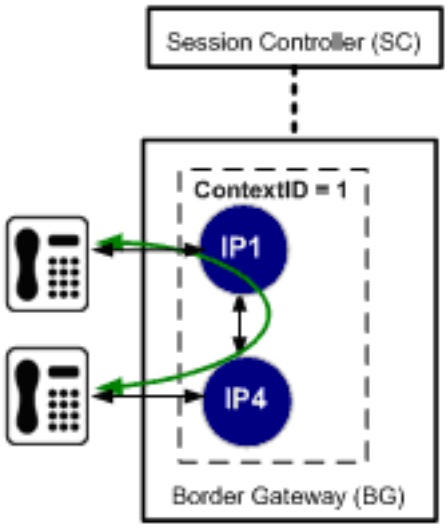

Case 2

When the SC is not able to set up a single context, as in case 1, it sets up 2 contexts, one per endpoint. If the local IP address of IP2 is the same as the remote IP address of P4 (IP2 and IP4 are on the same VLAN and physical interface), then the media will remain hairpinned within the BG.

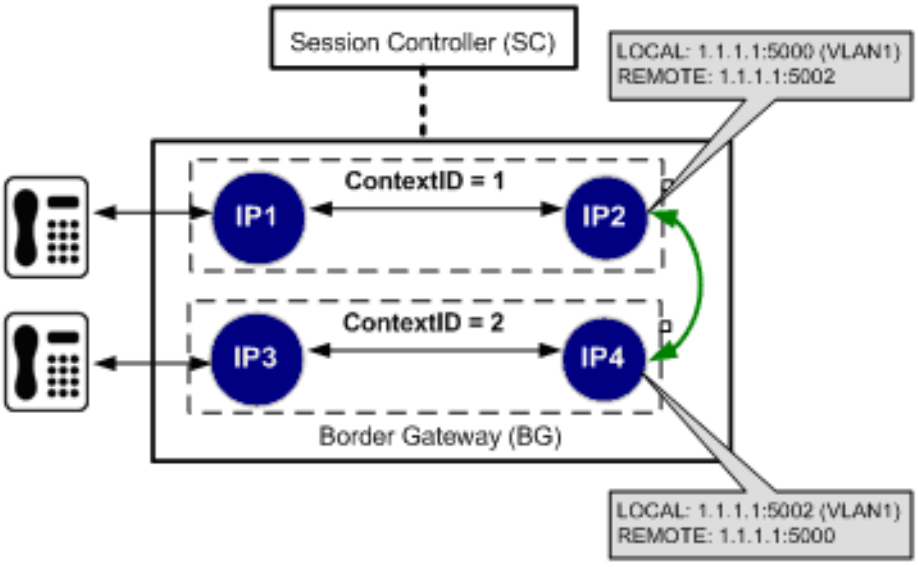

Case 3

When the SC is not able to set-up a single context, as shown in case-1, then the SC will set-up 2 Contexts, one per endpoint. If the Local IP address of IP2 is not the remote IP address of IP4, or IP2 and IP4 are on different VLANs, even though both addresses are owned by the BG), then the media will flow from the Net-Net BG onto the external network and then back into the Net-Net BG.

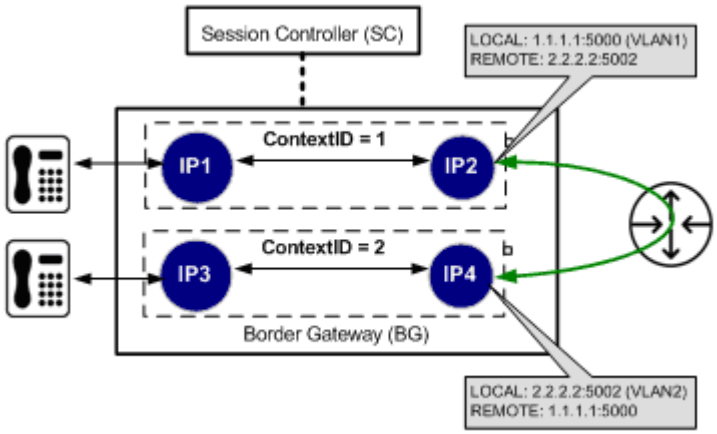

Hosted NAT Traversal

When the SC discovers that an endpoint is behind a NAT, it sends the ipnat/latch signal to the Net-Net BG. This indicates that the SDP provided by the SC in the remote descriptor cannot be reliably used as the egress destination for media sent by the termination. Therefore, the Net-Net BG will have to latch onto the source IP address and port of incoming RTP packets.

If the SC sends ipnat/latch for a particular termination, any SDP sent in the H.248 remote descriptor is ignored. The Net-Net BG performs a 2-step latching processes in this event. It first learns the source address and port from which an endpoint's RTP is received, and then uses this address and port as the egress destination for the media sent by the termination.

If NAT traversal is disabled using the ipnat latch {latch=off} signal, then the SDP present in the remote descriptor determines the destination of where the termination sends media. If no remote SDP is specified, the termination will stop sending media on the streams for which no remote SDP is specified.

Hosted NAT traversal is not enabled on the Net-Net BG, but on the SC.

RTCP NAT Traversal

When endpoints are behind a NAT, and their public port assignment is out of control of the calling application, there is no guarantee that the RTP and RTCP flows' ports will be contiguous. If NAT traversal is not enabled, the Net-Net BG uses one internal NAT entry for a unidirectional media flow. When an endpoint is behind a NAT, distinct NAT entries are required for both RTP and RTCP. You must enable RTCP for hosted NAT traversal parameter to force the Net-Net BG to allocate additional resources for this application. If NAT traversal is enabled, then the gm/rsb property is not supported.

To configured RTCP NAT traversal:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

```
ACMEPACKET(configure)# media-manager
```

3. Type **media-manager** and press <Enter>. The prompt changes to indicate you can configure individual parameters.

```
ACMEPACKET(media-manager-config)#
```

4. **hnt-rtcp**—Set this parameter to **enabled** for the Net-Net BG to use additional resources required for RTCP HNT traversal.
5. Save your work using the ACLI **done** command.

Re-latching

Ipnat/latch is treated as a transitory signal. If the Net-Net BG receives a latch signal from the SC while not in a latching state, the bidirectional media flow is terminated until the Net-Net BG receives media from an endpoint and the latching process begins again. It is important to note that if re-latching (which again can be accomplished by sending ipnat/latch {napt=latch} or {napt=re-latch} in a new H.248 MODIFY request) is enabled, the BG stops the media flows until it learns a new address from latching.

Source Address Filtering

The Net-Net BG supports source address filtering, which in conjunction with ipnat/latch is used to deploy restricted media latching. The Net-Net BG supports gm/saf and gm/spf properties, which it receives from an SC. When a source address filter is enabled and latching is enabled, the Net-Net BG applies source address filtering as specified by the SC.

Restricted Latching

The first packet received by the Net-Net BG, which meets the filtering criteria is latched onto and only media from that source will be accepted by the termination. The Net-Net BG will then send media back to that address. The following table describes the Net-Net BG's interaction between "gm", the RemoteDescriptor, "ipnapt", and BG configuration.

Remote Descriptor	Ipnapt/latch Signal	"gm" source filtering properties	BG Configuration	Result
Specified IP address Specified RTP port number	Not Included	Not Included	hnt-rtcp = disabled	• Accept RTP/RTCP packets from any address:port • Send RTP packets to address:port specified in RemoteDescriptor • Send RTCP packets to address:(port+1) specified in RemoteDescriptor
Specified IP address Specified RTP port number	Not Included	gm/saf = ON gm/sam = specified gm/spr = specified Note: port range is not allowed	hnt-rtcp = disabled	• Accept RTP/RTCP packets from the address:port specified by gm/sam and spr only • Send RTP packets to address:port specified in RemoteDescriptor • Send RTCP packets to address:(port+1) specified in RemoteDescriptor

Remote Descriptor	lpnapt/latch Signal	"gm" source filtering properties	BG Configuration	Result
Specified IP address Specified RTP port number	napt = LATCH	-	hnt-rtcp = enabled	• Accept RTP/RTCP packets only from the "latched" source address:port • Send RTP packets to source address:port of the first RTP packet received for the flow • Send RTCP packets to source address:port of the first RTCP packet received for the flow
Specified IP address Specified RTP port number	napt = LATCH	gm/saf = ON gm/sam = specified gm/spf and gm/spr = not allowed	hnt-rtcp = enabled	• Accept RTP/RTCP packets only from the "latched" source address:port • Send RTP packets to source address of the first RTP packet received for the flow that matches the criteria specified in gm/sam • Send RTCP packets to source address of the first RTCP packet received for the flow that matches criteria specified in gm/sam

Debugging

The following section explains the commands used to obtain BGF statistics which can be helpful for debugging purposes. The logfile for the BGF application is `log.bgfd`, which can be found in the expected location on the Net-Net BG's file system.

The H.248 message log files are available as the `h248msg.log` file. This file is created when the BG log level is set to trace or higher.

show bgfd statistics

This show command lists BG application statistics.

```
ACMEPACKET# show bgfd

13:54:03-148

State                     -- Period -- -------- Lifetime --------
                   Active    High   Total      Total  PerMax    High
BGF Contexts            0       0       0          0       0       0
BGF Streams             0       0       0          0       0       0
Media Pending           0       0       0          0       0       0
BGF Transactions            ---- Lifetime ----
                          Recent      Total  PerMax
Transactions rcvd            0          0       0
Dup Transactions rcvd        0          0       0
Replies Resent               0          0       0
```

```
Transactions sent           11         45       8
Transactions Resent         68        264      46
Dup Replies rcvd             0          0       0
Transactions timedout       11         43       8
```

show bgfd vbgs

This show command lists the statistics for a configured vBG. Including the MID displays only that vBGs statistics, whereas entering no argument displays the statistics for all configured vBGs.

```
ACMEPACKET# show bgfd vbgs [192.168.200.124]:2944

13:53:58-643
VBG [192.168.200.124]:2944 [RESTART IN PROGRESS]
State                     -- Period -- -------- Lifetime --------
                 Active    High   Total      Total  PerMax    High
VBG Contexts        0       0       0          0       0       0
VBG Streams         0       0       0          0       0       0
VBG Transactions            ---- Lifetime ----
                          Recent      Total  PerMax
Transactions rcvd            0          0       0
Dup Transactions rcvd        0          0       0
Replies Resent               0          0       0
Transactions sent            1         22       2
Transactions Resent         10        131       7
Transactions timedout        1         21       2
```

Historical Data Recording

Historical data recording (HDR) refers to a group of management features that allow you to configure the Net-Net BG to collect statistics about system operation and function, and then send those records to designated servers. System statistics, defined in detail below, are saved to a comma-separated value (CSV) files, which are then sent to the designated server(s).

Information types are grouped so that you can refer to a set of statistics by simply invoking their group name. Within each group, there are several metrics available.

How It Works

In the system configuration, you can enable HDR by first turning on the system's collection function, then choosing the records you want to capture, and finally setting up server(s) to which you want records sent.

The main collect configuration (found in the main system configuration) allows you to create global settings that:

- Turn the HDR function on and off
- Set the sample rate in seconds, or the time between sample individual collections
- Set the time in seconds in between individual pushes to designated servers (configured in the push receiver configuration accessed via the collect configuration)
- Set the time you want the collect to start and stop; time is entered in year, month, day, hours, minutes, and seconds

You also configure setting for each group of data you want to collect, and the push receiver (server) to which you want data sent.

About the CSV File

When you enable HDR and configure one or more servers to which you want records sent, data is transmitted in a CSV file in standard format. There is one CSV file per record group type, and the first record for each file is a header containing the field name for each attribute in that file.

Collection Interval and Push

In your HDR configuration, you set parameters that govern:

- The groups for which the Net-Net BG collects records
- How frequently the Net-Net BG collects records
- How frequently the Net-Net BG sends records off-box

Factoring in the number of groups for which you collect records, you can calculate the number of records that will be sent per push. The number of files that are sent off-box equals the number of groups for which the Net-Net BG is collecting records; there is always one additional record for each group, a header file containing the field name for each attribute.

The number of records in a file, then, equals the push interval divided by the sample interval time multiplied by the number of groups, plus one. Take the case, for example, where you set a push interval time of 60 seconds and a sample interval time of 5 seconds, with a group of ten records. With these settings, the Net-Net BG would send 120 group records and 10 header records (for a total of 130 records) for each push.

Note that after each push, the Net-Net BG clears (deletes) all records. The Net-Net BG also clears files on system reboot, and after three consecutive push failures.

Group Record Types

In the group-name parameter for the group-settings configuration, you can enter any one of the groups record type defined in the following table. You specify the collection object, and then all metrics for that groups are sent.

