February 1997

Ascend

GRF IP SWITCH

Architecture Guide

ASCEND

TABLE OF CONTENTS

PRODUCT OVERVIEW

GRF: ARCHITECTURE FOR HIGH-PERFORMANCE IP ROUTING

The GRF™ is a high-performance, standards-based IP switch designed for carriers, ISPs and on-line service providers that offer network access or backbone services. With unsurpassed price/performance and price/port-density, the GRF is the best solution for those demanding internetworking environments.

The unique GRF switching architecture is specially designed to handle network growth while offering consistent, high performance. By combining a switch that provides an aggregate bandwidth of up to 16 Gb/s with multi-ported intelligent IP Forwarding Media Cards, the GRF maintains predictable high performance, regardless of the dynamics of the network.

Ascend offers two GRF products:

- The GRF 400 has a compact design that supports up to four media cards in a single, 5.25-inch (13.335 cm) chassis.
- The GRF 1600 supports up to 16 media cards in a single, 21-inch chassis.
- Both GRF systems provide simultaneous support for industry-standard media types, such as HSSI, 10/100Base-T, ATM OC-3c (SDH/STM-1), IP/SONET OC-3C (IP/SDH/STM-1) using PPP and Frame Relay framing, ATM OC-12c (SDH/STM-4C) and FDDI/CDDI, allowing the GRF to be integrated easily into existing networks and enabling the development of new, high-performance networks.

GRF SYSTEM OVERVIEW - THE FIRST PRACTICAL IP SWITCH

Ascend established the following essential design objectives when creating the GRF IP Switch:

Design Objectives	How GRF Achieves Objectives
Compatibility with existing network infrastructures, including interoperability with conventional routers and LAN switches	Support for all popular media, signalling protocols, and both routing and switching
Full compliance with industry standards to eliminate any need for proprietary gateways or special client software	Support all major RFCs pertaining to Internet routing
IP next-hop address lookup fast enough to take advantage of the switching fabric's low-latency and high-throughput	Route lookup is performed by custom hardware (Quick Branch Routing Technology, or QBRT); Able to do a full route lookup in 150,000 route table in under 2.5 microseconds
Sustainable throughput that is independent of traffic characteristics (such as flows or cache hits)	Route lookup is so fast that route lookup is performed on every packet, no caching or traffic analysis is required
Wire-speed performance for all external LAN/WAN ports	Custom DMA engines automatically transfer data between buffer memory and the media
Support for a wide range of popular LAN and WAN media	Support for all popular media: FDDI/CDDI, Ethernet, ATM, HSSI, HIPPI and SONET/SDH; Customers have their choice
Support for ATM without an architectural dependence on ATM	No flows, just fast routing; The crosspoint switch is not cell-based; Customers are not locked into one media type
Linear scalability within each IP switch and in a network of IP switches	Each media card adds another route processor; Each card has its own dedicated data path through the switch core; There is no bus that must be shared
An order of magnitude improvement in price/performance	Line speed packet forwarding combined with the highest media density in the industry

The unparalleled performance levels achieved by Ascend's GRF IP switch are derived from the unique combination of a high-throughput switch fabric and a state-of-the-art hardware-assisted route table lookup. The exceptional packet-per-second (pps) performance of the GRF comes from its unique architectural attributes.

The true test of performance is the ability to maintain wire-speed packet forwarding on all LAN/WAN interfaces, independent of traffic characteristics or network topologies. The GRF easily handles a full complement of LAN/WAN interfaces, each operating at its maximum data rate.

Figure 1—GRF System Architecture

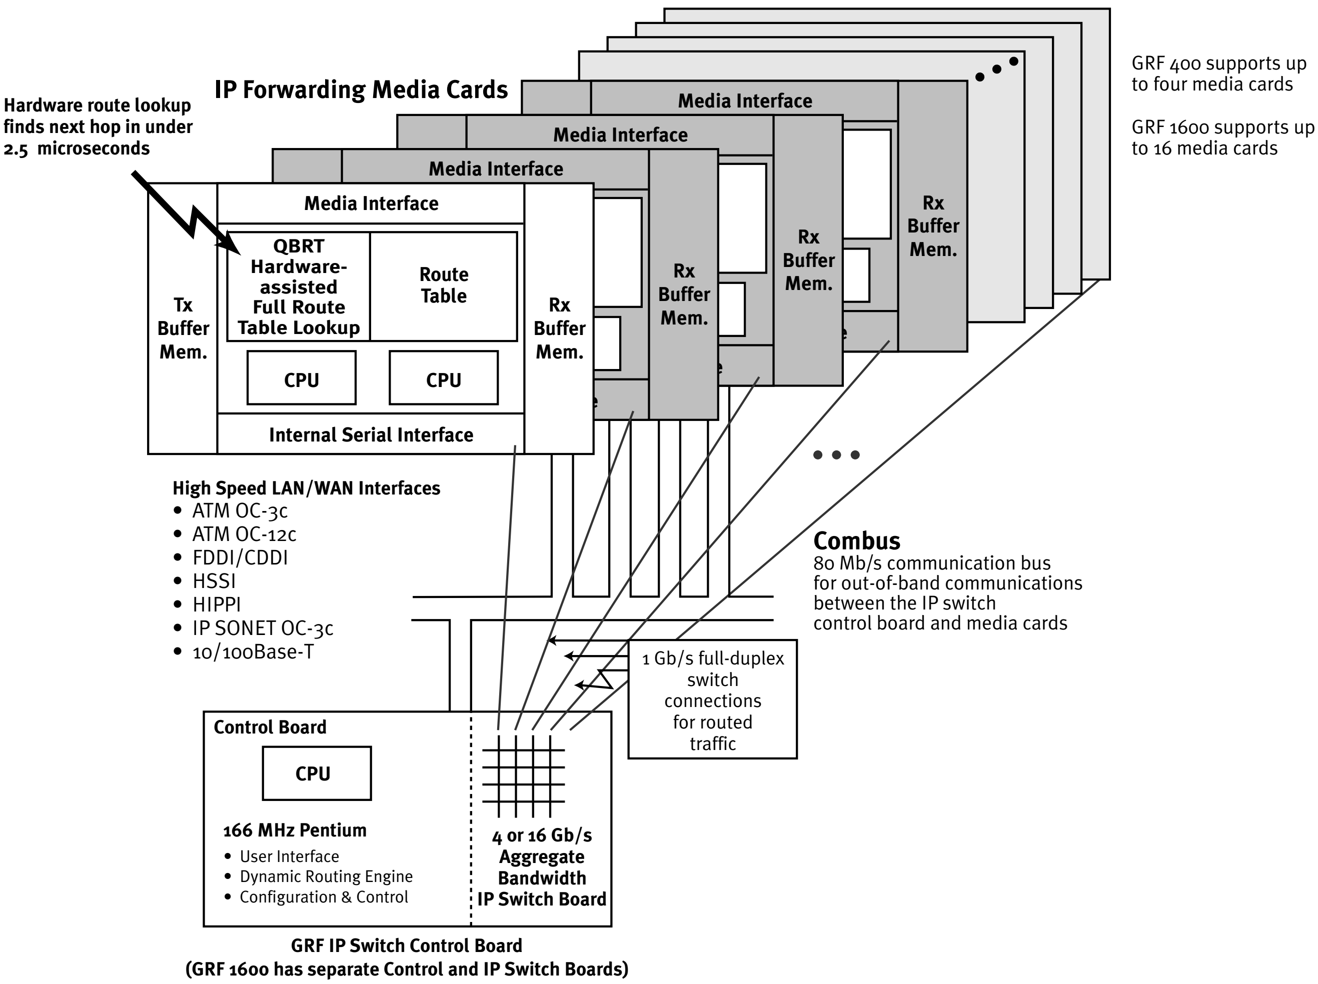

The GRF architecture consists of a unique combination of switching, route management and distributed IP forwarding.

On the GRF 400, the IP Switch Control Board provides dynamic route management processing and houses a 4 Gb/s crosspoint switch. On the GRF 1600, the Control Board provides dynamic route management processing and the IP Switch Board houses a 16 Gb/s crosspoint switch. The high-speed, non-blocking nature of the crosspoint switch fabric permits multiple data paths to operate simultaneously.

The GRF Route Manager runs a full suite of routing protocols using Ascend's embedded operating system. A 166 MHz Pentium processor off-loads the dynamic routing processing from the media cards, allowing them to forward packets without interruption.

The GRF IP Forwarding Media Cards contain an intelligent IP packet forwarding engine including an on-board full route table, CPUs and memory. Several high-speed processors run embedded software to perform dedicated tasks including buffer management and Layer-3 IP forwarding. Each card contains 8 MB of buffering: 4 MB for transmit and 4 MB for receive. The buffers enable speed matching between different media cards across the switch. All high-speed functions have been built using high-speed Field-Programmable Gate Arrays (FPGAs).