Collection Object	Metrics Included
General system statistics (system)	• CPU utilization • Memory utilization • Health score • Redundancy state • Current signaling sessions • Current signaling session rate (CPS) • CAM utilization media • CAM utilization ARP • I2C bus state • License capacity
Interface statistics (interface)	• Interface index • Name/description • Type • MTU • Speed • Physical address • Administrative status • Operational state • In last change • In octets • In unicast packets • In non-unicast packets • In discards • Out errors • Out octets • Out unicast packets • Out non-unicast packets • Out discards • Errors
Combined session agent statistics (session-agent)	• Hostname • System name • Status • Inbound active sessions • Inbound session rate (CPS) • Outbound active sessions • Outbound session rate (CPS) • Inbound sessions admitted • Inbound sessions not admitted • Inbound concurrent sessions high • Inbound average session rate (CPS) • Outbound sessions admitted • Outbound sessions not admitted • Outbound concurrent sessions high • Outbound average session rate (CPS) • Max burst rate (in and out) (CPS) • Total seizures • Total answered sessions • Answer/seizure ratio • Average one-way signaling latency (ms) • Maximum one-way signaling latency (ms)

Collection Object	Metrics Included
Session realm statistics (session-realm)	• Realm name • Inbound active sessions • Inbound session rate (CPS) • Outbound active sessions • Outbound session rate (CPS) • Inbound sessions admitted • Inbound sessions not admitted • Inbound concurrent sessions high • Inbound average session rate (CPS) • Outbound sessions admitted • Outbound sessions not admitted • Outbound concurrent sessions high • Outbound average session rate (CPS) • Max burst rate (in and out) (CPS) • Total seizures • Total answered sessions • Answer/seizure ratio • Average one-way signaling latency (ms) • Maximum one-way signaling latency (ms)
Environmental voltage statistics (voltage)	• Voltage type • Description • Current voltage (mv)
Environmental fan statistics (fan)	• Fan type • Description • Speed
Environmental temperature statistics (temperature)	• Type • Description • Value (Celsius)
SIP status statistics (sip-sessions)	• Sessions • Subscriptions • Dialogs • Call ID map • Rejections • ReInvites • Media sessions • Media pending • Client transaction • Server transaction • Response contexts • Saved contexts • Sockets • Requests dropped • DNS transactions • DNS sockets • DNS results • Session rate • Load rate

Collection Object	Metrics Included
SIP error/event statistics (sip-errors)	• SDP offer errors • SDP answer errors • Drop media errors • Transaction errors • Media expiration events • Early media expirations • Early media drops • Expired sessions • Multiple OK drops • Multiple OK terminations • Media failure drops • Non-AXK 2XX drops • Invalid requests
SIP policy/routing (sip-policy)	• Local policy lookups • Local policy hits • Local policy misses • Local policy drops • Agent group hits • Agent groups misses • No routes found • Missing dialog • Inbound SA constraints • Outbound SA constraints • Inbound REG SA constraints • Outbound REG SA constraints • Requests challenged • Challenge found • Challenge not found • Challenge dropped
SIP server transaction (sip-server)	• All states • Initial • Trying • Proceeding • Cancelled • Established • Completed • Confirmed • Terminated
SIP client transactions (sip-client)	• All states • Initial • Trying • Calling • Proceeding • Cancelled • EarlyMedia • Completed • SetMedia • Established • Terminated
SIP ACL status (sip-ACL-status)	• Total entries • Trusted • Blocked

Collection Object	Metrics Included
SIP ACL operations (sip-ACL-oper)	• ACL requests • Bad messages • Promotions • Demotions
SIP session status (sip-status)	• Sessions initial • Sessions early • Sessions established • Sessions terminated • Dialogs early • Dialogs confirmed • Dialogs terminated

ACLI Instructions and Examples

This section shows you how to configure HDR. You need to set up:

- The collection configuration to govern sample and push intervals, start and end times for collection
- Setting to support this feature across an HA node
- The group settings configuration that tells the Ne-Net 4000 what groups of records to collect, when to start and stop collecting them, and how often to sample for that group
- Push receivers that take the records the Net-Net BG sends

All HDR parameters are RTC-supported, so you can save and activate your configuration for them to take effect.

Accessing the HDR Configuration Parameters

You access the parameters that enable and support HDR using the ACLI **system-config** path.

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ```

2. Type **system** and press <Enter>.

   ```
   ACMEPACKET(configure)# system
   ACMEPACKET(system)#
   ```

3. Enter **collect** and press <Enter>. From here, you can type a question mark (?) to see individual parameters for the configuration.

   ```
   ACMEPACKET(system)# collect
   ACMEPACKET(collect)#
   ```

Global Collection Settings: Boot State, Collection Start and Stop, Sample and Push Intervals

You access the collection configuration through the ACLI **system-configuration** menu. Once in the collection configuration, you can establish the global settings for HDR collection.

Note that the push receiver is configured in a sub-configuration; refer to the Push Receiver Settings (33) for details.

To configure global settings for HDR support:

1. **boot-state**—Set this parameter to `enabled` to start group collection, or to disabled (default) to prevent the Net-Net BG from collecting HDR statistics. This parameter does not go into effect until the system is rebooted. You can also use the ACLI request collect start command to start collection; using this command, you can start collection for all groups, or for one specified group.
2. **sample-interval**—Enter the time in minutes for how often you want the Net-Net BG to sample data records. Leaving this parameter set to 0 (default) turns off the feature. The maximum value for this parameter is 120 minutes (2 hours).
3. **push-interval**—Enter the time in minutes for how often you want the Net-Net BG to send collected records to push receiver(s). The default is 0.
4. **start-time**—Enter the exact date and time (for your local timezone) when you want the Net-Net BG to start HDR collection; this time is either `now` (default) or a time in the future. Your entry must be in the format `yyyy-mm-dd-hh:mm:ss`, where: `yyyy` is the year, `mm` is the month, `dd` is the day, `hh` in the hour, `mm` is the minutes, and `ss` is the second.
5. **end-time**—Enter the exact date and time (for your local timezone) when you want the Net-Net BG to finish HDR collection; this time is either `never` (default) or a time in the future. Your entry must be in the format `yyyy-mm-dd-hh:mm:ss`, where: `yyyy` is the year, `mm` is the month, `dd` is the day, `hh` in the hour, `mm` is the minutes, and `ss` is the second. There is no default for this parameter.

HDR for an HA Node

If you are using the HDR feature on an HA node (or redundant pair of Net-Net BGs), then you need to make sure that several parameters in the collection configuration are set appropriately.

It is recommended that you do not change these parameters from their default for a normal HA node configuration. Therefore, if you need to change them to support HDR, you should do so with caution.

To configure parameters for HDR support across an HA node:

1. **red-collect-state**—Set the state of HA support for the collector function; HA support is `disabled` by default, so you must change this parameter to `enabled` if you want to use it.
2. **red-max-trans**—Enter the maximum number of HA synchronized transactions to maintain on the active system in the HA node. The valid range is 0 to 999999999, and the default is `1000`.
3. **red-sync-start-time**—Enter the amount of time in milliseconds that the active Net-Net BG checks to confirm that it is still the active system in the HA node. If the active system is still adequately healthy, this timer will simply reset itself. If for any reason the active has become the standby, it will start to checkpoint with the newly active system when this timer expires. The valid range is 0 to 999999999, and the default is `5000`.

4. **red-synch-comp-time**—Enter amount of time in milliseconds that determines how frequently after synchronization the standby Net-Net BG checkpoints with the active Net-Net BG. The first interval occurs after initial synchronizations of the systems; this is the timeout for subsequent synchronization requests. The valid range is 0 to 999999999, and the default is `1000`.

Collection Group Settings

You can configure multiple collection groups on your Net-Net BG; the names of these groups appear in the Group Record Types (28) section above. Collection group settings are accessible through the collection configuration.

Note that the sample collection interval, start time, and end time you set here override the ones established in the global collection settings. The largest value you can enter for an group's sample collection must be smaller than the global push interval value.

To configure collection group settings:

1. Access the collection group (**group-settings**) configuration by way of the collection configuration. Once

   ```
   ACMEPACKET(system-config)# collect
   ACMEPACKET(collect)# group-settings
   ```

2. **group-name**—Enter the group name corresponding to the records that you want to collect; there are 21 possible groups for which the Net-Net BG can collect data. The system group name is the default for this parameter; the other possible names to which you can refer are listed in the Group Record Types (28) table above.
3. **mode**—To turn HDR for this collection group on, set this parameter to `enabled`. By default, it is set to `disabled`.
4. **sample-interval**—Enter the time in minutes for how often you want the Net-Net BG to sample data records for the specified group. Leaving this parameter set to 0 (default) turns off the feature for this group. The maximum value for this parameter is 120 minutes (2 hours).
5. **start-time**—Enter the exact date and time (for your local timezone) when you want the Net-Net BG to start collecting records for this group; this time is either now or a time in the future. Your entry must be in the format `yyyy-mm-dd-hh:mm:ss`, where: `yyyy` is the year, `mm` is the month, `dd` is the day, `hh` in the hour, `mm` is the minutes, and `ss` is the second. There is no default for this parameter.
6. **end-time**—Enter the exact date and time (for your local timezone) when you want the Net-Net BG to stop collecting records for this group; this time is either never or a time in the future. Your entry must be in the format `yyyy-mm-dd-hh:mm:ss`, where: `yyyy` is the year, `mm` is the month, `dd` is the day, `hh` in the hour, `mm` is the minutes, and `ss` is the second. There is no default for this parameter.

Push Receiver Settings

You can configure multiple servers to receive the records that the Net-Net BG. Push receiver settings are accessible through the collection configuration.

If you configure more than one server, then the Net-Net BG sends data to all of the servers. If one server fails, the Net-Net BG generates an SNMP trap. In terms of clearing data, this means that if there are four servers configure and the Net-Net BG successfully pushes data to three of them, then it will clear the data.

To configure servers to act as push receivers for HDR data:

1. Access the collection group (**group-settings**) configuration by way of the collection configuration. Once

   ```
   ACMEPACKET(system-config)# collect
   ACMEPACKET(collect)# push-receiver
   ```

2. **address**—Enter the IP address or hostname of the push receiver (server) to which you want records sent. The default for this parameter is `0.0.0.0`.
3. **username**—Enter the username that the Net-Net BG will use when it tries to send records to this push server using FTP. There is no default for this parameter.
4. **password**—Enter the password (corresponding to the username) that the Net-Net BG will use when it tries to send records to this push server using FTP. There is no default for this parameter.

   This entry works differently from other ACLI configuration parameters. For other parameters, you enter the parameter name followed by a <Space> and the value; the you press <Enter>.

   4a. Type the parameter name **password**, and then press <Enter>.

   ```
   ACMEPACKET(push-receiver)# password
   ```

   4b. At the prompt, type the password to use when the Net-Net BG. The display does not echo the password you enter.

   ```
   Enter password: [enter the password]
   ```

   4c. The ACLI then asks you to enter the password again. If the passwords match, then you will be return to the user prompt to continue with configuring the push server.

   ```
   Enter password again: [enter the password again]
   ACMEPACKET(push-receiver)#
   ```

   If the passwords do not match, then you receive an error message and must set the password again.

   ```
   Error:  Password mismatch - aborted.
   ACMEPACKET(push-receiver)#
   ```

5. **data-store**—Enter the directory on the push receiver where you want collected data placed. There is no default for this parameter.

Controlling HDR from the Command Line

For added ease-of-use, you can stop and start record collection from the command line in Superuser Mode. You can stop and start record collection for the entire HDR process, or you can specify a group name for which you want to stop and start collection.

When you stop collection using this command, it will not restart it using this command.

To start record collection from the command line:

1. In Superuser mode, type the ACLI **request collect** command, and the word **start**. If you press <Enter> at this point, the Net-Net BG will stop all record collection. If you continue with the command-line entry to specific a group-name and press <Enter>, collection for that record group only will be stopped.

   ```
   ACMEPACKET# request collect start voltage
   ```

To stop record collection from the command line:

1. In Superuser mode, type the ACLI **request collect** command, and the word **stop**. If you press <Enter> at this point, the Net-Net BG will stop all record collection. If you continue with the command-line entry to specific a group-name and press <Enter>, collection for that record group only will be stopped.

   ```
   ACMEPACKET# request collect stop voltage
   ```

To restart record collection from the command line:

1. In Superuser mode, type the ACLI **request collect** command, and the word **restart**. If you press <Enter> at this point, the Net-Net BG will restart all record collection. If you continue with the command-line entry to specific a group-name and press <Enter>, collection for that record group only will be stopped.

   ```
   ACMEPACKET# request collect restart voltage
   ```

To delete all collection data files resident on the Net-Net BG:

1. In Superuser mode, type the ACLI **request collect** command, and the word **purge**. If you press <Enter> at this point, the Net-Net BG will restart all record collection. Then press <Enter>, collection for that record group only will be stopped.

   ```
   ACMEPACKET# request collect purge
   ```

To display all active collection groups and push servers:

1. In Superuser mode, type the ACLI **request collect** command, and the word **status**. If you press <Enter> at this point, the Net-Net BG will restart all record collection. Then press <Enter>, collection for that record group only will be stopped.

   ```
   ACMEPACKET# request collect status
   ```

2 BG Basic Configuration Support

Overview

The previous chapter details how to configure your Net-Net BG, and H.248 connection to an SC. In addition to the configuration steps previously described, you need to create a baseline configuration that establishes required logical functions on the Net-Net BG. This section explains necessary configurations.

All configurations listed here are explained fully in the Net-Net 4000 ACLI Configuration Guide, Release Version S-CX6.4.0. Please refer to that document for all necessary explanations and procedures.

Getting Started

In the Net-Net 4000 ACLI Configuration Guide, Release Version S-CX6.4.0, refer to the Getting Started chapter to learn how to configure the following aspects of the system.

- Bootparameters
- Network Time
- Net-Net BG software licensing

System Configuration

In the Net-Net 4000 ACLI Configuration Guide, Release Version S-CX6.4.0, refer to the System Configuration chapter to learn how to configure the following aspects of the system.

- General System Information
- Physical Interfaces
- Network Interfaces
- System Alarm levels
- Syslog and Process log servers and logging facilities
- Media Manager

Realm Configuration

In the Net-Net 4000 ACLI Configuration Guide, Release Version S-CX6.4.0, refer to the Realms and Nested Realms chapter to learn how to configure the following aspects of the system.

- Realms
- Steering pools

ToS Marking

You can configure the Net-Net BG to set ToS values for traffic exiting the system. You can also configure the SC the tell the Net-Net BG how to mark outbound traffic. Until the SC configures the Net-Net BG for ToS marking, the Net-Net BG will perform ToS marking according to configuration. Once the SC sends ds/dscp properties, the Net-Net BG will use them and not revert to internal configurations.

3 Updates

Release 4.1.3 Additions

The following sections note additional BG functionality added in Release Version 4.1.3.

Net-Net BG Licensing Information

When you use the Net-Net BG, you need to enable the features defined in the table below.