The GRF Quick Branch Routing Technology (QBRT) hardware-assisted route table lookup provides consistent, high-performance route table lookup, regardless of traffic patterns. Each card contains a full route table of up to 150,000 routes. The QBRT will find the next hop from the route table in less than 2.5 microseconds. This is 100 times faster than software-driven route table lookups.

GRF IP Switch - Features and Benefits

Table 1 – Features and Benefits

Feature	Feature Description	Benefit
Standards-based Architecture	• The GRF is designed using industry standards • No proprietary features, special systems or software modifications to host system are required to achieve the benefits offered by the GRF	Assures interoperability
Layer-3 Switch Architecture	Provides a low-latency, high-bandwidth path for packet forwarding	Meets expected performance needs resulting from the explosive Internet growth
IP Forwarding Media Cards	• Each media card contains an independent, intelligent IP packet forwarding engine • Layer-3 switching decisions are local to each card	Distributed functions provide superior system performance and packet throughput
High Port Density Delivered in a Small Form Factor	GRF cards have a higher usable port density than the competition	Connectivity per dollar-cubic-foot is maximized
Large Route Table	GRF supports up to 150,000 route table entries in hardware	Accommodates the predicted Internet address space growth with no change
Hardware-Assisted Full Route Table Lookup (QBRT)	Route table lookup is implemented in hardware (no cache)	• The GRF performs Layer-3 decisions at switching speeds, with no degradation in large route-table environments • Hardware-assisted route table lookup is 100 times faster than software implementations
Low System Price	Typical POP configurations are less expensive than the competition	Lower cost of ownership
Scalable Capacity	• The GRF can be expanded to its fullest physical capacity without degrading performance • Its advanced architecture dedicates one gigabit of internal bandwidth to each incremental media card	• Protects prior investment • The GRF predictably and truly accommodates growth to its fullest physical capacity
Very Dense Low-Profile Packaging	The GRF is designed to be very compact The GRF 400 measures 5 1/4" (3u) by 19" The GRF 1600 measure 21" (12U) by 19"	The GRF fits into the restricted space available at a POP
SNMP Compliant	The GRF is accessible via standard SNMP management packages	Fits into standard management strategies
Reliable Hardware Design	• Redundant (N+1), load-balancing power supplies • Hot-swappable power supplies • Redundant, load balancing fans (GRF 1600) • Hot-swappable fan drawer (GRF 1600) • System ignores failed cards without impacting remainder of system • Hot-swappable IP Forwarding Media Cards • Thermal protection	High level of system availability

Feature	Feature Description	Benefit
BGP4 Support	The GRF supports BGP4 including all the latest enhancements: Route Reflections, MEDs Communities, DPAs, Flat Route Dampening, Weighted Route Dampening, Confederations, NextHop-Self and Static Routing as an IGP	• Ensures ability to expand backbone topology • Creates a more-robust network environment • Interconnectivity and redundant paths to multiple NAPs • Coexists with existing router infrastructure
Unmatched High-Performance Media Suite	The GRF provides the most robust suite of high-performance media support in the industry	• Ensures flexible growth options • Gives customers a choice
Two GRF IP Switch models offered	The GRF is available in two models: The GRF 400 supports up to four IP Forwarding Media Cards and the GRF 1600 supports up to 16 IP Forwarding Media Cards	• System size/price matched to customer requirements • Lower cost of ownership

Applications

Internet Protocol (IP) traffic levels have grown exponentially over the past decade for the Internet, and more recently, for intranets. The underlying LANs and WANs have grown accordingly in both size and complexity. Whenever growth reached a level where overall performance degraded, the network was partitioned into smaller segments, which in turn were relinked by new, modified IP internetworking equipment. This dance of scale continues today: a relentless cycle of growth, partitioning and relinking as network architects race to keep pace with the exponential growth in demand.

But in several environments, particularly in large-scale public and private backbones, traffic levels have reached the point where conventional routers are overwhelmed, and proposed routing enhancements are almost instantly rendered ineffective. The core of nearly all performance problems in modern networks lies in the conventional router's basic design. The GRF's 16 Gb/s crosspoint switch fabric and hardware-assist table lookup technology offer a high-performance, highly-scalable solution for the Internet infrastructure.

NAP

The GRF IP switch is ideal for colocation with other network devices in a NAP. It supports a robust suite of routing protocols including RIP1, RIP2, BGP4, OSPF and IS-IS, and up to 30 logical peering sessions without degrading its high-performance, packet forwarding rates. Peering sessions can be implemented over HSSI/DS3 (PPP, Frame Relay), OC-3c (PPP, ATM framing Frame Relay framing), Fast Ethernet, FDDI/CDDI and HIPPI media.

In addition to fitting into current NAP topologies easily, the GRF's leading-edge technology enables new, simpler and more cost-effective network design options to meet the future needs of the expanding Internet.

Figure 2—NAP-New Options with the GRF

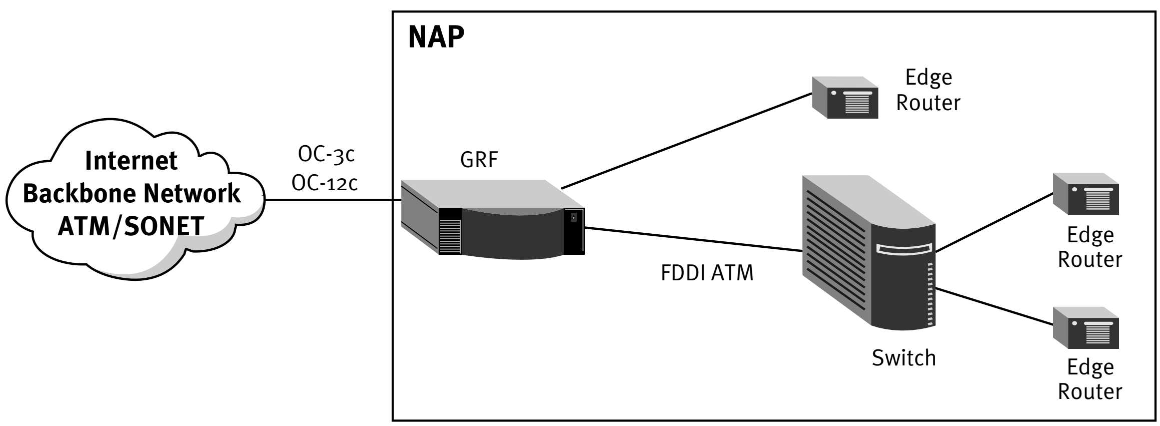

BACKBONE

The GRF IP switch is well suited for providing backbone services over high-speed media because of its high aggregate packet forwarding rates and its support for multiple high-speed media. This facilitates building new high-speed backbones well positioned for the future. The GRF can complement or replace existing Frame Relay switches in the backbone network topologies.

The GRF's high aggregate bandwidth enables interconnection of specialized Autonomous Systems (AS) such as Virtual Private Networks (VPNs), Virtual Access Networks (VANs) and private backbones within a single system. This allows better management of transit traffic and simplifies network design.

Figure 3—GRF in the Backbone

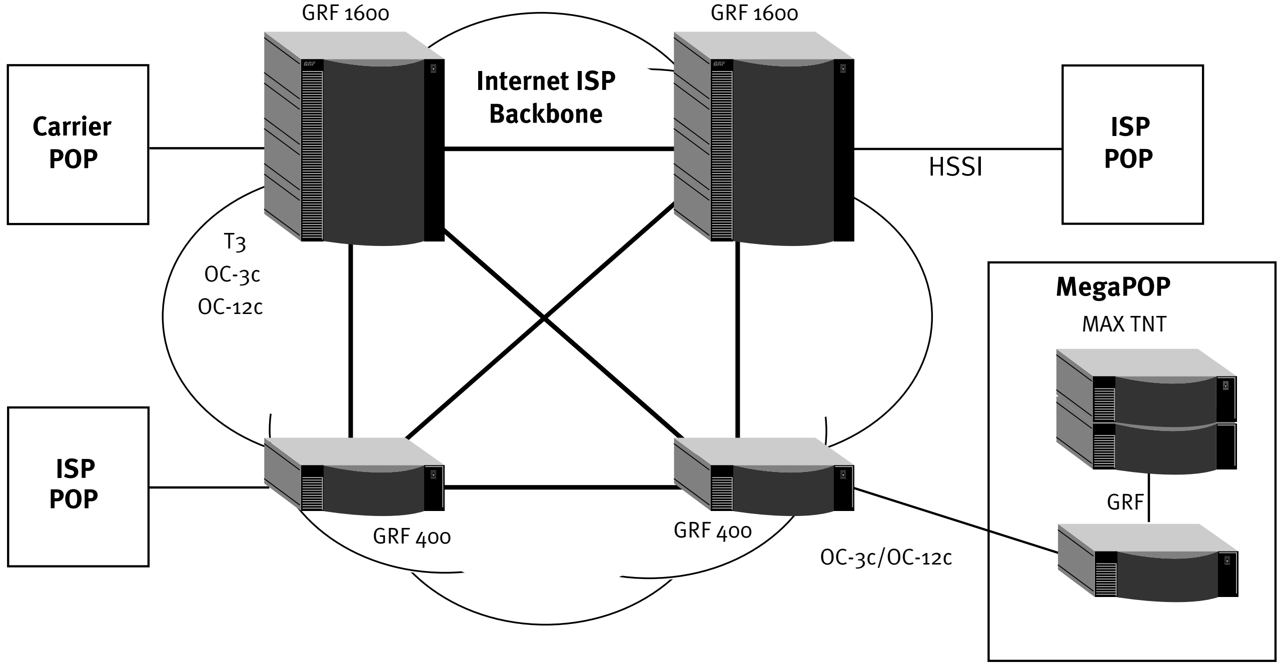

MEGAPOP™

The GRF combined with the MAX TNT™ WAN access switch provides unprecedented packet aggregation services. This powerful solution supports hundreds of thousands of users in under eight feet of rack space.

A single GRF can easily handle the aggregation capacity of a fully configured MAX TNT. This combination is simpler and more cost-effective than using multiple conventional routers. The unmatched density of this solution dramatically reduces rack space requirements.

Figure 4—MegaPOP Solutions Using the GRF and MAX TNT

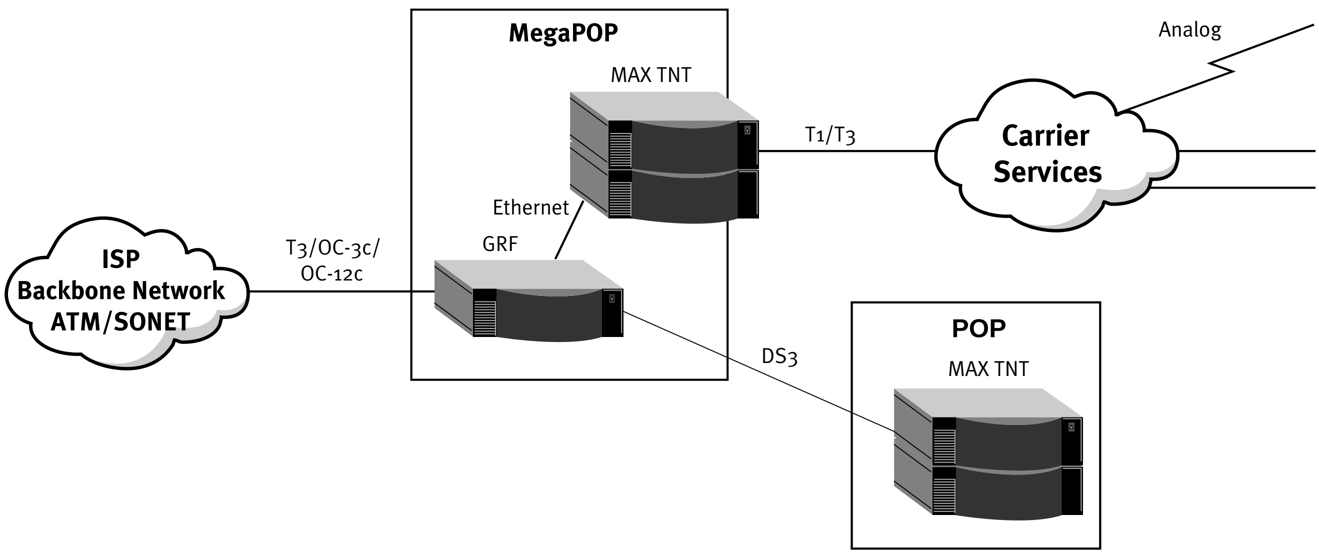

SERVER FARMS

The GRF has the bandwidth to colocate with content servers on high-speed lines while maintaining access to the backbone and providing aggregation services. It easily fits into an ISP's current network topology with support for high-speed media type such as OC-3c, OC-12c and even HIPPI using standard routing protocols.

Figure 5—Server Farm Diagram

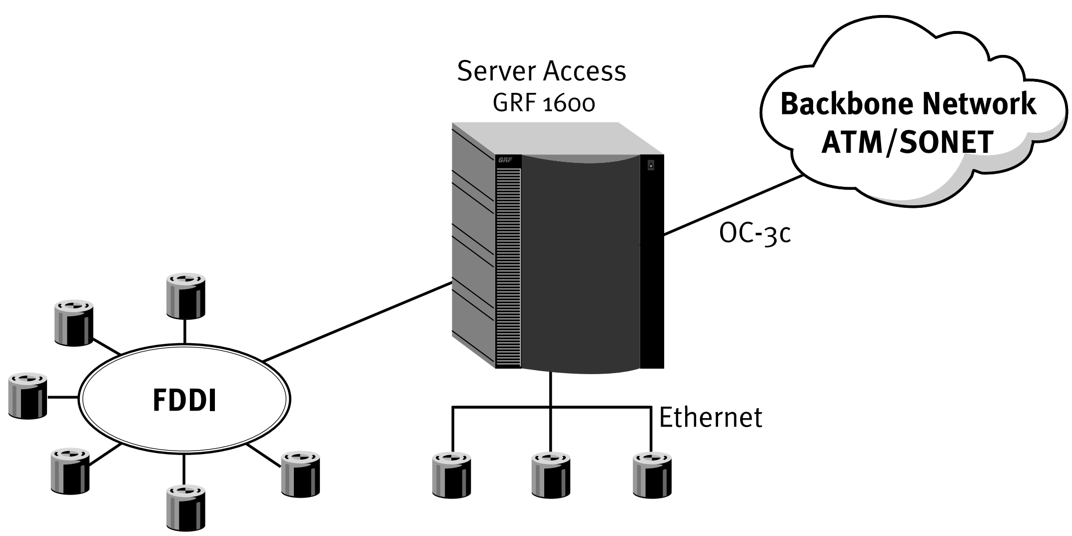

GRF SOFTWARE

Software functions within the GRF IP Switch are distributed between the Route Manager on the IP Switch/Control Board and the individual media cards. The Route Manager runs Ascend's embedded operating system (an enhanced, proprietary version of BSD/OS v2.1) and a number of daemons, including an enhanced version of GateD, that perform various user interface, routing and management functions.