For more information about how to obtain and install licenses, refer to the *Net-Net ACLI 4000 Configuration Guide*'s *Getting Started* chapter.

License	Description
ACP	Enables the Net-Net BG to respond to ACP requests. Required for Net-Net EMS use.
H.248 BG	Enables the Net-Net 4000 to operate as a Net-Net Border Gateway.
HA	Enables two Net-Net BGs to work as an HA node so that, in case of failover, one system can take over for the other. The two systems paired as an HA node checkpoint configuration, signaling state, and media.
LI	Enables lawful intercept use (for media only) of the X3 LI package on the Net-Net BG.
QoS	Enables measurement for QoS (jitter, packet latency, and packet loss) on the Net-Net BG.

SNMP Additions

This section describes new traps and tables added for SNMP support for these areas:

- Media supervision trap—When configured to do so, the system generates the `apSysMgmtMediaSupervisionTimerExpTrap` if the media supervisor (or flow guard) timer expires.
- H.248 control association traps—The system generates the `apSysMgmntH248AssociationLostTrap` if the H.248 control association on either the session controller or the border gateway in a pair goes out of service for any reason besides a configuration change. Once the control association is re-established (back in service), the system sends a trap communication the loss has been cleared; this trap is the `apSysMgmntH248AssociationLostClearTrap`.
- R-factor below threshold traps—You can set a threshold for the r-factor of calls, much like you set thresholds for CPU utilization or memory usage. You configure the r-factor threshold in the system configuration, using the alarm-threshold parameter by using the r-factor type and setting the threshold value and alarm severity.

When a call's r-factor falls below the threshold you configure, the system sends the `apSysMgmtRFactorBelowThresholdTrap`. To clear the condition, the call's r-factor must rise above the least severe threshold for a time equal to two periods as set in the **vq-qos-interval** parameter, at which time the system sends the `apSysMgmtRFactorBelowThresholdClearTrap` trap.

Definitions and OIDs: Objects

This section provides samples of the SNMP objects and information about their OIDs.

These are the definitions:

```
apSysMgmtBorderGatewayIdOBJECT-TYPE
SYNTAX          DisplayString (SIZE (1..255))
MAX-ACCESSaccessible-for-notify
STATUS   current
DESCRIPTION
"The identifier of a Border Gateway."
::= { apSysMgmtMonitorObjects 32 }

apSysMgmtRFactorOBJECT-TYPE
SYNTAXDisplayString (SIZE (1..255))
MAX-ACCESSaccessible-for-notify
STATUScurrent
DESCRIPTION
"The R-factor value of the call that has exceeded a
configured threshold."
::= { apSysMgmtMonitorObjects 33 }

apSysMgmtCallIdOBJECT-TYPE
SYNTAX          DisplayString (SIZE (1..255))
MAX-ACCESSaccessible-for-notify
STATUS   current
DESCRIPTION
"A call correlation identifier."
::= { apSysMgmtMonitorObjects 34 }
```

These are the OIDs:

Object Name	OID
apSysMgmtBorderGatewayID	1.3.6.1.4.1.9148.3.2.5.32
apSysMgmtRFactor	1.3.6.1.4.1.9148.3.2.5.33
apSysMgmtCallId	1.3.6.1.4.1.9148.3.2.5.34

Definitions and OIDs: Traps

This section provides samples of the SNMP traps and information about their OIDs.

These are the definitions:

```
apSysMgmtMediaSupervisionTimerExpTrapNOTIFICATION-TYPE
OBJECTS{ apSysMgmtCallId }
STATUScurrent
DESCRIPTION
" The trap will be generated when a media supervision timer
has expired.  This behavior is disabled by default but may
be enabled by changing the 'media-supervision-traps'
```

```
parameter of the 'media-manager' configuration element.  The
included object is the call identifer for the call which had
the timer expire."
::= { apSystemManagementMonitors 34 }

apSysMgmntH248AssociationLostTrapNOTIFICATION-TYPE
OBJECTS{ apSysMgmtBorderGatewayId }
STATUScurrent
DESCRIPTION
" This trap will be generated when an H248 control
association between a border gateway and session controller
is lost.  The included object is the border gateway
identifier."
::= { apSystemManagementMonitors 35 }

apSysMgmntH248AssociationLostClearTrapNOTIFICATION-TYPE
OBJECTS{ apSysMgmtBorderGatewayId }
STATUScurrent
DESCRIPTION
" This trap will be generated when an H248 control
association between a border gateway and session controller
has been restored.  The included object is the border
gateway identifier."
::= { apSystemManagementMonitors 36 }

apSysMgmtRFactorBelowThresholdTrapNOTIFICATION-TYPE
OBJECTS{ apSysMgmtRFactor,
    apSysMgmtCallId }
STATUScurrent
DESCRIPTION
" This trap will be generated when a call using the
vq-qos statistics gathering function determines that
the R-factor has gone below a configured threshold.
The included objects are the current RFactor value
and a call identifier."
::= { apSystemManagementMonitors 37 }

apSysMgmtRFactorBelowThresholdClearTrapNOTIFICATION-TYPE
OBJECTS{ apSysMgmtCallId }
STATUScurrent
DESCRIPTION
" This trap will be generated when a call using the
vq-qos statistics gathering function determines that
the R-factor has recovered after having gone below a
configurable threshold.  The included object is a call
identifer"
::= { apSystemManagementMonitors 38 }
```

These are the OIDs:

Trap Name	OID
apSysMgmtMediaSupervisionTimerExpTrap	1.3.6.1.4.1.9148.3.2.6.0.34
apSysMgmtH248AssociationLostTrap	1.3.6.1.4.1.9148.3.2.6.0.35
apSysMgmtH248AssociationLostClearTrap	1.3.6.1.4.1.9148.3.2.6.0.36
apSysMgmtRFactorBelowThresholdTrap	1.3.6.1.4.1.9148.3.2.6.0.37
apSysMgmtRFactorBelowThresholdClearTrap	1.3.6.1.4.1.9148.3.2.6.0.38

ACLI Instructions and Examples

This section describes the configuration steps you must take to enable media supervision and r-factor SNMP support.

To enable media supervision SNMP support:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Type **media-manager** again and press <Enter>.

   ```
   ACMEPACKET(media-manager)# media-manager
   ACMEPACKET(media-manager-config)#
   ```

4. **media-supervisor-traps**—Set this parameter to enabled if you want to use the media supervisor traps to report if the H.248 control association on either the session controller or the border gateway in a pair goes out of service for any reason besides a configuration change. Once the control association is re-established (back in service), the system sends a trap communication the loss has been cleared. By default, this parameter (and use of these traps) is disabled.

To set the r-factor threshold and use SNMP support:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **system** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# system
   ACMEPACKET(system)#
   ```

3. Type **system-config** and press <Enter>.

   ```
   ACMEPACKET(system)# system-config
   ACMEPACKET(system-config)#
   ```

4. Type **alarm-threshold** and press <Enter>.

   ```
   ACMEPACKET(system-config)# alarm-threshold
   ACMEPACKET(alarm-threshold)#
   ```

5. **type**—Set this parameter to `rfactor` to specify the type of event threshold you want to set. This parameter is empty by default.
6. **severity**—Set the severity of the alarm generated when a call crosses the r-factor threshold. Valid values are: `minor` (default), `major`, and `critical`.
7. **value**—Set the percentage that represents the threshold value; when this value is crossed, the alarm will be triggered. There is no default for this parameter, and valid values are from `1` to `99`.

vBG Statistics

This section shows an example of and explains the statistics you can display using the ACLI **show bgfd vbgs** command.

You can use the ACLI **show bgfd vbgs** command to see information about the Net-Net BG's vBG transactions according the H.248 message types. The display is broken down by message type, and it also contains summary data about all message types appearing at the top of the display. When you use the ACLI **show bgfd vbgs all** command, the systems returns comprehensive information about all MIDs. But you can also limit the display to show only information about a specific MID by using the name of that MID as the last argument in the command. For any MID for which the system has no data, the `---< NO DATA AVAILABLE >---` message appears along with the MID name.

When you look at the display, you will note that sections are divided into "server" and "client" categories. The following table defines for which message types the Net-Net BG is a server and for which it is a client.

Server	Client
ADD	AUDITVAL
MODIFY	AUDITCAP
SUBTRACT	NOTIFY- • CONGESTION • HANGING-TERM • INACTIVITY • OTHER SVCCHANGE- • FAILOVER • FORCED • GRGACEFUL • DISCONNECTED • HANDOFF • NONROOT-FORCED • NONROOT-RESTART • NONROOT-OTHER

```
ACMEPACKET# show bgfd vbgs all

11:26:58-117

VBG [192.168.0.1]:2944 [ASSOCIATION IN SERVICE]

State                    -- Period -- -------- Lifetime --------

                Active    High   Total      Total  PerMax    High
VBG Contexts         0       1       3          3       2       1
VBG Streams          0       1       3          3       2       1
VBG Transactions           ---- Lifetime ----
                         Recent      Total  PerMax

Transactions rcvd             11         11       6
Dup Transactions rcvd          0          0       0
Replies Sent                  11         11       6
Replies Resent                 0          0       0
Transactions Rx Timedout       0          0       0
Transactions sent              0          1       1
Transactions Resent            2          3       2
Replies Rcvd                   1          1       1
Transactions Tx timedout       0          0       0


ADD (11:26:58-117)

                 --------- Server --------  --------- Client --------

Command            Recent     Total  PerMax   Recent     Total  PerMax
                  ------  ---------  ------   ------  ---------  ------
Requests                6          6       6       0          0       0
Retransmissions         0          0       0       0          0       0
Successful              6          6       6       0          0       0
Response Retrans        0          0       0       0          0       0
Transaction Timeouts    -          -       -       0          0       0


MODIFY (11:26:58-117)
                 --------- Server --------  --------- Client --------

Command            Recent     Total  PerMax   Recent     Total  PerMax
                  ------  ---------  ------   ------  ---------  ------
Requests                5          5       5       0          0
Retransmissions         0          0       0       0          0       0
Successful              5          5       5       0          0       0
Response Retrans        0          0       0       0          0       0
Transaction Timeouts    -          -       -       0          0       0
```

```
SUBTRACT (11:26:58-117)

                 --------- Server --------   --------- Client --------
Command            Recent     Total  PerMax   Recent     Total  PerMax
                   ------  ---------  ------   ------  ---------  ------
Requests                6          6       6        0          0       0
Retransmissions         0          0       0        0          0       0
Successful              6          6       6        0          0       0
Response Retrans        0          0       0        0          0       0
Transaction Timeouts    -          -       -        0          0       0

SVCCHANGE-RESTART (11:26:58-117)

                 --------- Server --------   --------- Client --------
Command            Recent     Total  PerMax   Recent     Total  PerMax
                   ------  ---------  ------   ------  ---------  ------
Requests                0          0       0        1          1       1
Retransmissions         0          0       0        3          3       3
Successful              0          0       0        1          1       1
Response Retrans        0          0       0        0          0       0
Transaction Timeouts    -          -       -        0          0       0


---< NO DATA AVAILABLE >----<MOVE>
---< NO DATA AVAILABLE >----<AUDITVAL>
---< NO DATA AVAILABLE >----<AUDITCAP>
---< NO DATA AVAILABLE >----<NOTIFY-CONGESTION>
---< NO DATA AVAILABLE >----<NOTIFY-HANGING-TERM>
---< NO DATA AVAILABLE >----<NOTIFY-INACTIVITY>
---< NO DATA AVAILABLE >----<NOTIFY-OTHER>
---< NO DATA AVAILABLE >----<SVCCHANGE-FAILOVER>
---< NO DATA AVAILABLE >----<SVCCHANGE-FORCED>
---< NO DATA AVAILABLE >----<SVCCHANGE-GRACEFUL>
---< NO DATA AVAILABLE >----<SVCCHANGE-DISCONNECTED>
---< NO DATA AVAILABLE >----<SVCCHANGE-HANDOFF>
---< NO DATA AVAILABLE >----<SVCCHANGE-NONROOT-FORCED>
---< NO DATA AVAILABLE >----<SVCCHANGE-NONROOT-RESTART>
---< NO DATA AVAILABLE >----<SVCCHANGE-NONROOT-OTHER>
---< NO DATA AVAILABLE >----<Other>
```

Hanging Termination

The Net-Net BG supports the Hanging Termination Detection package, specified in ITU-T H.248.36, and can notify its session controller of this capability. This means that the Net-Net BG can identify potential mismatches between the information in the record of Context and Termination identities between itself and the session controller. Upon mismatch detection, the Net-Net BG notifies the session controller, which subsequently subtracts the relevant TerminationID.

How It Works

The Net-Net BG checks each incoming ADD and MODIFY command for an events descriptor with the hanging termination event. When it finds this event, the Net-Net BG then looks for a provisioned `timerx` value—the time between the last message exchanged and the generation of the hanging termination event—which it then uses to start the hanging termination time unless it subsequently finds a `timerx` parameter in the incoming message.

If the timer expires prior to the Net-Net BG receiving any message for the termination, the Net-Net BG will send a NOTIFY message to the MGC. It re-initializes the timer once it sends the message. The Net-Net SBC restarts the timer with the `timerx` value if it receives a message for the termination before the timer expires.

ACLI Instructions and Examples

You enable hanging termination support by configuring the `timerx` parameter in the virtual border gateway configuration.

To enable hanging termination support:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter>.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Enter **vbg-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# vbg-config
   ACMEPACKET(vbg-config)#
   ```

4. **timerx**—Enter the `timerx` value in seconds for the Net-Net BG when it receives a hanging termination event without a `timerx` value of its own. This is also enabled on any termination for which the `timerx` value is received.

   The default for this parameter is 0, which disables this functionality.

5. Save and activate your configuration.

Management via Front Interfaces

You can configure your Net-Net BG and Net-Net SC to use the front interfaces for management traffic, i.e. traffic that customarily travels over the system's rear interfaces. The front interfaces support these types of traffic:

- ICMP (with host-in-path, or HIP, interface configured)
- Telnet (with HIP interface configured)
- FTP (with HIP interface configured)
- SNMP (traps and read-write functionality, with HIP interface configured)
- Process log
- Syslog
- SSH/SFTP

Note that support for several of these all types requires you to configure the HIP interface. The HIP parameters are effectively firewall functions that open the well-known ports for specified services on front interfaces.

ACLI Instructions and Examples

This section shows you how to configure your front interfaces for management traffic, including configuration for the HIP.

If you are familiar with using the HIP and associated functionality on the Net-Net SBC, then these setting will be familiar to you.

To configure management service functionality on a front interface, you must define the IP addresses on the front physical interfaces of your Net-Net SBC where you will receive management traffic. Adding HIP entries automatically opens the well-known port associated with a service.

To allow management traffic over the Net-Net BG's media interfaces:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **system** and press <Enter>.

   ```
   ACMEPACKET(configure)# system
   ACMEPACKET(system)#
   ```

3. Type **system-config** and press <Enter>.

   ```
   ACMEPACKET(system)# network-interface
   ACMEPACKET(network-interface)#
   ```

4. **add-hip-ip**—Set all possible IP address(es) on which you want the Net-Net SBC to accept administrative traffic. Entries in this element are IP addresses of front panel network interfaces. This parameter can accept multiple IP addresses. You can later remove this entry by typing **remove-hip-ip** followed by the appropriate IP address.

5. **add-ftp-ip**—Set the IP address where ports 20 and 21 are opened. This lets standard FTP packets enter the Net-Net SBC and reach the host. You can later remove this entry by typing **remove-ftp-ip** followed by the appropriate IP address.

6. **add-icmp-ip**—Set the IP addresses to pass standard ping packets to the host; this parameter can accommodate multiple ping IP addresses. You can later remove this entry by typing **remove-icmp-ip** followed by the appropriate IP address.

   When you configure multiple ICMP ping addresses in for a network interface, you must also configure the host-in-path addresses in the hip-ip-list for each IMCP address. For security, if the ICMP address and the hip-ip-list are not added for an address, the Net-Net 4000 hardware discards ICMP requests or responses for the address.

   To remove multiple IP addresses at one time, type the **remove-icmp-ip** and a <Space>, open quotation mark ("), the IP addresses you want removed from the list each separated by a space, close quotation mark ("), and then press <Enter>.

   ```
   ACMEPACKET (network-interface)# remove-icmp-ip "142.214.5.34
   124.8.67.3"
   ```

7. **add-snmp-ip**—Set the IP address where port 161 is opened. This lets SNMP traffic enter the Net-Net SBC and reach the host. You can later remove this entry by typing **remove-snmp-ip** followed by the appropriate IP address.

8. **add-telnet-ip**—Set the IP address where port 23 is opened for Telnet access. You can later remove this entry by typing **remove-telnet-ip** followed by the appropriate IP address.

9. **add-ssh-ip**—Set the IP address for SSH access, permitting remote SFTP packets to come into the system over the front interface. SSH and SFTP share this address because the SSH server offers SFTP as a service; thus they can share a configured address. In addition, SSH and SFTP use TCP destination port 22.

   You can later remove this entry by typing **remove-ssh-ip** followed by the appropriate IP address.

Interim QoS Statistics Filename

On the Net-Net decomposed model, you can configure periodic logging for quality of service (QoS) statistics that are stored to a CSV file stored on the system. The file storing this data uses a name indicating the system's target name, creation date, and timestamp rounded to the nearest time unit you set as that file rotation time. This mechanism for timestamping files consistent with the frequency at which the file are rotated.

In addition, you have the option of either pushing the CSV files containing QoS statistics to a designated location, or you can pull them from the Net-Net BG.

Timestamp Style

The CSV's file name contains the following information:

- Target name of the Net-Net BG—The name you set in the system's boot configuration **target name** parameter; it is the name that appears in the ACLI system prompt (both in User and Superuser modes).
- Creation date—The date the file was created in the format `YYYY-MM-DD`, where `Y` is the year, `M` the month, and `D` the day.
- Timestamp—Time of day the file was created in the format `hh-mm`, where `h` is the hour and `m` the minute. In other Net-Net SBC releases, the time shown is the exact time of file creation. However, now the Net-Net BG uses a time representing the nearest multiple of the value you set in the file rotation time interval (set in the accounting configuration), rounded down. That is, if you set a file rotation time of five minutes, and a file is created at 33 minutes after the hour, then the 30 would be the time appearing in the timestamp.
- Numerical suffix—When more than one CSV file is created during the file rotate interval, the Net-Net BG adds a number reflecting how many times the file has been updated. The oldest file has no numerical suffix, and then the numbers increment so the highest number appended to the file name reflects the fact that it is the most newly-created file. This number appears at the very end of the file name.

Your file name might look like this:

```
QoS_BGF-2007-12-05_16-55.csv.3
```

In this sample:

- The target name of the Net-Net BG is `QoS_BGF.`
- The creation date is `2007-12-05.`
- The timestamp is `16-55`.
- The numerical suffix is `3`.

More than one file is created, and the numerical suffix used, when the size of a file exceeds the maximum file size you set in the accounting configuration. That is, rotation is based on file size. Note that if there are multiple files, then the first one will contain less than five minutes of information, while all the subsequent files will contain the amount of information collected over the time you set in the file rotate time parameter. For example, if the time you set in that parameter is five minutes and multiple files are created in that period, then the first and oldest file will have no suffix, the next file in the sequence will bearing a `.1` numerical suffix, and so on.

FTP Push and Pull

As with the Net-Net SBC, you can use the FTP push mechanism to send collected CSV files to a remote server you designate. For more information about the Net-Net SBC' FTP push feature, refer to the *Net-Net 4000 Accounting Guide*; there you will find a detailed explanation of the mechanism and how to configure it. In addition, however, you must also set the **vq-qos-stats** parameter in the media manager configuration to enable FTP push of the CSV files on your Net-Net BG.

If you want, you can also pull those files from the Net-Net BG to place them at a chosen location. Rather than configuring how often you want the Net-Net BG to push files to a remote server, you configure how often you want to rotate stored files on the Net-Net BG (before you presumably pull them). The files are rotated according the file rotate time interval you set in the accounting configuration. If the number of stored files exceeds the maximum number of files you have configured during an interval, then the Net-Net BG removes a file with the oldest timestamp until the number of file is within the configured maximum number.

Pushing Files from the Command Line

You can push files from the command line rather than waiting for the system to do so at the configured interval using the ACLI **notify** command. When you use this command for interim QoS statistics, you must specify the type of file (`CSV-QOS`). Used with the file type `CSV-CDR`, the system pushes the RADIUS CDRs; used without any file type defined, the system attempts to push both types of files.

To push CSV QoS files using the ACLI notify **command:**

1. In Superuser mode at the command line, type **notify pusher pushdata CSV-QOS**. Then press <Enter>.

   ```
   ACMEPACKET# notify pusher pushdata CSV-QOS
   ```

System ACLs

You can configure a system access control list (ACL) for your Net-Net BG that determines what traffic the Net-Net BG allows over its management interface (wancom0). By specifying who has access to the Net-Net BG via the management interface, you can provide DoS protection (not policing) for this interface.

Using a list of IP addresses and subnets that are allowable as packet sources, you can configure what traffic the Net-Net BG accepts and what it denies. All IP packets arriving on the management interface are subject to these rules; if it does not match your configuration for system ACL, then the Net-Net BG drops it.

Note, however, that all IP addresses configured in the SNMP community table are automatically permitted.

ACLI Instructions and Examples

The new subconfiguration **system-access-list** is now part of the system configuration. For each entry, you must define an IP source address and mask; you can specify either the individual source host or a unique source subnet.

If you do not configure this list, then there will be no ACL/DoS protection for the Net-Net BG's management interface.

Adding an ACL for the Management Interface

You access the **system-access-list** via system path, where you set an IP address and netmask. You can configure multiple system ACLs using this configuration.

To add an ACL for the management interface:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **system** and press <Enter> to access the signaling-level configuration elements.

   ```
   ACMEPACKET(configure)# system
   ACMEPACKET(system)#
   ```

3. Type **system-access-list** and press <Enter>.

   ```
   ACMEPACKET(system)# system-access-list
   ACMEPACKET(system-access-list)#
   ```

4. **source-address**—Enter the IP address representing for the source network for which you want to allow traffic over the management interface.

5. **netmask**—Enter the netmask portion of the source network for the traffic you want to allow. The netmask is in dotted decimal notation.

Management

This section provide information about system ACL removal, and about viewing system ACL statistics and configurations.

Notes on Deleting System ACLs

If you delete a system ACL from your configuration, the Net-Net BG checks whether or not there are any active FTP or Telnet client was granted access when the entry was being removed. If such a client were active during ACL removal, the Net-Net BG would warn you about the condition and ask you to confirm the deletion. If you confirm the deletion, then the Net-Net BG's session with the active client is suspended.

The following example shows you how the warning message and confirmation appear. For this example, and ACLI has been deleted, and the user is activating the configuration that reflects the change.

```
ACMEPACKET# activate-config
Object deleted will cause service disruption:
 system-access-list: identifier=172.30.0.24

 ** WARNING: Removal of this system-ACL entry will result
             in the lockout of a current FTP client

Changes could affect service, continue (y/n) y

Activate-Config received, processing.
```

Viewing System ACL Configurations

The system-access-list configuration has been added to the list of configurations you can display using the show configuration and show running-config ACLI commands. It will display each system ACL entry.

```
ACMEPACKET# show running-config system-access-list
system-access-list
        dest-address                    165.31.24.2
        netmask                         225.225.0.0
        last-modified-date              2007-04-30 13:00:02
system-access-list
        dest-address                    175.12.4.2
        netmask                         225.225.225.0
        last-modified-date              2007-04-30 13:00:21
task done
```

Viewing System ACL Statistics

You can display statistics for system ACLs using the **show ip stats** ACLI command. Two new entries have been added to let you see the total number of ACL denials and the last ACL denial the Net-Net BG made.

```
ACMEPACKET# show ip stats
               total            3170
              badsum               0
            tooshort               0
            toosmall               0
             badhlen               0
              badlen               0
          infragments              0
          fragdropped              0
          fragtimeout              0
              forward              0
          fastforward              0
          cantforward             14
         redirectsent              0
     unknownprotocol               0
           delivered            1923
            localout             855
           nobuffers               0
         reassembled               0
          fragmented               0
        outfragments               0
            cantfrag               0
          badoptions               0
             noroute               0
             badvers               0
              rawout               0
             toolong               0
           notmember               0
               nogif               0
             badaddr               0
          acl-denials            1233
   last-srcip-denied    174.35.60.72

ACMEPACKET#
```

V3 Spoofing

The Net-Net BG has the ability to perform V3 spoofing—that is, to advertise it is using the H.248 version 3 protocol in ServiceChange messages—by enabling a parameter in the virtual border gateway configuration. In fact, the Net-Net BG uses H.248 version 2. Despite the fact the Net-Net BG has this ability to spoof protocol versions, it is recommended that you do not use it; instead, your most reliable deployment will occur using version 2, the officially supported protocol.

When the Net-Net BG registers with a session controller, the initial H.248 version is 1, but the Net-Net BG advertises the highest version it can support. The SC responds without specifying a version number (in the services descriptor) if it can support the version the Net-Net BG requests. If the SC cannot support the requested version number, it specifies the highest version it can support in the services descriptor.

ACLI Instructions and Examples: New Configurations

This section shows you how to enable your Net-Net BG to communicate support for H.248 version 3 where no vBG functionality has previously been configured.

To enable V3 spoofing on your Net-Net BG for new configurations only:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Type **vbg-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# vbg-config
   ACMEPACKET(vbg-config)#
   ```

4. **options**—Set the options parameter by typing **options**, a <Space>, the option name **faux-version=x** (where **x** is the H.248 version number you want to advertise; for V3 spoofing, enter `3`) with a "plus" sign in front of it. Then press <Enter>.

   ```
   ACMEPACKET(vbg-config)# options +faux-version=3
   ```

   If you type **options** and then the option value for either of these entries without the "plus" sign, you will overwrite any previously configured options. In order to append the new option to the configuration's options list, you must prepend the new option with a "plus" sign as shown in the previous example.

ACLI Instructions and Examples: Changing V3 Spoofing Support in Existing Configurations

This section describes the process for adding V3 spoofing support to an existing configuration. To change the messaging the Net-Net BG currently uses, you must take the control association out of service and then re-establish it. The next time the Net-Net BG establishes a control association with an SC, it will advertise the version you configure.

The following steps summarize the process:

1. Add/remove/modify the faux-version option in the vBG configuration
2. Set the state of the vBG to disabled (which will break the control association once the configuration is saved and activated)
3. Save and activate your configuration
4. Set the vBG state to enabled (which will re-establish the control association once the configuration is saved and activated)

5. Save and activate your configuration

The process below provides detailed instructions for how to perform each step using the ACLI.

To add V3 spoofing to an existing vBG configuration:

```
In Superuser mode, type configure terminal and press <Enter>.
ACMEPACKET# configure terminal
ACMEPACKET(configure)#
```

6. Type **media-manager** and press <Enter> to access the media-manager path.

```
ACMEPACKET(configure)# media-manager
ACMEPACKET(media-manager)#
```

7. Type **vbg-config** and press <Enter>.

```
ACMEPACKET(media-manager)# vbg-config
ACMEPACKET(vbg-config)#
```

8. **options**—Set the options parameter by typing **options**, a <Space>, the option name **faux-version=x** (where **x** is the H.248 version number you want to advertise; for V3 spoofing, enter `3`) with a "plus" sign in front of it. Then press <Enter>.