Figure 6—GRF IP Switch Software Architecture Diagram

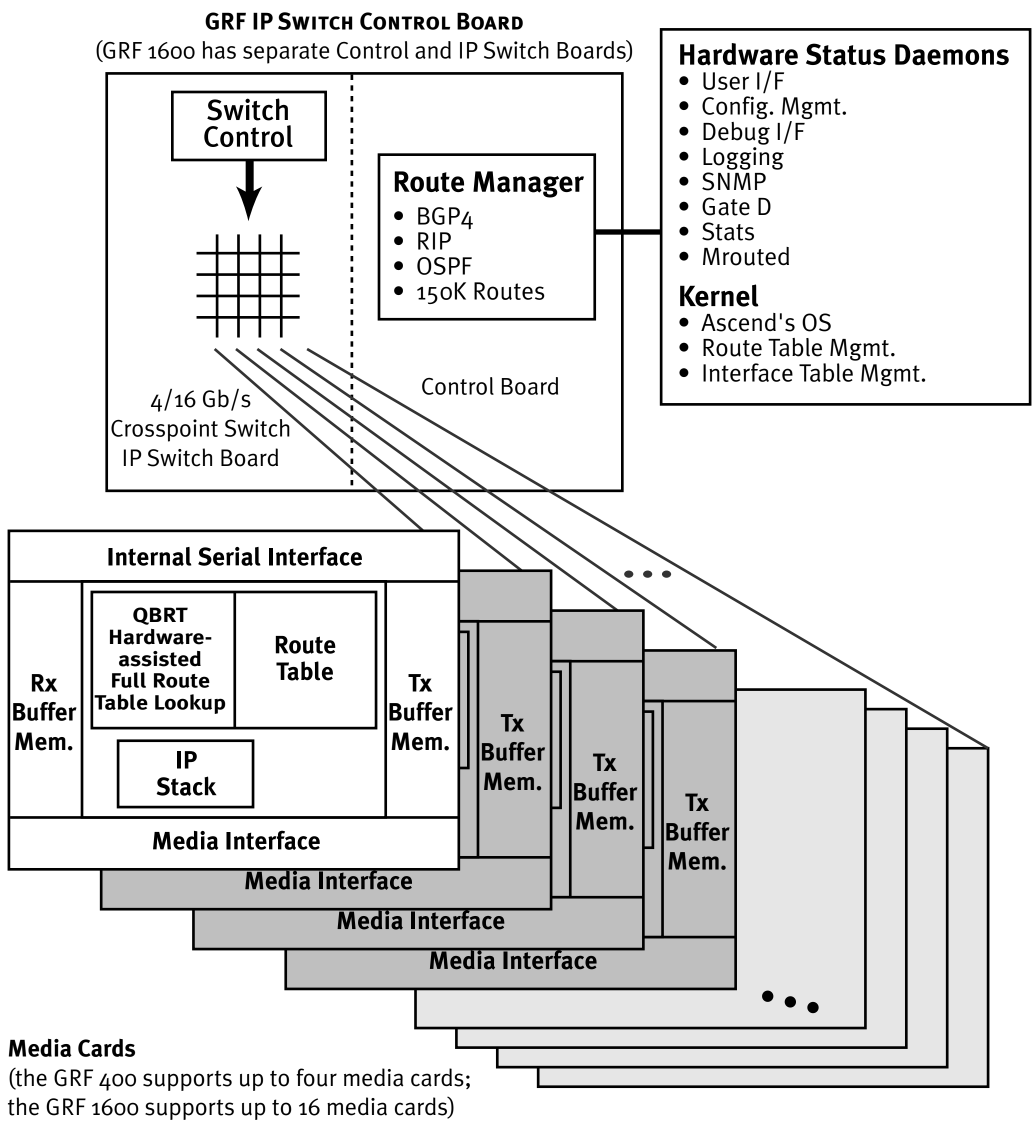

Modifications have been made to the kernel so that interfaces present in the router chassis are logically treated as belonging to the embedded operating system. While the Route Manager updates the system routing tables and performs other administrative functions, the intelligent processors on each media card perform all other IP routing and forwarding functions. This design supports efficient distributed processing of all IP switch forwarding and routing functions.

IP PROTOCOL

The GRF IP switch supports IP datagram routing and forwarding between standard LAN and WAN media. The implementation of IP routing is comprehensive and conforms with the IP Version 4 routing specifications described in the Internet RFCs. The entire TCP/IP protocol suite is supported, including TCP, UDP, ICMP and MTU discovery.

Each media card is its own independent IP forwarding engine and contains a forwarding table of up to 150,000 entries. Address Resolution Protocol (ARP) is also resident on each card. IP packets are buffered in large transmit and receive buffers from which they are streamed across the switch fabric to the destination media card. Any difference in MTU size (large MTU to smaller MTU) is handled by packet fragmentation as specified in the IP standard. Hardware on the destination media is responsible for any media-specific framing.

Data Forwarding

Individual media cards maintain their own forwarding tables, perform lookups and autonomously pass datagrams to other media cards for export, without intervention of the Route Manager or any other control source.

Route Table Implementation

Critical to providing sustained performance in a highly-dynamic environment is Ascend's cacheless route table and hardware-assisted route lookup implementation. Each card carries a complete copy of the route table and can support up to 150,000 entries.

Keeping pace with significant advances in routing, the GRF IP Switch also supports variable-length subnet masking and supernetting/route aggregation.

Subnet Masking/Supernetting

Variable length subnet masking is a classless addressing scheme for interdomain IP packet routing. It is a way to more efficiently manage the current 32-bit IP addressing method. Subnet masks let sites configure the size of their subnets based on site needs, not on the arbitrary Class A, B, or C originally set up in Internet addressing. Class-based addressing uses subnet masks aligned on 8-bit boundaries. The new addressing method allows the network portion of an IP address to be separated from the host portion of the address at any point within the 32-bit address space. This expanded boundary is called the "netmask" and is explicitly provided to the router along with the network address information.

The dynamic routing protocols OSPF, BGP4 and RIPv2 and explicit static routes support classless addressing.

Subnet masking offers a number of benefits in extending the current address space. By eliminating implicit 8-bit netmask assignments, addresses can now be assigned from any unused portion of the entire 32-bit address range rather than from within a specific subset of the space previously called a class. Since it hides multiple subnets under a single network address, this method is called supernetting.

Classless addressing allows the network administrator to further apportion an assigned address block into smaller network (or host) segments based on powers of two (two, four, eight, or 16 networks, for example). Knowledge of the apportioned segments need not be communicated to exterior peers. They need only a single pointer to the entire address block. Subnet masking better utilizes precious IP address space and results in smaller routing tables.

IP Multicast

The GRF supports IP multicast routing per RFC 1112 and components of RFCs 1301 and 1469. The implementation includes the IP multicast 3.5 kernel modifications, DVMRP dynamic route support and mrouted, Version 3, and is supported over CDDI, FDDI, and Ethernet media.

Data that arrives at one GRF interface is duplicated and forwarded to multiple destination interfaces. The multicast packet's destination address is a Class D address. A lookup of the multicast route table returns a list of Virtual Interfaces (VIFs) to which the packet is to be sent.

Host extensions for IP multicasting as described in RFC 1112 are also provided. The Route Manager acts as a host and uses the Internet Group Management Protocol (IGMP), Version 2, to add and delete its membership in multicast groups. Accordingly, the Route Manager "joins" the appropriate routing protocol host groups for OSPF and RIP.

The media cards support multicast as follows:

FDDI/CDDI media card supports multicast messages between nodes; routes and forwards multicast datagrams.

ATM media card supports multicast messages between nodes; multicast datagrams are multicast to a list of IP destinations as configured in the command-line interface (CLI).

10/100Base-T supports multicast messages between nodes, routes and forwards multicast datagrams.

HIPPI media card supports multicast messages between nodes.

HSSI media card supports multicast messages between nodes.

IP/SONET media card supports multicast messages between nodes.

MBONE MBONE datagrams are supported and routed.

Filtering

IP filtering supports specific permit or deny permissions for each instance of a filter. The filters can be configured on a per-interface basis, to look at every inbound packet, every outbound packet, or just the inbound packets destined for the router itself (e.g. dynamic routing packets). This third type of filtering is designated as "Into-me" filtering. Criteria within each filter may include any combination of the following:

- Protocol (ICMP, TCP, UDP)
- Source address
- Destination address
- Protocol port number (single number, or range, or ranges) for TCP and UDP
- Established TCP connections
- Packet header logging

For packet header logging, the filtering system provides the ability to act as a "sniffer," capturing and forwarding portions of packets on the receive side of a given interface. The configuration options available on a statistics sniffer are:

- Filter (which may be a "match all" filter to look at every packet)
- Rate of forward (one packet every n)
- Target machine IP address and UDP port number
- Amount of each packet to capture
- Select to immediately send the UDP packet containing the sniffed header to the target machine or batch a number of headers in the UDP packet before sending

Filtering is performed autonomously by each media card; it is not a function of a central, shared resource.

Precedence Handling

Precedence handling prioritizes delivery of dynamic routing update packets, even when the transmitting media card is congested. Conventional routers designed using shared resources can become overloaded. When overloaded, they drop packets, including dynamic routing update packets. When these update packets are dropped, the next router thinks the transmitting router is down and changes its route table to route around the assumed "network outage." Such events contribute to network instability.

To ensure that dynamic routing update packets are not dropped, the GRF uses precedence features to avoid this instability:

- The GRF dynamic routing agent sets a TOS value in the IP headers of the dynamic routing update packets it generates, which communicates to the media card a high-priority status for the packet.
- The media card maintains a user-configurable threshold of transmit buffers that always remain available for high-priority traffic, ensuring that dynamic routing update packets are forwarded during congested conditions.

ROUTING PROTOCOLS

The GRF IP Switch supports static and dynamic routing. In static routing environments the GRF supports standard IP packet routing with full IP packet forwarding, header validation and processing, ICMP processing and default routing with static routing. The network administrator configures the IP address tables manually.

Dynamic IP routing information is exchanged with neighboring (peer) routers through the dynamic routing agent, an enhanced version of GateD. The GRF supports common dynamic routing protocols including:

Routing Protocols Supported

- RIP v1/v2, (also EGP, BGP 3/4)
- OSPF/OSPF Multicast
- Integrated IS-IS
- BGP4 Mods:
    - Route Reflection
    - Communities
    - MEDs
    - Confederations
    - DPAs
    - Flat Route Dampening
    - Weighted Route Dampening
    - NextHop-Self
    - Static Routing as an IGP
- Multicast DVMRP (see page 10)

Ascend's implementation of dynamic routing protocols is standard and interoperates with other network devices employing the same protocols. These protocols support standard distance/vector and link/state algorithms, as well as interior and exterior route table exchanges.

OSPF, BGP4 and RIP v2 support methods of exchanging network and netmask information between routers. These protocols support supernetting and route aggregation. Although the GRF supports CIDR, older protocols such as RIP v1 and EGP do not.

Dynamic Routing Protocol Overview

Dynamic routing protocols determine the "best" route to each destination and distribute routing information among the systems on a network. Dynamic routing protocols are divided into two general groups: interior protocols and exterior protocols. The GRF supports both types.

Interior Routing Protocols

Interior protocols are used to exchange reachability information within an Autonomous System (AS), which is a collection of interconnected routers managed by one administrator. They are referred to as Interior Gateway Protocols (IGPs). The GRF supports RIP, OSPF and IS-IS.