```
ACMEPACKET(vbg-config)# options +faux-version=3
```

   If you type **options** and then the option value for either of these entries without the "plus" sign, you will overwrite any previously configured options. In order to append the new option to the configuration's options list, you must prepend the new option with a "plus" sign as shown in the previous example.

9. **state**—Set the state of the vBG to `disabled`.

10. Save and activate your configuration.

```
ACMEPACKET# save-config
Save-Config received, processing.
waiting 1200 for request to finish
Request to 'SAVE-CONFIG' has Finished,
Save complete
Currently active and saved configurations do not match!
To sync & activate, run 'activate-config' or 'reboot activate'.
ACMEPACKET# activate-config
Activate-Config received, processing.
waiting 120000 for request to finish
Request to 'ACTIVATE-CONFIG' has Finished,
Activate Complete
ACMEPACKET#
```

11. **state**—Set the state of the vBG to `enabled`.

12. Save and activate your configuration.

```
ACMEPACKET# save-config
Save-Config received, processing.
waiting 1200 for request to finish
Request to 'SAVE-CONFIG' has Finished,
Save complete
Currently active and saved configurations do not match!
To sync & activate, run 'activate-config' or 'reboot activate'.
ACMEPACKET# activate-config
Activate-Config received, processing.
waiting 120000 for request to finish
Request to 'ACTIVATE-CONFIG' has Finished,
```

```
Activate Complete
ACMEPACKET#
```

Media Hairpinning across VLANs and Physical Interfaces

You can configure your Net-Net BG to hairpin media across VLANs and across physical interfaces when the termination identifiers are in the same realm. However, the termination identifier does not have to use the same VLAN or physical interface. When you enable this feature, it means that the Net-Net BG can hairpin media (redirect media within the itself) between two endpoints whose local and remote SDP are inverse reflections of one another.

Consider the example in the following diagram. IP2 and IP4 have local and remote IP addresses that are inverse reflections of one another. That is, IP2's local address is IP4's remote address, and IP2's remote address is IP4's local address.

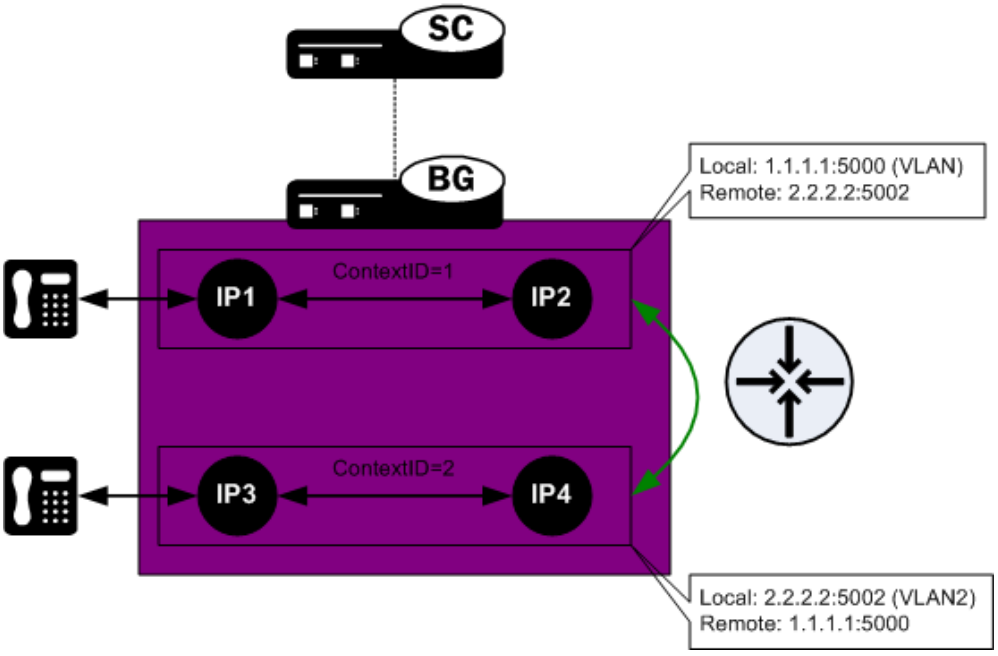

How It Works

Your Net-Net BG maintains a table of flow entries, each entry being keyed by IP port and interface. The interface contains further information about the physical slot and port, in addition to VLAN information. However, this feature uses the concept of "hairpin identifiers," which circumvent physical and logical interfaces. To create hairpin media flows across interfaces, the Net-Net BG uses hairpin identifiers (virtual information) instead of physical interface data, which makes the physical interface information independent of the flows as long as hairpin identifiers are configured correctly. This allows the Net-Net BG to link media flows across physical interfaces.

Overlapping IP Addresses Restrictions

Despite the fact that the hairpin identifier replaces physical and logical interfaces as part of the key to search for flows, IP addresses and ports remain part of the key. Thus, steering pools IP addresses cannot overlap for realms with the same hairpin identifier.

This condition means that you cannot create multiple steering pools with the same IP address and port ranges if the realms associated with those steering pools are configured with the same hairpin identifier. This is the case regardless of differences between the physical and logical interface definitions for the logical interfaces defined for logical interfaces configured with that IP address.

Consider sample scenarios illustrating where you cannot use overlapping IP addresses and ports when you use the hairpin identifier.

- You configuration might contain steering pools assigned to the same realm. Further, these steering pools might use their network interface parameter to designate different physical or logical interfaces with the same IP address.
- Your configuration might contain steering pools assigned to different realms that share the same hairpin identifier. For example, you might have Realm A and Realm B. Realm A is configured with a hairpin identifier value of 2, and a network interface NIA:0 that has 192.168.5.3 as an IP address. Similarly, Realm B is configured with a hairpin identifier value of 2, and a network interface NIB:5 that also has 192.168.5.3 as an IP address.

  In addition, Realm A is associated with a steering pool configuration set to 192.168.5.3, using ports 1000 through 2000. Realm B is associated with a steering pool configuration also set to 192.168.5.3, but using ports 3000 through 4000. Although the physical and logical interface information is replaced with the hairpin identifier, the difference in port ranges differentiates sufficiently between the between the ports.

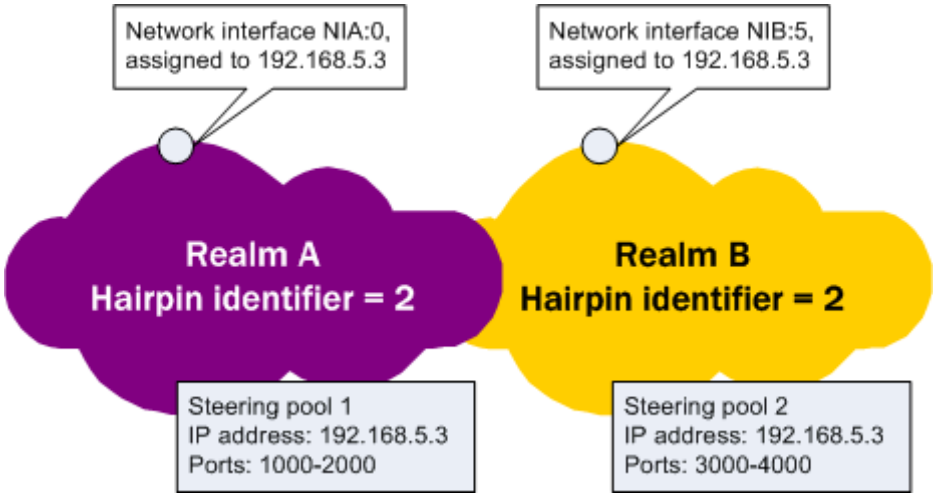

ACLI Instructions and Examples

This section shows you how to configure your Net-Net BG for media hairpinning across VLANs and physical interfaces.

To configure the hairpin identifier for a realm:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter>.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Type **realm-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# realm-config
   ACMEPACKET(realm-config)#
   ```

4. Select the realm where you want to want to add SIP per user CAC.

   ```
   ACMEPACKET(realm-config)# select
   ```

5. **hairpin-id**—Enter the identifier for hairpin media flows for this realm. The default is 0, which disables the hairpinning media across VLANs and physical interfaces. Valid values are in the range of 1 to 65534.

6. Save and activate your configuration.

H.248 Subtract Enhancements

This section describes the following enhancements to the Net-Net BG's H.248 support:

- SUBTRACT request with ContextID=Specific & Termination ID=ALL—The Net-Net BG can support a SUBTRACT request with a Specific ContextID and a wildcarded Termination ID (ALL). This wildcarded Termination ID can include either an empty Audit Descriptor, and Audit Descriptor with a statistics token, or no Audit Descriptor.
- SUBTRACT reply with no statistics—When it receives a SUBTRACT request with an Empty Audit Descriptor, the Net-Net BG responds with a SUBTRACT reply that has no Statistics Descriptor. When it receives a SUBTRACT request without an Audit Descriptor (or with an Audit Descriptor that has the Statistics token included), the Net-Net BG responds with a SUBTRACT reply with a Statistics Descriptor as currently observed.

Backwards and Forwards Octet Strings for ds/dscp Parameters

By default, the Net-Net BG supports the ds/dscp parameter (DiffServ Code Point or TOS bits) by allowing either hexidecimal or decimal values. If the value is proceeded by 0x, then the Net-Net BG parses it as hexidecimal; if not, then it is parsed as a decimal value. In the case, the most significant digit is always on the left.

However, ETSI TI-SPAN has defined a package for ds/dscp stating that the parameter should be interpreted as an octet string. So, in addition to its default behavior, the Net-Net BG supports a vBG configuration options that allow two more ways for parsing to occur. These options are:

- **hex-dscp**—Enables the Net-Net BG to interpret the ds/dscp parameter in an H.248 message as a hexidecimal value, where the most significant nibble is on the left (i.e., the leftmost character of value) and the least significant on the right.
- **reverse-hex-dscp**—Enables the Net-Net BG to interpret the ds/dscp parameter in an H.248 message as a hexidecimal value, where the least significant nibble is on the left and the most significant is on the right(i.e., contained in the rightmost character of value).

If both of these options are set, then the **reverse-hex-dscp** overrides the **hex-dscp** option.

The two options account for the fact that the ETSI TI-SPAN definition leaves room for interpretation about whether the most significant hexidecimal digit (as well as the binary digits specified therein) should appear on the left or the right. As you can see, the options you configure to control ds/dscp treatment account for both big-endian and little-endian octet string interpretation of the definition.

ACLI Instructions and Examples

This section shows you how to configure the vBG configuration options that support backwards and forwards octet string for the ds/dscp parameter. Note that if both of these options are set, then the **reverse-hex-dscp** overrides the **hex-dscp** option. And if you do not set an option for ds/dscp treatment, then the Net-Net BG resorts to its default behavior.

To configure support for interpreting the ds/dscp as a big-endian octet string:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Type **vbg-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# vbg-config
   ACMEPACKET(vbg-config)#
   ```

   If you are adding support to an existing configuration, then you need to select the configuration before you edit it.

4. **options**—Set the options parameter by typing **options**, a <Space>, the option name **hex-dscp** with a "plus" sign in front of it. Then press <Enter>.

   ```
   ACMEPACKET(vbg-config)# options +hex-dscp
   ```

   If you type **options** and then the option value for either of these entries without the "plus" sign, you will overwrite any previously configured options. In order to append the new option to the configuration's options list, you must prepend the new option with a "plus" sign as shown in the previous example.

5. Save and activate your configuration.

To configure support for interpreting the ds/dscp as a little-endian octet string:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ```

```
ACMEPACKET(media-manager)#
```

3. Type **vbg-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# vbg-config
   ACMEPACKET(vbg-config)#
   ```

   If you are adding support to an existing configuration, then you need to select the configuration before you edit it.

4. **options**—Set the options parameter by typing **options**, a <Space>, the option name **reverse-hex-dscp** with a "plus" sign in front of it. Then press <Enter>.

   ```
   ACMEPACKET(vbg-config)# options +reverse-hex-dscp
   ```

   If you type **options** and then the option value for either of these entries without the "plus" sign, you will overwrite any previously configured options. In order to append the new option to the configuration's options list, you must prepend the new option with a "plus" sign as shown in the previous example.

5. Save and activate your configuration.

Failover ServiceChange Command

During the process of transitioning from one state to another, the Net-Net BG communicates its service state. You can configure the ServiceChange method sent when a failover occurs between two Net-Net BGs deployed as an HA node.

ACLI Instructions and Examples

You configure the ServiceChange method sent when a failover occurs between two Net-Net BGs deployed as an HA node by setting an option in the vBG configuration. When you set this option, the Net-Net BG sends the Disconnected method instead of Restart.

To set the ServiceChange method to Disconnected:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **media-manager** and press <Enter> to access the media-manager path.

   ```
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)#
   ```

3. Type **vbg-config** and press <Enter>.

   ```
   ACMEPACKET(media-manager)# vbg-config
   ACMEPACKET(vbg-config)#
   ```

4. **options**—Set the options parameter by typing **options**, a <Space>, the option name **send-disconnected-on-failover** with a "plus" sign in front of it. Then press <Enter>.

   ```
   ACMEPACKET(vbg-config)# options +send-disconnected-on-failover
   ```

   If you type **options** and then the option value for either of these entries without the "plus" sign, you will overwrite any previously configured options. In order to append the new option to the configuration's options list, you must prepend the new option with a "plus" sign as shown in the previous example.

5. Save and activate your configuration.

Release S-C[x]6.2.0 Additions

The following sections note additional BG functionality added in Release Version S-CX6.4.0.

About Your Net-Net Border Gateway and IPv6

The Net-Net BG supports IPv6. Ideally, IPv6 support would be a simple matter of configuring IP addresses of the version type you want in the configurations where you want them. While this is the case for some configuration areas, in others you will need to take care with—for example—the format of your IPv6 address entries or where parameters must be configured with IP addresses of the same version type.

This section explains the changes to the ACLI of which you need to be aware as you start to use IPv6 on your BG. Note that not all configurations and their parameters are available for IPv6 use.

Licensing

IPv6 is a licensed feature on the Net-Net 3800 and Net-Net 4500. If you want to add this license to a system, then contact your sales engineering for information related to the license. Once you have the license information, refer to the Getting Started chapter of the Net-Net 4000 ACLI Configuration Guide of instructions about how to add a license.

You do not need to take action if you are working with a new system with which the IPv6 license was purchased.

Updated ACLI Help Text

As you complete configuration work and perform monitoring tasks on your system, you might note that there have been changes to the help text to reflect the addition of IPv6 support. These changes are minor, but nonetheless reflect feature support.

In the ACLI that supports only IPv4, there are many references to that version as the accepted value for a configuration parameter or other IPv4-specific languages. For IPv6 support, these references have been edited. For example, rather than providing help that refers specifically to IPv4 addresses when explaining what values are accepted in an ACLI configuration parameter, you will now see an <ipAddr> note.

IPv6 Address Configuration

This section calls out the configurations and parameters for which you can enter IPv6 addresses. In this first IPv6 implementation, the complete range of system configurations and their parameters are available for IPv6 use.

The Net-Net BG follows RFC 3513 its definition of IPv6 address representations. Quoting from that RFC, these are the two forms supported:

- The preferred form is x:x:x:x:x:x:x:x, where the 'x's are the hexadecimal values of the eight 16-bit pieces of the address. Examples:

  FEDC:BA98:7654:3210:FEDC:BA98:7654:3210

  1080:0:0:0:8:800:200C:417A

  Note that it is not necessary to write the leading zeros in an individual field, but there must be at least one numeral in every field (except for the case described in 2.).

- Due to some methods of allocating certain styles of IPv6 addresses, it will be common for addresses to contain long strings of zero bits. In order to make writing addresses containing zero bits easier a special syntax is available to

compress the zeros. The use of "::" indicates one or more groups of 16 bits of zeros. The "::" can only appear once in an address. The "::" can also be used to compress leading or trailing zeros in an address. For example, the following addresses: 1080:0:0:0:8:800:200C:417A a unicast address FF01:0:0:0:0:0:0:101 a multicast address

0:0:0:0:0:0:0:1 the loopback address

0:0:0:0:0:0:0:0 the unspecified addresses

may be represented as:

1080::8:800:200C:417A a unicast address

FF01::101 a multicast address

::1 the loopback address

:: the unspecified addresses

Access Control

These are the IPv6-enabled parameters in the **access-control** configuration.

Parameter	Entry Format
source-address	<ip-address>[/<num-bits>][:<port>[/<port-bits>]]
destination-address	<ip-address>[/<num-bits>][:<port>[/<port-bits>]]

Host Route

These are the IPv6-enabled parameters in the **host-route** configuration.

Parameter	Entry Format
dest-network	<ipv4> | <ipv6>
netmask	<ipv4> | <ipv6>
gateway	<ipv4> | <ipv6>

Local Policy

These are the IPv6-enabled parameters in the **local-policy** configuration.

Parameter	Entry Format
from-address	<ipv4> | <ipv6> | POTS Number, E.164 Number, hostname, wildcard
to-address	<ipv4> | <ipv6> | POTS Number, E.164 Number, hostname, wildcard

Network Interface

These are the IPv6-enabled parameters in the **network-interface** configuration.

Parameter	Entry Format
hostname	<ipv4> | <ipv6> | hostname
ip-address	<ipv4> | <ipv6>
pri-utility-addr	<ipv4> | <ipv6>
sec-utility-addr	<ipv4> | <ipv6>
netmask	<ipv4> | <ipv6>
gateway	<ipv4> | <ipv6>
sec-gateway	<ipv4> | <ipv6>
dns-ip-primary	<ipv4> | <ipv6>
dns-ip-backup1	<ipv4> | <ipv6>
dns-ip-backup2	<ipv4> | <ipv6>
add-hip-ip	<ipv4> | <ipv6>
remove-hip-ip	<ipv4> | <ipv6>
add-icmp-ip	<ipv4> | <ipv6>
remove-icmp-ip	<ipv4> | <ipv6>

Realm Configuration

These are the IPv6-enabled parameters in the **realm-config**.

Parameter	Entry Format
addr-prefix	[<ipv4> | <ipv6>]/prefix

Session Agent

These are the IPv6-enabled parameters in the **session-agent** configuration.

Parameter	Entry Format
hostname	<ipv4> | <ipv6>
ip-address	<ipv4> | <ipv6>

SIP Configuration

These are the IPv6-enabled parameters in the **session-config**.

Parameter	Entry Format
registrar-host	<ipv4> | <ipv6> | hostname | *

SIP Interface>SIP Ports

These are the IPv6-enabled parameters in the **sip-interface>sip-ports** configuration.

Parameter	Entry Format
address	<ipv4> | <ipv6>

Steering Pool

These are the IPv6-enabled parameters in the **steering-pool** configuration.

Parameter	Entry Format
ip-address	<ipv4> | <ipv6>

System Configuration

These are the IPv6-enabled parameters in the **system-config**.

Parameter	Entry Format
default-v6-gateway	<ipv6>

IPv6 Default Gateway

In the system configuration, you configure a default gateway—a parameter that now has its own IPv6 equivalent.

To configure an IPv6 default gateway:

1. In Superuser mode, type **configure terminal** and press <Enter>.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)#
   ```

2. Type **system** and press <Enter>.

   ```
   ACMEPACKET(configure)# system
   ACMEPACKET(system)#
   ```

3. Type **system-config** and press <Enter>.

```
ACMEPACKET(system)# system-config
ACMEPACKET(system-config)#
```

4. **default-v6-gateway**—Set the IPv6 default gateway for this Net-Net SBC. This is the IPv6 egress gateway for traffic without an explicit destination. The application of your Net-Net SBC determines the configuration of this parameter.
5. Save your work.

Network Interfaces and IPv6

You set many IP addresses in the network interface, one of which is the specific IP address for that network interface and others that are related to different types of management traffic. This section outlines rules you must follow for these entries.

- For the **network-interface ip-address** parameter, you can set a single IP address. When you are working with an IPv6-enabled system, however, note that all other addresses related to that network-interface IP address must be of the same version.
- Heterogeneous address family configuration is prevented for the **dns-ip-primary**, **dns-ip-backup1**, and **dns-ip-backup2** parameters.
- For HIP addresses (**add-hip-ip**), you can use either IPv4 or IPv6 entries.
- For ICMP addresses (**add-icmp-ip**), you can use either IPv4 or IPv6 entries.
- For Telnet (**add-telnet-ip**), FTP (**add-ftp-ip**), and SNMP (**add-snmp-ip**), you are not allowed to use IPv6; your entries MUST use IPv4.

IPv6 Reassembly and Fragmentation Support

As it does for IPv4, the Net-Net SBC supports reassembly and fragmentation for large signaling packets when you enable IPV6 on your system.

The Net-Net SBC takes incoming fragments and stores them until it receives the first fragment containing a Layer 4 header. With that header information, the Net-Net SBC performs a look-up so it can forward the packets to its application layer. Then the packets are re-assembled at the applications layer. Media fragments, however, are not reassembled and are instead forwarded to the egress interface.

On the egress side, the Net-Net SBC takes large signaling messages and encodes it into fragment datagrams before it transmits them.

Note that large SIP INVITE messages should be sent over TCP. If you want to modify that behavior, you can use the SIP interface's option parameter **max-udp-length=xx** for each SIP interface where you expect to receive large INVITE packets.

Other than enabling IPv6 on your Net-Net SBC, there is no configuration for IPv6 reassembly and fragmentation support. It is enabled automatically.

Access Control List Support

The Net-Net SBC supports IPv6 for access control lists in two ways:

- For static access control lists that you configure in the **access-control** configuration, your entries can follow IPv6 form. Further, this configuration supports a prefix that enables wildcarding the source IP address.
- Dynamic ACLs are also supported; the Net-Net SBC will create ACLs for offending IPv6 endpoints.

Data Entry

When you set the **source-address** and **destination-address** parameters in the **access-control** configuration, you will use a slightly different format for IPv6 than for IPv4.

For the **source-address**, your IPv4 entry takes the following format: `<ip-address>[/<num-bits>][:<port>[/<port-bits>]]`. And for the **destination-address**, your IPv4 entry takes this format: `<ip-address>[:<port>[/<port-bits>]]`.

Since the colon (:) in the IPv4 format leads to ambiguity in IPv6, your IPv6 entries for these settings must have the address encased in brackets (`[]`): `[7777::11]/64:5000/14`.

In addition, IPv6 entries are allowed up to 128 bits for their prefix lengths.

The following is an example access control configuration set up with IPv6 addresses.

```
ACMEPACKET(access-control)# done
access-control
        realm-id                        net7777
        description
        source-address                  7777::11/64:5060/8
        destination-address             8888::11:5060/8
        application-protocol            SIP
        transport-protocol              ALL
        access                          deny
        average-rate-limit              0
        trust-level                     none
        minimum-reserved-bandwidth      0
        invalid-signal-threshold        10
        maximum-signal-threshold        0
        untrusted-signal-threshold      0
        deny-period                     30
```

DNS Support

The Net-Net SBC supports the DNS resolution of IPv6 addresses; in other words, it can request the AAAA record type (per RFC 1886) in DNS requests. In addition, the Net-Net SBC can make DNS requests over IPv6 transport so that it can operate in networks that host IPv6 DNS servers.

For mixed IPv4-IPv6 networks, the Net-Net SBC follows these rules:

- If the realm associated with the name resolution is an IPv6 realm, the Net-Net SBC will send the query out using the AAAA record type.
- If the realm associated with the name resolution is an IPv4 realm, the Net-Net SBC will send the query out using the A record type.

In addition, heterogeneous address family configuration is prevented for the **dns-ip-primary**, **dns-ip-backup1**, and **dns-ip-backup2** parameters.

RADIUS Support for IPv6

The Net-Net BG's RADIUS support now includes:

- RADIUS CDR generation for SIPv6-SIPv6 and SIPv6-SIPv4 calls
- IPv6-based addresses in RADIUS CDR attributes

This means that for the CDR attributes in existence prior to the introduction of IPv6 to the Net-Net 3800/4500 are mapped to the type `ipaddr`, which indicates four-byte field. The sixteen-byte requirement for IPv6 addresses is now supported, and there are a parallel set of attributes with the type `ipv6addr`. Attributes 155-170 are reserved for the IPv6 addresses.

NAS addresses use the number 95 to specify the NAS-IPV6-Address attribute. And local CDRs now contain IPv6 addresses.

Supporting RADIUS VSAs

The following VSAs have been added to the RADIUS dictionary to support IPv6.