RIP

Routing Information Protocol (RIP), Version 1 and Version 2, is the most-commonly used interior protocol. RIP selects the route with the lowest number of router hops between the current router and the destination as the "best route."

The longest path accepted by RIP is 15 hops. If the route is longer than 15 hops, a destination is considered unreachable and the GRF, adhering to RIP protocol, will discard the route. RIP assumes the best route is the one that uses the fewest hops. RIP does not make use of information such as line speeds or congestion conditions and therefore is usually used in simple networks.

The RIP Version 1 protocol is described in RFC 1058 and the RIP Version 2 protocol is described in RFC 1388.

OSPF

Open Shortest Path First (OSPF) is a link-state protocol. The state and condition of links between routes are used to determine the "best route." OSPF is better suited than RIP for complex networks using many routers. OSPF supports equal "cost" multipath routing. As long as one route is as "good" (in terms of link state) as another, packets can make use of both. This feature can sometimes alleviate congestion.

OSPF is described in RFC 1583, the MIB is defined in RFC 1253. Other related documents are RFC 1245, RFC 1246 and RFC 1370.

Integrated IS-IS

Intermediate System to Intermediate System (IS-IS) is also a link state interior gateway protocol and was originally developed for routing International Organization for Standardization/Connectionless Network Protocols (ISO/CLNP).

In IS-IS, the network is partitioned into "routing domains." IS-IS intradomain routing is configured and organized hierarchically so that a large domain may be administratively divided into smaller areas. It uses Level 1 Intermediate Systems within areas and Level 2 Intermediate Systems between areas. The Route Manager determines the boundaries of routing domains by setting some links to be exterior links. If a link is marked as "exterior," no IS-IS routing messages are sent on that link.

Routing between administrative domains is handled using BGP.

Exterior Routing Protocols

Exterior protocols are used to exchange routing information between ASes. This enables packets headed for destinations only reachable by "foreign" networks to be passed to the proper router in the "foreign" network. Exterior protocols are only required when an AS must exchange routing information with another AS. Routers within an AS run an interior routing protocol like RIP or OSPF. Only those gateways that connect one AS to another AS need to run an exterior routing protocol. The GRF supports today's predominant exterior protocol, BGP.

BGP

Border Gateway Protocol (BGP) is today's exterior protocol of choice. BGP exchanges reachability information between Autonomous Systems. BGP uses path attributes to provide more information about each route as an aid in selecting the best route. BGP supports nonhierarchical topologies and can be used to implement complex network structures, including autonomous systems.

BGP Version 1 is described in RFC 1105, Version 2 in RFC 1163, Version 3 in RFC 1267, and Version 4 in RFC 1771. The Version 3 MIB is described in RFC 1269. The three documents, RFC 1164, RFC 1268 and RFC 1772, describe the application of Versions 2, 3 and 4 in the Internet.

The recent explosion of Internet users, concomitant IP address assignments and associated network topology diversity has forced modifications to the original BGP. Most of the modifications were devised to cope with one or another aspect of topological and associated administrative complexity.

Ascend uses an enhanced version of BGP which supports the following features:

Route Reflection

Generally, all border routers in a single AS need to be internal peers with each other. Often all routers within an AS need to be internal peers of all border routers in the AS. In large networks, this "full mesh" connectivity can become costly and difficult to manage.

To help address this scaling problem, BGP supports route reflection for internal peer groups (with BGP Version 4 only). When using route reflection, some computers are allowed to "re-advertise" routes to other internal peers, called "route reflectors." A typical use of route reflection might involve a "core" backbone of fully meshed routers, some of which act as route reflectors for routers that are not part of the core group. All the routers in the fully meshed group peer directly with all other routers in the group.

Refer to the route reflection specification document (RFC 1966) for further details.

Communities

The Communities attribute allows the administrator of a Routing Domain to tag groups of routes with a community tag. The tag consists of two octets for the AS and two octets for the Community ID. The Community attribute is passed from routing domain to routing domain, maintaining the logical grouping of these routes. A set of routes may have more than one community tag in its Community attribute. Communities of routers facilitate the management of routes.

Multi Exit Discriminators (MEDs)

Multi Exit Discriminators (MEDs) allow the administrator of a Routing Domain to choose among various paths into a neighboring AS. This attribute is used only for decision making in choosing the best route to the AS. If all other factors for a path to a given AS are "equal" the path with the lower MED metric value takes preference over other paths. It acts as a preferential tie-breaker. The value of the MED attribute is a four-octet unsigned integer.

This attribute is not propagated to other neighboring ASes. However, this attribute may be propagated to other BGP speakers within the same AS.

Confederations

A confederation is a group of routing domains. Confederations accommodate larger networks by aggregating route information. Without Confederations, the complete list of ASes along the path to a destination must be reported in the routing update message. With Confederations, if several ASes in the path are all in the same Confederation, only the Confederation need be listed.

DPA

The Destination Preference Attribute (DPA) is a feature that can be used by an autonomous system to specify a global preference in its routing announcement via BGP. This allows the upstream BGP router to use the preference to favor a certain path for return traffic.

Flat Route Dampening

If new routes are shared too fast, path instability can result. Routers end up devoting more CPU cycles to route calculations than to packet forwarding, performance suffers, networks get congested, links drop and networks go down. A cascade of instability across ASes can result in loss of network access to millions of users. Flat Route Dampening holds a route for a specified amount of time in the routing database before it is exported to its BGP peers.

Weighted Route Dampening

Network reachability oscillations can develop with frequencies related to interactions between router architecture and route calculation algorithms. Weighted Route Dampening treats routes that are going up and down at some specific rate as unreachable. It differs from Flat Route Dampening by suppressing routes in a manner that adapts to the frequency and duration that a particular route appears to be flapping. The more a route flaps during a period of time, the longer it will be suppressed, dampening the instability and freeing up routing CPU cycles.

NextHop-Self

NextHop-Self routing changes the value with the NextHop field of the BGP header to be the IP of the local interface. When using this feature, all BGP peering neighbors believe that the local machine is the next hop to reach all of the routes that are advertised.

NextHop-Self is useful in the following circumstance:

If a router (A) is BGP peering with another router (B) in a neighboring AS via an NBMA cloud, and a route is advertised from B to A that originates from a different router (C) in the same AS, router B has access via a circuit in the NBMA cloud, and A does not; the route will be unreachable from A (routing will fail). Thus, if one sets the NextHop-Self attribute on B, A will believe that the route to C will be via B and routing will succeed.

Static Routing as an IGP

Group-type “routing” is supported for internal BGP peering sessions. Group-type routing now uses an IGP (RIP, OSPF) or static routes to resolve the next hop.

LINK-LEVEL FUNCTIONS

Frame Relay

The GRF supports both switching and routing over Frame Relay networks on the HSSI and IP/SONET OC-3c IP Forwarding Media Cards. On the same link, some virtual circuits can be configured for switching and others for routing. Only Ascend offers Frame Relay at OC-3c speeds.

Routing

The GRF encapsulates IP packets in Frame Relay in accordance with RFC 1490. Inverse ARP and multiple logical interfaces are supported.

Switching

The GRF can operate as a Frame Relay switch, but does not merely switch Frame Relay packets with the speed of a conventional router. It switches at full-speed, like a true Frame Relay switch.

Signaling

If only routing is being done (no switching), then the GRF can use UNI signaling. If any circuits on a link are configured for switching, then the GRF uses NNI signaling. The GRF supports both Annex A and Annex D or can be configured to run without signaling.

SYSTEM MANAGEMENT

SNMP

The GRF currently supports the Simple Network Management Protocol (SNMP) Version 1, which provides a mechanism for remote query or setting operational parameters for the device. Media cards collect network statistics, such as the number of octets coming into and leaving the card, the number of octets discarded, the Maximum Transmission Unit (MTU), and the status of the card. These statistics are reported to network management packages such as HP OpenView or SunNet Manager via the Router Manager on the IP Switch Control Board. Network administrators can also make SNMP connections to the GRF across any active IP interface (media card). In addition to collecting statistics and other management information, the SNMP agent is also capable of issuing traps.

The GRF supports MIB II (RFC 1223) and the FDDI MIB (RFC 1285), SONET MIB (RFC 1595), and HIPPI MIB. In addition, the GRF supports a number of GRF Enterprise MIBs.

MIBs

SNMP management stations can display and alter “objects” in the managed systems. Objects are defined formally in a MIB: a document in a standard language that is both human and machine readable. One MIB, commonly called MIB-II, describes common objects of any system that supports the TCP/IP protocol suite. Managed objects comprise many aspects of the system, including counters, statistics, configuration parameters and tables of items related to the system’s components.

Internet-Standard Management Model

Agent software on the managed system receives read or write requests from the manager software on the management stations. The agent then either carries out the request and returns the result, or responds with an error message explaining why the request could not be carried out. The agent may also send information (known as a trap) to the manager without a request when a significant event or serious problem occurs.

MAINTENANCE

Command-Line Configuration Tools

The GRF IP Switch command-line interface configuration tool has the same look and feel as Ascend's MAX TNT.

RADIUS

The GRF supports the client side of Remote Authentication Dial In User Service (RADIUS) for administrative authentication. Once configured, the client GRF sends authentication requests to a RADIUS server and handles access to the GRF based on the server's response.

Secure ID

The Secure ID client gives the GRF the ability to authenticate administrators connecting to the GRF when used in conjunction with Security Dynamic's ACE/Server.

Self-Tests

The GRF supports a robust set of self-test diagnostics that aid in monitoring the status of the equipment:

1. IP Switch/Control Board power-on self-test: At power-on, the IP Switch/Control Board tests itself and program memory. The processor then performs a series of checks of on-board components.
2. IP Forwarding Media Card self-test: While the IP Switch/Control Board checks are ongoing, media cards are held in reset. When the control board tests are complete, the media cards perform the following tests:
   - Memory tests
   - Loopback and diagnostic tests of their media connections
   - Loopback and diagnostic tests of the serial lines
   - Verifying their communications bus interface via loopback

At this point the GRF is ready to operate and the appropriate LED is illuminated. System boot is automatically performed at system power-on. Reset can be performed manually by pressing the reset button the IP Switch/Control Board. Under software direction, the Route Manager on the IP Switch/Control Board can reset individual Media Cards.

IP FORWARDING MEDIA CARD SOFTWARE

GRF IP Forwarding Media Card software features are detailed on pages 23 through 29.

GRF Hardware

CHASSIS

The GRF IP Switch is ideal for large-scale public and private backbone applications. The GRF 400 is packaged in a 5.25" high, rack-mountable chassis and is configurable with four slots that accept plug-in media cards for external LAN/WAN interfaces. The GRF 1600 is 21" high and supports up to 16 media cards.

Figure 7—The GRF 400 Chassis

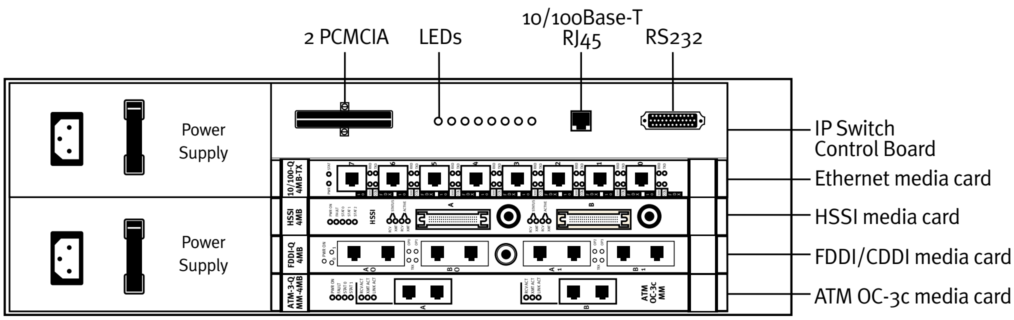

Figure 8—The GRF 1600 Chassis

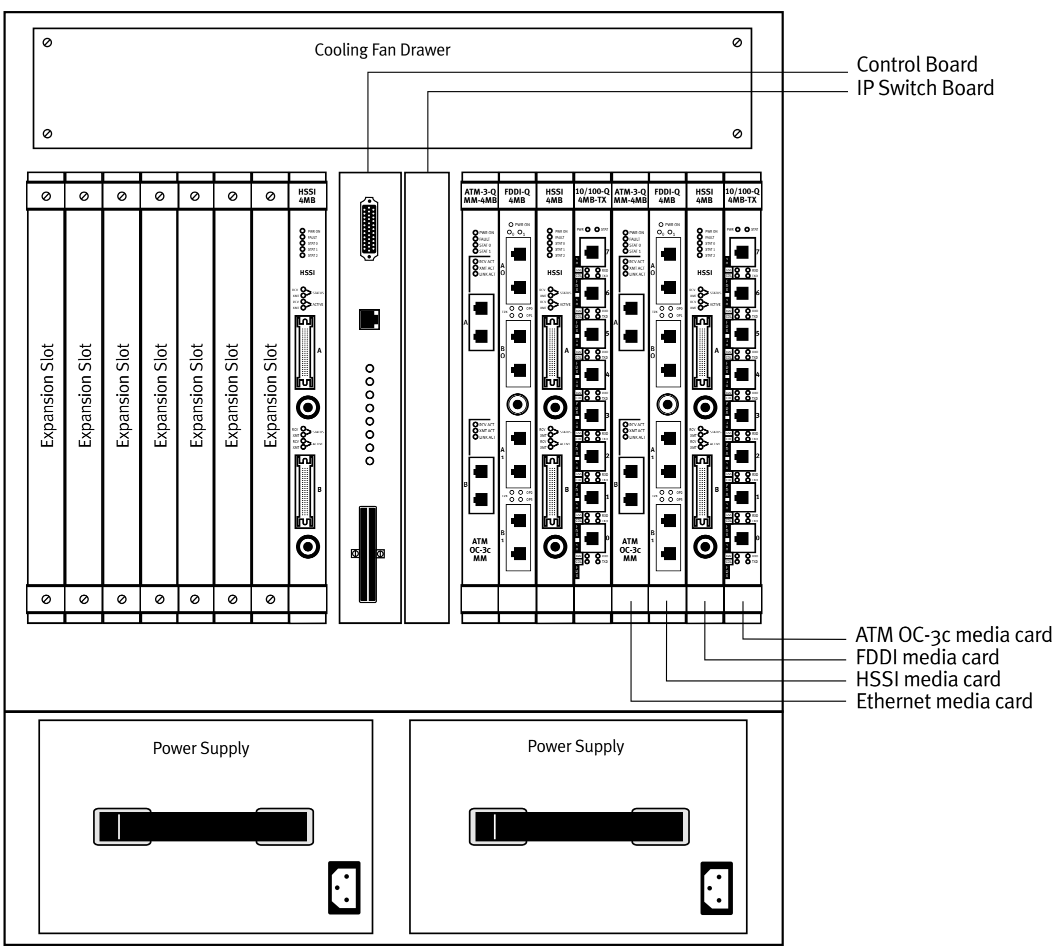

The GRF can be configured with single or dual power supplies. With a second power supply, the GRF distributes the power load across both power supplies and if a power supply fails, it can be swapped out while the system remains on-line.

Passive High-Speed Backplane

A passive high-speed backplane serves as a low cost interconnection between the media cards and the IP Switch/Control Board.

IP SWITCH CONTROL BOARD

The GRF IP Switch Control Board on the GRF 400, and the combination of the Control Board and the IP Switch Board on the GRF 1600, contains an Intel Pentium 166 MHz processor, which runs the Ascend embedded operating system and houses the crosspoint switch. The high-speed, non-blocking nature of the crosspoint switch fabric permits multiple data paths to operate simultaneously. The Route Manager runs a full suite of routing protocols using Ascend's embedded operating system.

Switch Fabric

At the heart of the GRF IP Switch is a serial crosspoint switch fabric. The GRF 400 features a 4 Gb/s switch, and the GRF 1600 features a 16 Gb/s switch. The switch fabric is media independent; it is not an ATM, FDDI or Ethernet switch. This switch fabric is located on the IP Switch Control Board. The high-speed, non-blocking nature of the switch fabric permits up to four multiple data paths to operate simultaneously, each with a non-blocking, bidirectional capacity of 1 Gb/s. Path setup time is less than 250 ns. Since the switch is non-blocking, media cards do not share data path resources and an idle switch data path is always accessible. Because the connection path between each input switch port and output switch port is point-to-point, no resources are shared and bandwidth is maintained for each path. Therefore, no special data transformations are required to carry the IP packets over the switch.

Crosspoint Switch Features:

- Non-blocking design with 250 ns path setup
- Independent 1 Gb/s data paths with hardware-assisted parallel route table lookups ensure that:
  - 4 Gb/s aggregate forwarding rate is achievable with GRF 400
  - 16 Gb/s aggregate forwarding rate is achievable with GRF 1600
  - linear scaling as new cards are added
  - wire-speed performance on all LAN/WAN interfaces (up to 32 interfaces with four Ethernet boards)
- Low latency

Address and switch control signals dynamically connect paths within the switch, while currently connected paths actively and independently transfer data at the 1 Gb/s rate. The GRF has no requirement for synchronized transfer speeds through media card input and output ports, and supports the transfer rate specific to any connected media.

Switch Control

The Switch Control handles requests from media cards to make a connection through the switch to another switch port. A special bus is used to transmit only these connection requests and acknowledgments. There is no exchange of data or card status via this bus. Switch control logic determines whether or not the target switch port is available. If it is, switch control enables the connection and signals the request-maker that the connection has been made.