```
Acme-Flow-In-Src-IPv6_Addr_FS1_F        155     ipv6addr      Acme
Acme-Flow-In-Dst-IPv6_Addr_FS1_F        156     ipv6addr      Acme
Acme-Flow-Out-Src-IPv6_Addr_FS1_F       157     ipv6addr      Acme
Acme-Flow-Out-Dst-IPv6_Addr_FS1_F       158     ipv6addr      Acme
Acme-Flow-In-Src-IPv6_Addr_FS1_R        159     ipv6addr      Acme
Acme-Flow-In-Dst-IPv6_Addr_FS1_R        160     ipv6addr      Acme
Acme-Flow-Out-Src-IPv6_Addr_FS1_R       161     ipv6addr      Acme
Acme-Flow-Out-Dst-IPv6_Addr_FS1_R       162     ipv6addr      Acme
Acme-Flow-In-Src-IPv6_Addr_FS2_F        163     ipv6addr      Acme
Acme-Flow-In-Dst-IPv6_Addr_FS2_F        164     ipv6addr      Acme
Acme-Flow-Out-Src-IPv6_Addr_FS2_F       165     ipv6addr      Acme
Acme-Flow-Out-Dst-IPv6_Addr_FS2_F       166     ipv6addr      Acme
Acme-Flow-In-Src-IPv6_Addr_FS2_R        167     ipv6addr      Acme
Acme-Flow-In-Dst-IPv6_Addr_FS2_R        168     ipv6addr      Acme
Acme-Flow-Out-Src-IPv6_Addr_FS2_R       169     ipv6addr      Acme
Acme-Flow-Out-Dst-IPv6_Addr_FS2_R       170     ipv6addr      Acme
```

Hide-Media-Update

The Real-Time Transport Protocol uses timestamps, sequence numbers, and endpoint identifiers (referred to as SSRCs) to enable call participants to correlate incoming streams. In certain instances, the SSRC can change unexpectedly with this change also being reflected by unexpected variations in the timestamp and sequence numbers. Such unexpected variations can cause certain VoIP terminals to freeze, rendering them inoperable.

The following call flow and illustrations help to visualize the scenario described above. In this case, a call is initially established between UE#1 and UE#2. The call transits a BG which is controlled by a decomposed SBC. TerminationID IP#1, towards UE#1, and, terminationID IP#2, towards UE#2, are created within the same Context ID.

During the call, the decomposed SBC MODIFYs the remote descriptor of TerminationID IP#2, and redirects the media stream to UE#3. The media between the BG IP#1 towards UE#1 remains unchanged.

At this point, that is after the redirection from UE#2 to UE#3, it is important that the media sent towards UE#1 appears to be seamless, which can be guaranteed if the SSRC is maintained, and continuity of timestamps and sequence numbers is maintained.

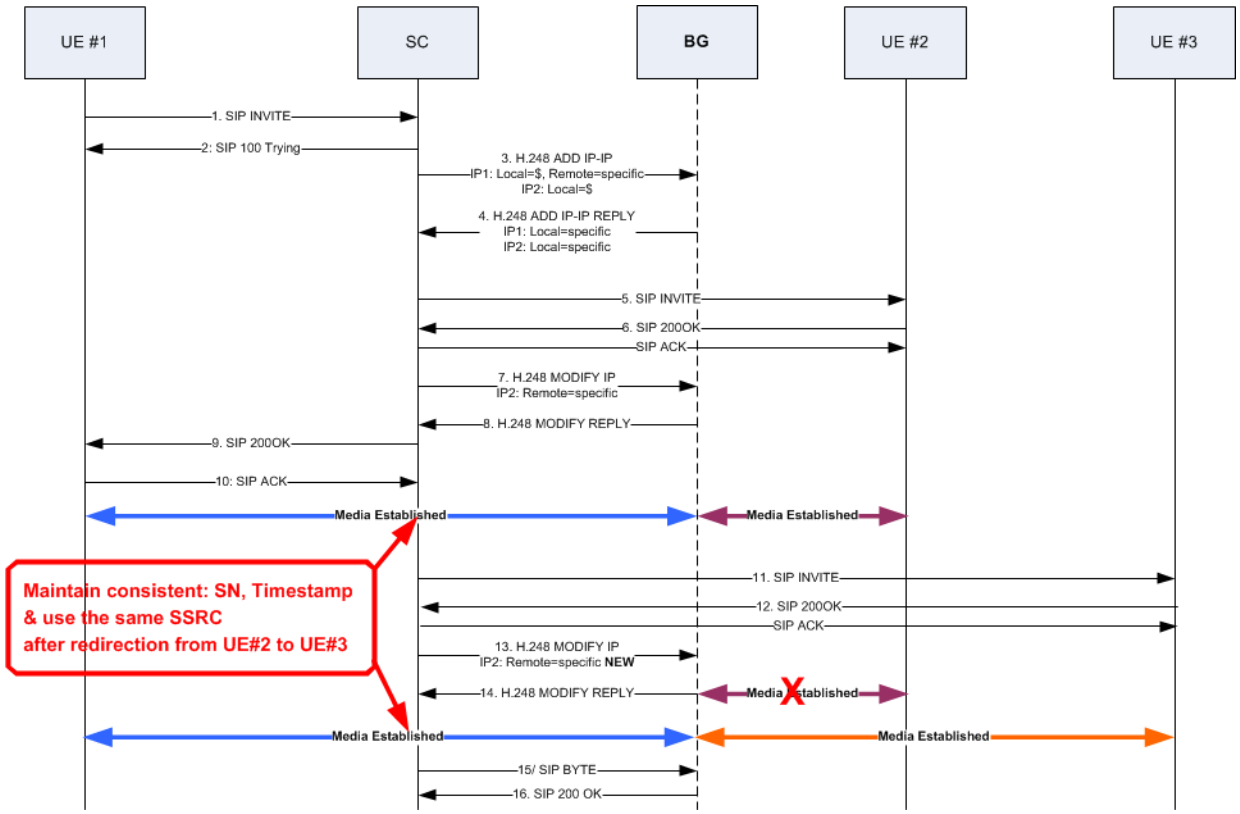

The following two illustrations depict how the media signalled in the above call flow is redirected via the BG from UE#2 to UE#3.

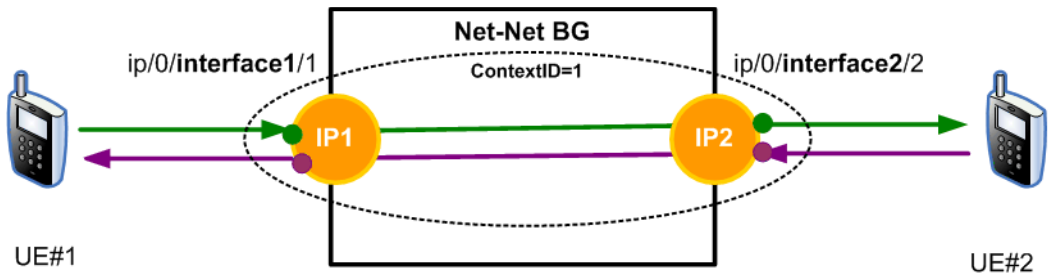

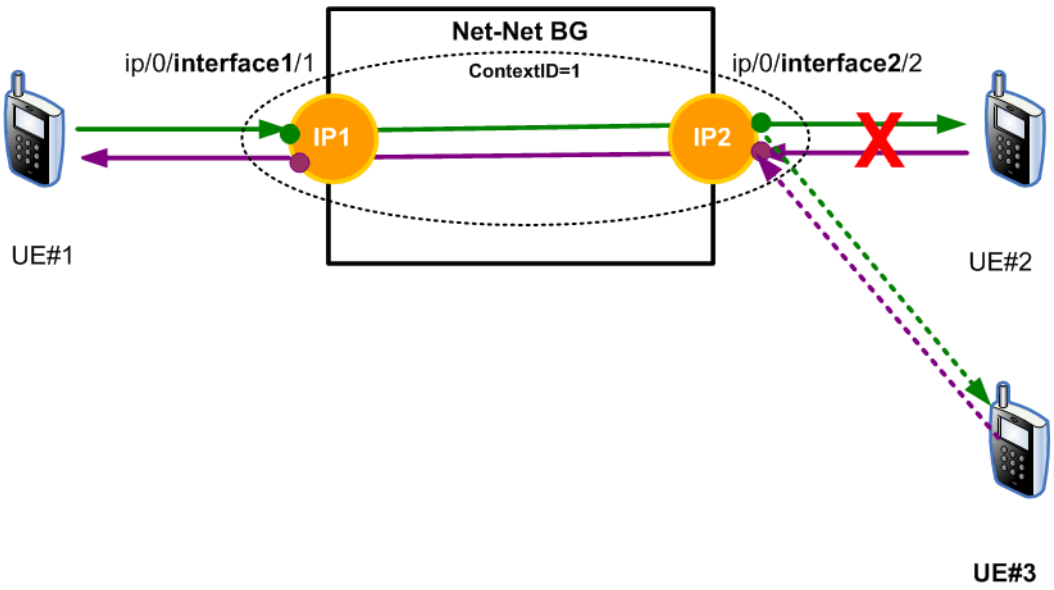

The following chart shows the relationship between HMU and realms/BG virtual interfaces.

Ingress Realm HMU Enabled	Egress Realm HMU Enabled	Result
No	No	HMU not used
No	Yes	HMU applied to response RTP
Yes	No	HMU not applied to response RTP
Yes	Yes	HMU applied to response RTP

Configuring hide-media-update

Use the following procedure to enable hide-media-update, which is enabled for specific realms. By default hide-media-update is disabled.

1. From *superuser* mode use the following ACLI command sequence to access *realm-config* configuration mode.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)# realm-config
   ACMEPACKET(realm-config)#
   ```

2. Use the **select** command to specify the target realm.
3. Use the **hide-egress-media-update** parameter with to enable hide-media-update support.

   ```
   ACMEPACKET(realm-config)# hide-egress-media-update enabled
   ACMEPACKET(realm-config)#
   ```

4. Use **done**, **exit**, and **verify-config** to complete this configuration.

Media Mirroring

Version 6.2.0 provides a new Media Mirroring H.248 package that enables the decomposed SBC to command the BG to replicate the current media stream and transport the replicated stream to a media recording server identified by its IP address. Such instructions are communicated via the H.248 interface, and can be issued during call
set-up (ADD command) or mid-call (MODIFY command).

The triggering information will be conveyed from the decomposed SBC in an H.248 LocalControlDescriptor, which is specified on a per termination basis.

H248 descriptor example:

```
LocalControl {
     Mode = SendReceive,
     gm/rsb = ON,
     mm/destIP= 192.168.202.9,
     mm/srcInt = core
}
```

Upon receiving instructions from the decomposed SBC, the BG replicates all RTP and RTCP streams associated with the specified termination, and directs the streams to the media recording server identified by the mm/destIP parameter in the LocalControlDescriptor. The associated mm/srcInt parameter identifies the realm through which the replicated media streams are sent. Replication of RTP content is fully supported; replication of RTCP content is on a best-effort basis. Generally RTCP content is not replicated in peering scenarios, or in access scenarios with Hosted NAT.

When a termination is designated for media mirroring, all streams related to that termination are replicated.Consequently, if a termination is identified by an ADD command, and a subsequent MODIFY command adds a new stream that new stream is replicated. Media mirroring continues through the termination lifetime and ceases when the termination is SUBTRACTed from the context.

License Requirements

Media mirroring functionality requires the presence of the Media Mirroring license. In the absence of this license, functionality is not available.

Configuring Media Mirroring

No BG configuration is required to implement Media Mirroring.

Blocking RTCP Traffic

Under certain conditions it may be advantageous or required to block the entrance of Real-Time Transport Control Protocol (RTCP) traffic from a specific realm or realms. Use the following procedure to do so.

1. From *superuser* mode use the following ACLI command sequence to access *realm-config* configuration mode.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)# realm-config
   ACMEPACKET(realm-config)#
   ```

2. Use the **select** command to specify the target realm.

3. Use the **block-rtcp** command with to block introduction of RTCP traffic into the selected realm.

```
ACMEPACKET(realm-config)# block-rtcp enabled
ACMEPACKET(realm-config)#
```

4. Use **done**, **exit**, and **verify-config** to complete this configuration.

Message Session Relay Protocol

With Version S-CX6.4.0, the BG supports SIP-based Message Session Relay Protocol (MSRP) sessions over a TCP transport layer. MSRP sessions are initiated by Session Description Protocol (SDP) messages exchanged via the Session Initiation Protocol (SIP) offer/answer model. MSRP usage with SDP and SIP is described in Section 8 of RFC 4975, *The Message Relay Protocol*, and in RFC 4145, *TCP-Based Media Transport in the Session Description Protocol (SDP)*.

Within an IMS deployment, the BG serves as a relay function steering MSRP data flows from incoming to outgoing realms under the supervision of the SBC. Specifically the BG:

1. Supports TCP media transport to include setup and termination of incoming and outgoing TCP connections. The BG does not participate in end-to-end TCP connections, nor does it terminate any TCP traffic; TCP packets are simply relayed across the BG in the same fashion as UDP packets.
2. Supports the media type "*message*".
3. Supports the transport protocol "*TCP/MSRP*".
4. Supports the SDP "*path*" attribute.
5. Supports the same media-related functions (for example, DSCP marking, rate limiting, DoS protection, etc.) for both UDP and TCP.
6. Supports highly-available (HA) TCP service; TCP media flows are maintained across a BG failover.

The following two sections illustrate MSRP signalling — end-to-end signalling between the MSRP originator and responder, and signalling between the decomposed SBC and BG used to establish MSRP data flows.

MSRP End-to-End Signalling

A sample end-to-end RFC 4975-compliant Offer/Answer SDP exchange for an MSRP session is shown below.

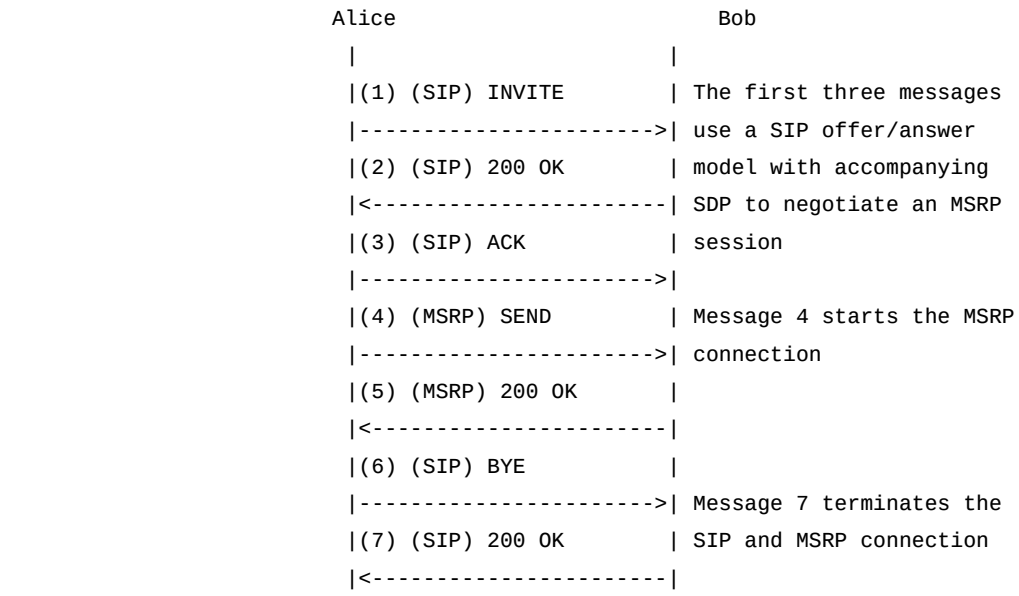

1. Alice sends an INVITE request with accompanying SDP to Bob.

```
Alice->Bob (SIP):

INVITE sip:bob@example.com
SDP:
v=0
o=alice 2890844557 2890844559 IN IP4 alicepc.example.com
s=
c=IN IP4 alicepc.example.com
m=message 7777 TCP/MSRP *
a=accept-types:text/plain
a=path:msrp://alicepc.example.com:7777/iau39soe2843z;tcp
```

   The SDP media (m) line is defined in RFC 4975, and adheres to the format

   m=<media> <port> <protocol> <format-list>

   MSRP operations require the following values:

   media = *message*

   protocol = *TCP/MSRP*

   format-list =*

   port = the TCP port (7777 in the SDP example, although any valid port number can be specified) monitored by the message originator for a response to the SDP offer

   The required SDP attributes, *accept-types* and *path*, are also defined in RFC 4975.

   accept-types contains a list of media types that the message originator is willing to receive. It may contain zero or more registered media-types, or an * wildcard character in a space-delimited string.

path contains the MSRP URI of the message originator. An MSRP URI is constructed as shown below.

scheme = *msrp*

//

address = IP address of the message originator, or

= FQDN of the message originator

:

port = the port (7777 in the SDP example, although any valid port number can be specified) monitored by the message originator for MSRP responses

session-id = a random local value generated by the message originator used to produce an ephemeral MSRP URI lasting only for the duration of the current MSRP session

;

protocol = *tcp*

2. Bob accepts the SDP offer, generates a local session-id (contained in his MSRP URI specified by the *path* attribute), and issues a 200 OK response to Alice.