Route Manager

Router management takes place on the IP Switch Control Board, based on a 166 MHz Pentium processor. It is responsible for system monitoring, configuration management and the user interface.

Specific functions of the Route Manager:

- Processes all dynamic routing packets.
- Synchronizes the route tables on the media cards.
- Controls the media cards - issues interrupts and resets to individual media cards and downloads executable programs and connections.
- Provides external connections - provides one serial RS232, two PCMCIA and one Ethernet connection.
- Supervises - monitors temperature, fan and power supply units. If temperature levels are violated, it issues an alarm. If the temperature continues to rise to a second threshold, it cuts power to all media cards and puts switch control in reset until the temperature returns to an acceptable level. If a power supply fails, the supply is inhibited, the router manager system is notified of the failure and the fault LED is lit and remains on.

Combus

The GRF communications bus (Combus) is an “out-of-band” data path for configuration, control and monitoring of media cards. The Combus connects the IP Switch/Control Board to the media cards independent of the switch connection to each card. It is not used for routed data between media cards. Route update packets received on any media card are also sent across the Combus to the Route Manager and, therefore, do not have to compete with normal IP traffic. The Combus is FIFO-buffered. The Arbitration Logic is on the IP Switch/Control Board.

IP FORWARDING MEDIA CARDS

External LAN/WAN interfaces are provided by an assortment of plug-in IP Forwarding Media Cards. Each media card has its own dedicated 1 Gb/s path in the crosspoint switch fabric, thereby assuring linear scalability to the switch’s capacity of 4 Gb/s. Up to four media cards can be inserted into the GRF 400, for a total of up to 32 external interfaces (when using the Ethernet media cards). With the GRF 1600, up to 16 media cards can be inserted for a total of up to 128 external interfaces (when using the Ethernet media cards).

Every media card is an IP router attached to the internal switching engine. Each media card contains 8 MB of input/output buffering, a complete next-hop route table, hardware-assisted route table lookup and IP forwarding software.

Within the GRF distributed architecture, each intelligent IP Forwarding Media Card makes independent Layer-3 decisions. The cards contain a large route table supporting up to 150,000 routes and memory resources to deliver the performance and capacity demanded by large network such as the Internet.

The GRF dedicates 1 Gb/s of internal bandwidth to each IP Forwarding Media Card, enabling use of every port on each card, and every slot on the GRF, without performance degradation. The GRF scales linearly with every additional IP Forwarding Media Card.

An IP Forwarding Media Card is easy to install and maintain. From the IP Switch Control Board, the card can be reset or put on hold. And since each card is independent; the GRF will continue to operate if a card problem occurs. A card can be replaced while the system remains on-line and the new card will automatically configure itself.

Common Architecture

All GRF IP Forwarding Media Cards share a common basic architecture: Ascend’s unique hardware-assisted route table lookup, a serial switch interface, large input/output buffers, and on-board header and route processing. The command-line interface allows for dynamic reconfiguration of media cards.

Figure 9—IP Forwarding Media Cards

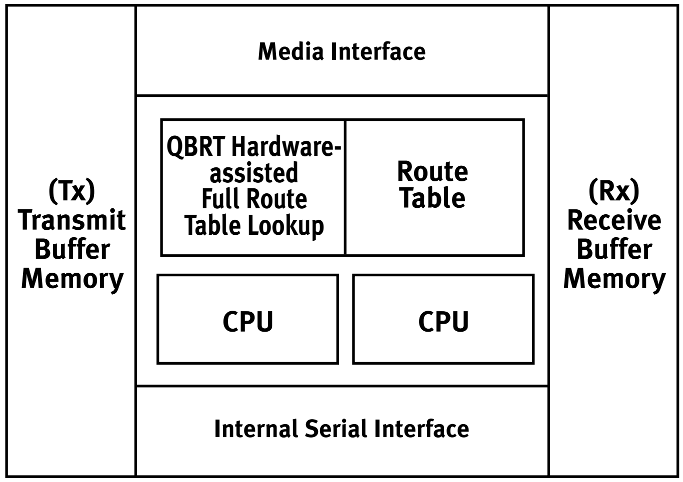

QBRT Hardware-Assisted Route Table Lookup

The ATM OC-3c, 10/100Base-T, FDDI/CDDI, IP/SONET OC-3c (PPP and Frame Relay framing) and ATM OC-12c media cards feature Ascend's Quick Branch Routing Technology (QBRT) hardware-assisted route table lookup. Route lookup times range from 1-2.5 µs with up to 150,000 next-hop routes in the table. The largest tables on the Internet today have only around 50,000 next-hop routes. (The HSSI media card is able to maintain line speed packet forwarding using a software-based full route table lookup.)

The benefit of this architecture is that the entire route table can be stored locally on the media card and searched quickly. In the traditional cached route table method, a small number of routes can be stored and searched locally. However, when a large number of routes is desired, or the kind of traffic one would see on the Internet backbone arise, caching is inadequate. Inevitably, cache misses occur, and route table lookups are performed at a limited, central, shared resource. Thus, performance is severely reduced. QBRT provides consistent, high-performance route table lookup of large tables, regardless of traffic patterns.

Performance is enhanced even further with parallel processing of table lookups occurring on each media card, another technique that helps assure linear scalability. The route manager on the controller board, which also contains the switch fabric, maintains the master route table and distributes updates simultaneously to all installed media cards, even as the cards continue their forwarding function.

Internal Switch Interface

Each media card in the GRF has a full-duplex, 1 Gb/s interface to the internal crosspoint switch. Custom hardware moves the data into and out of buffer memory, independent and in parallel with other media card operations, performs the serial-to-parallel conversions, and generates and verifies CRCs to prevent data corruption across the switch.

Data Buffering

Each media card has a pair of 4 MB RAM buffer memories to support fully buffered transfers. Within the GRF, network buffering serves three purposes:

- Speed decoupling
- Congestion smoothing (independent of TCP flow control)
- Distance decoupling (supports large - 10's of millisecond - depending on WAN media)

Speed decoupling smooths differences in data transfer rates between high-speed and lower-speed media, and between hosts with different performance and workload levels. Input and output media interface (ports) both have buffers to buffer data on both sides of the switch.

Distance decoupling and congestion smoothing are achieved by data buffering during long-distance data transfer. When the GRF is used in wide area networking environments, data transfer can be affected by time delays (over satellite and other long-distance carriers) required for long distance transmission. The GRF distributed architecture is capable of supporting long, global bandwidth delay products.

On-Board Header and Route Processing - Key to GRF's High Performance

Each media card has at least one processor used for IP header and route processing. Each media card processes its own data independently from any other media card, and all route decisions are made local to the card. This is the key to the GRF's performance and scalability. Each time a media card is added to the system, another independent router engine is added. In a conventional router, each new interface only increases the load on the shared resources with concomitant diminution of extant media. In a GRF with 16 media cards, up to 16 route lookups can be done simultaneously. In a conventional router, only one route lookup can be done at a time, regardless of how many media cards are added. When a cache miss occurs, the route lookup is performed by a central resource using software techniques. The added work often interferes with route-efficient calculations.

IP FORWARDING MEDIA CARDS

- 2-port 52 Mbps/45 Mbps DS-3 HSSI (High-Speed Serial Interface)
- 4-port 100 Mbps FDDI (Fiber Distributed Data Interface)
- 4-port 100 Mbps CDDI (Copper Interface)
- 2-port 155 Mbps OC-3c ATM (Asynchronous Transfer Mode UNI 3.0/3.1)
- 4-port 10/100 Mbps Ethernet
- 8-port 10/100 Mbps Ethernet
- 1-port 800 Mbps HIPPI (High-Performance Parallel Interface)
- 1-port 155 Mbps OC-3c SONET (with Frame Relay and PPP Framing)
- 1-port 622 Mbps OC-12c ATM

Table 2 – HSSI Features

HSSI Features	
Attachment Density	Two Interfaces (Bidirectional)
Transfer Rate (PLL will support any rate between 0 and 52 Mbps)	52 Mbps (OC-1), 45 Mbps (DS3), 34 Mbps (E3), 8 Mbps (E2)
Processor	Two 40 MHz Fujitsu Sparc processors (one for transmit side of card and one for receive)
Data Buffering	8 MB - 4 MB Input/4 MB Output
Aggregate Throughput Performance	Up to 70K pps
Route Table Size	150,000
Standard IP Routing	Performs IP datagram forwarding, header validation and processing and ICMP processing
Support for Multiple Framing Protocols	• Frame Relay • Point-to-Point Protocol (PPP) • High-level Data Link Control (HDLC)
Frame Relay	• Performs IP encapsulation per RFC 1490 • Automatically adds PVCs via LMI (optional) • Supports Annex A, Annex D, LMI-Standard and LMI-Disabled • Supports Inverse ARP • Provides 256 logical interfaces per card, independently configurable as point-to-point (single PVC) or multi-access (multiple PVCs) • Supports entire set of 906 unreserved PVCs per physical interface (1,812 per card) • UNI and NNI Signaling
PPP	• Complies with IETF RFCs 1661 and 1662 • Supports Link Quality Monitoring • LCP and IPCP options: - maximum receive unit - quality protocol - magic number - IP address
HDLC	• Full-duplex, synchronous transfer mode • Cisco-compatible
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by any SNMP manager • MIBs include: MIB II (RFC 1223) and the GRF enterprise MIB
Hot-swap Capability	Yes