```
Bob->Alice (SIP):

SIP/2.0 200 OK
SDP:
v=0
o=bob 2890844612 2890844616 IN IP4 bob.example.com
s=
c=IN IP4 bob.example.com
m=message 8888 TCP/MSRP *
a=accept-types:text/plain
a=path:msrp://bob.example.com:8888/9di4eae923wzd;tcp
Bob->Alice (SIP):
```

The port parameter in the Media line indicates that Bob listens for MSRP messages on TCP port 8888

3. Alice ACKs Bob's response, establishing a SIP session between the two.

```
Alice->Bob (SIP):

ACK sip:bob@example.com
```

4. Alice initiates an MSRP session with an MSRP SEND request to Bob.

```
Alice->Bob (MSRP):

MSRP d93kswow SEND
To-Path: msrp://bob.example.com:8888/9di4eae923wzd;tcp
From-Path: msrp://alicepc.example.com:7777/iau39soe2843z;tcp
Message-ID: 12339sdqwer
Byte-Range: 1-16/16
Content-Type: text/plain
Hi, I'm Alice!
-------d93kswow$
```

All MSRP requests begin with the MSRP start line, which contains three elements.

protocol-id = *MSRP*

transaction-id = an ephemeral transaction identifier (*d93kswow* in the following MSRP example) used to correlate MSRP requests and responses, and to frame the contents of the MSRP message

method = *SEND* (MSRP method that supports data transfer)

The MSRP start line is followed by the To-Path and From-Path headers, which contain destination and source addresses — the MSRP URIs exchanged during the MSRP negotiation.

The Message-ID header contains a random string generated by the message originator. This ephemeral value whose lifetime is measured from message start to message end, is used to correlate MSRP status reports with a specific message, and to re-assemble MSRP message fragments (chunks in MSRP terminology).

The Byte-Range header contains the message length, in bytes, and the specific byte range carried by this message. Contents of this header are generally of interest only if the message has been fragmented.

The Content-Type header describes the message type, and must conform to the results of the MSRP negotiation.

The actual message follows the Content-Type header.

Finally, the SEND request is closed with an end-line of seven hyphens, the transaction-id, and a

$ to indicate that this request contains the end of a complete message, or

+ to indicate that this request does not contain the message end

5. Bob acknowledges receipt with an MSRP 200 OK response to Alice.

```
Bob->Alice (MSRP):

MSRP d93kswow 200 OK
To-Path: msrp://alicepc.example.com:7777/iau39soe2843z;tcp
From-Path: msrp://bob.example.com:8888/9di4eae923wzd;tcp
-------d93kswow$
```

Note that the response includes the initiator-originated transaction-id, *d93kswow*.

6. Alice sends a BYE request to Bob.

```
Alice->Bob (SIP):

BYE sip:bob@example.com
```

Alice sends a BYE request to terminate the SIP session and MSRP sessions. Alice can, of course, send more SEND requests to Bob before sending the BYE.

7. Bob sends a 200 OK response to Alice.

```
Bob->Alice (SIP):

SIP/2.0 200 OK
```

The SIP session and the MSRP session are terminated.

The following are examples of H.248 messages used in the setup of MSRP.

Request:

```
MEGACO/2 [192.168.200.211]:2944
Transaction=2 {
 Context=$ {
  Add=ip/$/net1/$ {
   Media {
    Stream=1 {
     LocalControl {
      Mode=Inactive
     },
     Local {v=0
c=IN IP4 $
m=message $ TCP/MSRP *
}
    }
   }
  },
  Add=ip/$/net200/$ {
   Media {
    Stream=1 {
     LocalControl {
      Mode=SendReceive
     },
     Local {v=0
c=IN IP4 $
m=message $ TCP/MSRP *
},
     Remote {v=0
c=IN IP4 192.168.200.211
m=message 23 TCP/MSRP *
}
    }
   }
  }
 }
}
```

Reply:

```
MEGACO/2 [192.168.200.212]:2944
Reply=2{
 Context=4{
  Add=ip/1/net1/6{
   Media{
```

```
Stream=1{Local{
v=0
c=IN IP4 192.168.1.212
m=message 10002 TCP/MSRP *
}}
   }
  },
  Add=ip/1/net200/8{
   Media{
Stream=1{Local{
v=0
c=IN IP4 192.168.200.212
m=message 10002 TCP/MSRP *
}}
   }
  }
 }
}
```

Version S-CX6.4.0 provides support for H.248 MODIFY messages that can change the media and transport type initially passed between the decomposed SBC and the BG. This capability is required in certain deployments where the SBC does not inform the BG of the *requested* media and transport type proposed by the session initiator, but rather waits for the SDP response, which definitively identifies the negotiated media and type, to be returned by the session responder.

In these deployments, the SBC sends, and the BG receives an initial *default* H.248 ADD message with the media type identified as audio/UDP. After the requested media and transport type have been negotiated, the SBC sends the BG and H.248 MODIFY message which contains the agreed-upon media type and transport — MSRP/TCP in the case of an MSRP message session.

ACLI Instructions and Examples

Use the following procedure to enable the BG for MSRP support.

1. From *superuser* mode use the following ACLI command sequence to access *bgf-config* configuration mode.

   ```
   ACMEPACKET# configure terminal
   ACMEPACKET(configure)# media-manager
   ACMEPACKET(media-manager)# bgf-config
   ACMEPACKET(bgf-config)#
   ```

2. **options**—Use the **options** command with the **msrp-star-support** argument to enable MSRP support.

   ```
   ACMEPACKET(bgf-config)# options +msrp-star-support
   ACMEPACKET(bgf-config)#
   ```

3. Use **done**, **exit**, and **verify-config** to complete MSRP configuration.

RTCP Flows/H.248 Termination Modes

With Version S-CX6.4.0, the BG changes previous behavior which allowed the establishment of Real-Time Control Protocol (RTCP) flows only when Real-Time Protocol (RTP) flows were allowed. This new behavior enables RTCP flows through a ContextID independent of the termination mode (*inactive*, *recvonly*, *sendonly*, *sendrecv*) on either of the H.248 terminations with the context. Consequently, the BG now supports bi-directional RTCP exchange across all permutations or termination modes.

Version S-CX6.4.0 provides an ACLI **options** command to revert to the prior default behavior (RTCP only when RTP is allowed) if required by local network conditions.

Note: This command is made available on a temporary basis, and is subject to deprecation.

Use the following procedure to revert to prior behavior which allowed RTCP flows only when RTP flows were allowed.

1. From *superuser* mode use the following ACLI command sequence to access *bgf-config* configuration mode.

```
ACMEPACKET# configure terminal
ACMEPACKET(configure)# media-manager
ACMEPACKET(media-manager)# bgf-config
ACMEPACKET(bgf-config)#
```

2. **options**—Use the **options** command with the **srtp-only-when-rtp** argument to enable MSRP support.

```
ACMEPACKET(bgf-config)# options +rtcp-only-when-rtp
ACMEPACKET(bgf-config)#
```

Use **done**, **exit**, and **verify-config** to complete the restoration of previous behavior.

Relatch Support

In topologies where an endpoint is behind a firewall or NAT device, the BG has no knowledge of the source IP address and port. This is because intervening firewalls or NATs substitute local address and port values for the original source address available only within the SDP body. In order for the BG to relay RTP traffic to the endpoint, the NAT address and port must be learned from the first RTP message received. This process of setting the RTP flows to the source address/port from the first RTP message is called *latching*. Once a flow is latched to a source address/port, it will not forward traffic received from any other source address/port.

Using a SIP audio session as an example, when an endpoint is transferred, put on hold, or the audio flow is modified in any fashion, the previously latched flows are deleted and new latching flows reflecting the new media parameters are added. These new flows will latch to the first RTP message that is received. A problem occurs if the original endpoint continues to send RTP after the flows are modified, In that case the original endpoint can latch to the old address/port and prevent traffic from being received from a new address/port.

To address this potential problem, Version S-CX6.4.0 supports the relatch signal defined in H.248.37, *Gateway control protocol: IP NAPT traversal package*.

Relatch Signal

When the Latch signal with an assigned value of RELATCH, the BG will check the incoming flow for a change of

- source IP address and same or different port number, or
- same IP address and a different port number

Restricted latching, as defined in H.248.37, is also supported if *gm/sam=x.x.x.x* (source address mask field) and *gm/saf=on* (source address filtering field) are provided to the BG. If provided, the BG only latches onto an RTP packet matching the corresponding IP source address mask.

The relatching process does not cause transmission of packets from the old source address and port to be blocked prior to relatching to a new source address or port unless received *gm/saf* and *gm/sam* fields specify a source address filter that blocks packets from the old IP address. Once relatched to a new source address and port, packets arriving from the old address and port are to be blocked as described in Section 6.6.3 of H.248.37.

Relatch Signal

Relatching is desirable for all protocols, so it is implemented in a generic way that can be utilized for all media regardless of the signaling protocol. Any MODIFY command will trigger relatching if the modified flow has already been latched and the new endpoint is behind a NAT. The previous implementation of the H.248 relatch signal on the SBC was to reset the latching on the flow. With Version S-CX6.4.0 (and later versions), the relatch (or latch) H.248 is now used to denote that the endpoint is behind a NAT and thus triggers a generic relatching behavior.

Relatch NatFlows

When the BG receives a MODIFY command from the decomposed SBC, indicating that a modified flow has already been latched or that a new endpoint is marked as being behind a NAT, or H.248 has re/latch set, the flow needs to perform relatching. Previously, the BG would delete the old latched NatFlow and install a new latching NatFlow.

For relatching, the old source address and port information is saved and then the old NatFlow is reestablished along with the new latching NatFlow — resulting in the installation of old and new NatFlows. Note that if the media-manager configuration **hnt-rtcp** is enabled, the old RTCP NatFlow is also restored and a new latching RTCP NatFlow is added. The following diagram shows the installed NatFlows for relatching.

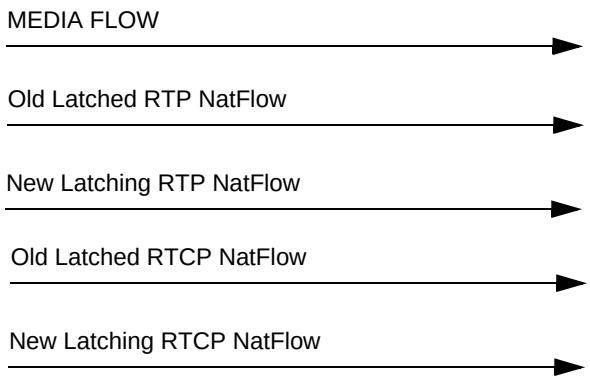

Latching with Relatch NatFlows

By reestablishing the old NatFlows that had already been latched, any existing RTP traffic from the old latched source can continue to be forwarded from the old source IP address/port. If traffic is received from a new source address or port does not match existing latched NatFlow), that traffic is forwarded on the new latching NatFlow. This NatFlow latches to the new source IP address/port, and then the old NatFlow is deleted so that traffic from the old source no longer is forwarded. In the following example an old source (192.168.1.1:6000) is sending to the destination (192.168.200.1:6000) and the new source (192.168.1.2:6000) begins sending.

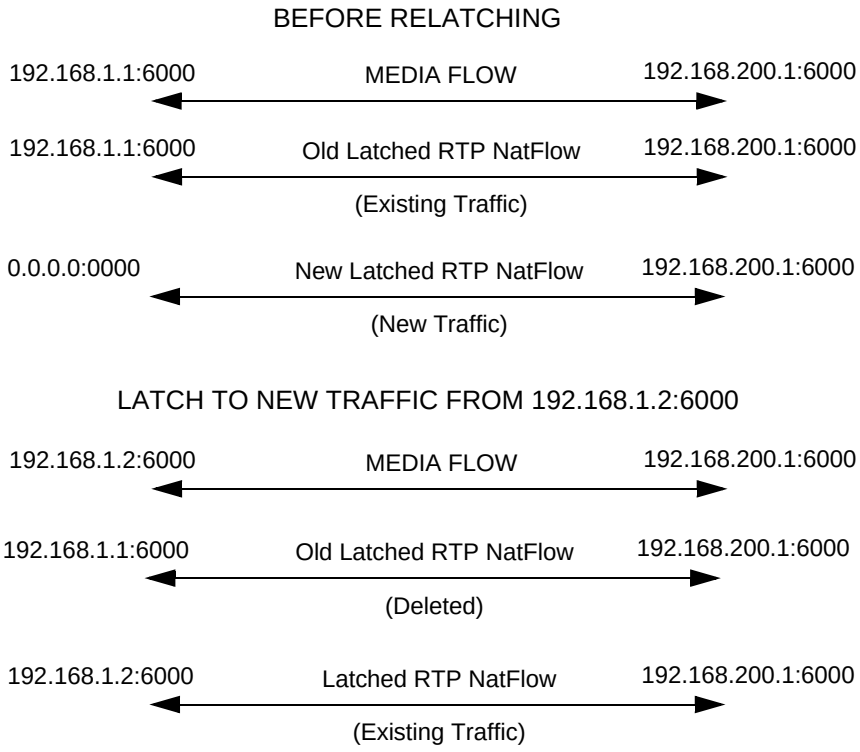

Timeout on Relatch NatFlows

As relatching requires double the amount of NatFlows, the lifespan of these flows are subject to specific timer limitations. The initial guard timer for a new latching NatFlow is hardcoded to 30 seconds. If the new NatFlow is not latched within that time, the NatFlow times out and is deleted. The old reestablished NatFlow remains in place to service existing traffic.

ACLI Instructions and Examples

Relatch support does not require user configuration.