Table 3 – FDDI/CDDI Features

FDDI/CDDI Features	
CDDI	FDDI board is available for the Copper Interface (CDDI)
Attachment Density	Four Interfaces
Transfer Rate	100 Mbps
Transmission Distance (CDDI)	100 meters (CDDI)
Processors	Dual 40 MHz Fujitsu Sparc processors, QBRT hardware-assisted route table lookup engine
Data Buffering	8 MB - 4 MB Input/4 MB Output
Aggregate Throughput Performance	Up to 136K pps
Route Table Size	150,000
Standard IP Routing	Performs IP datagram forwarding, header validation and processing and ICMP processing
Flexible Attachment Options	The four ports on an individual card can be configured as: • two dual-attach interfaces • one dual-attach interface and two single-attach interfaces • four single-attach interfaces (each interface has unique IP and MAC addresses)
Full FDDI Multicast	• Supports multicast messages between nodes, and routes and forwards multicast datagrams
ARP on FDDI	• ARP table and implementation per RFC 826 and RFC 1390 • Supports Proxy ARP
SMT	Station Management Protocol (Version 7.3), ANSI X3.229
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by any SNMP manager • MIBs include: MIB II (RFC 1223), the FDDI MIB (RFC 1512) and the GRF enterprise MIB
Hot-swap Capability	Yes

Table 4 – ATM OC-3c Features

ATM OC-3c Features	
Attachment Density	Two Interfaces (bi-directional)
Transfer Rate	155 Mbps
Processor	Dual 40 MHz Fujitsu Sparc processors, QBRT hardware-assisted route table lookup engine
Data Buffering	8 MB - 4 MB Input/4 MB Output
Aggregate Throughput Performance	Up to 130K pps
Route Table Size	150,000
Standard IP Routing	Performs IP datagram forwarding, header validation and processing, and ICMP processing
IP over ATM	Encapsulation of IP datagrams and routing on ATM networks using AAL-5 and LLC/SNAP, as recommended in RFC 1483 and RFC 1577
Traffic shaping - UNI Quality of Service (QoS)	Includes peak rate, priority and UNI-specified traffic shaping parameters
UNI 3.0 and 3.1 Signaling	• Configurable for either UNI 3.0 or 3.1 per interface • Enables dynamic linking to/from GRF via SVCs • Implements inverse ARP (InATMARP)
Virtual Circuits	512 Active VCs per connector; 1,024 active VCs per card
Logical Interface	Provides 256 interfaces per card, independently configurable as point-to-point (single PVC) or multi-access (multiple PVCs)
ATM Multicast	Implements IP multicast extensions in RFC 1112
ANSI Standards	Complies with ANSI T1.105-1991 and ANSI T1E1.2/93-020R1
SDH (STM-1)	Complies with Synchronous Digital Hierarchy standards
Cables	Single-mode and multimode options
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by any SNMP manager • MIBs include: MIB II (RFC 1223) and the GRF enterprise MIB
HIPPI over ATM (Tunneling)	• Extends HIPPI long-distance using public carrier services • Tunneling transparent to HIPPI hosts • Transfers all HIPPI protocols, IP, IPI-3, etc., per ANSI X311, Project 1026
Hot-swap capability	Yes

Table 5 – 10/100Base-T Features

10/100Base-T Features	
Attachment Density	Four or Eight 10/100Base-T Autosensing Interfaces
Transfer Rate	100 Mbps or 10 Mbps
Processor	40 MHz Fujitsu Sparc Processors, QBRT hardware-assisted route table lookup engine
Data Buffering	8 MB - 4 MB Input/4 MB Output
Aggregate Throughput Performance	Up to 120K pps
Full Ethernet Multicast	• Supports multicast messages between nodes, and routes and forwards multicast datagrams
Route Table Size	150,000
Standard IP Routing	Performs IP datagram forwarding, header validation and processing and ICMP processing
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by any SNMP manager • MIBs include: MIB II (RFC 1223), the GRF enterprise MIB
Hot-swap capability	Yes

Table 6 – HIPPI Features

HIPPI Features	
Attachment Density	One Interface (Dual-Simplex)
Transfer Rate	800 Mbps
Processor	50 MHz TI C30
Data Buffering	8 MB - 4 MB Input/4 MB Output
Forwarding Performance	Up to 10K pps
Route Table Size	70,000
Standard IP Routing	Performs IP datagram forwarding, header validation and processing, and ICMP processing
Complete ANSI Standard Compliance	HIPPI-SC, HIPPI-PH, HIPPI-FP and HIPPI-LE
HIPPI Switch Emulation	Switches any HIPPI protocol, using raw mode
IP Over HIPPI	Encapsulation of IP datagrams per HIPPI-FP, HIPPI-LE and LLC/SNAP, routing support according to RFC 1374
I-Field-based Dynamic Mode Change	Mode dynamically changes between raw HIPPI switching and IP routing
HIPPI over ATM (tunneling)	• Extends HIPPI long-distance using public carrier services, tunneling transparent to HIPPI hosts • Transfers all HIPPI protocols, IP, IPI-3, etc., per ANSI X311 project 1026
ARP on HIPPI	ARP data maintained manually via user-level command
HIPPI Multicast	• Implements OSPF protocol RFC 1583 and IP extensions in RFC 1112 • Forwards multicast datagrams between HIPPI nodes
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by any SNMP manager • MIBs include: MIB II (RFC 1223), the GRF enterprise MIB, HIPPI-MIB (HIPPI end-point MIB) and HIPPSW-MIB (experimental for a HIPPI switch)
Hot-swap Capability	Yes

Table 7 – IP/SONET OC-3c Features

IP/SONET OC-3c Features	
Standard IP Routing	Performs IP datagram forwarding, header validation and processing and ICMP
Attachment Density	• One interface (Bi-directional) • Supports an APS 1+1 Architecture Switching
Transfer Rate	155 Mbps
Processor	Dual 40 MHz Fujitsu Sparc processors, QBRT hardware-assisted route table lookup engine
Data Buffering	8 MB - 4 MB Input/4 MB Output
Aggregate Throughput Performance	Up to 150K pps
Route Table Size	150,000
Support for Multiple Framing Tools	• Frame Relay • Point-to-Point Protocol (PPP)
Frame Relay	• Performs IP encapsulation per RFC 1490 • Automatically add PVCs via LMI (optional) • Supports Annex A, Annex D, LMI-Standard and LMI-disabled • Provides 256 interfaces per card, independently configurable as point-to-point (single PVC) or multi-access (multiple PVCs) • Supports entire set of 906 unreserved PVCs per physical interface • UNI and NNI
PPP	• Complies with IETF RFCs 1661, 1662 • Supports Link Quality Monitoring • LCP and IPCP options: - maximum receive unit - quality protocol - magic number - IP address
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by the SNMP manager • MIBs include: MIB II (RFC 1223), PPP MIB(RFCs 1471, 1473) and the GRF enterprise MIB
Hot-swap Capability	Yes

Table 8 – ATM OC-12c Features

ATM OC-12c Features	
Attachment Density	One Interface (Bi-directional)
Transfer Rate	622 Mbps
Processor	50 MHz Fujitsu Sparc processor, hardware-assisted route table lookup engine
Data Buffering	8 MB - 4 MB Input/4 MB Output
Aggregate Throughput Performance	Up to 600K pps (300 transmit/300 receive)
Route Table Size	150,000
IP over ATM	Encapsulation of IP datagrams forwarding, header validation and processing, and ICMP processing
UNI 3.0 and 3.1 Signaling	• Configurable for either UNI 3.0 or 3.1 per interface • Supports SVCs • Implements ATM ARP and inverse ARP (InATMARP)
Logical Interface	Provides 256 interfaces per card, independently configurable as point-to-point (single PVC) or multi-access (multiple PVCs)
SNMP Monitoring	• Traffic on the card is monitored by the SNMP agent (RFC 1157), which can in turn be monitored by the SNMP manager • MIBs include: MIB II (RFC1223) and the GRF enterprise MIB
HIPPI over ATM (Tunneling)	• Extends HIPPI long-distance using public carrier services • Tunneling transparent to HIPPI hosts • Transfers all HIPPI protocols, IP, IPI-3, etc., per ANSI X311, Project 1026
Traffic Shaping	Up to 16 traffic profiles per interface
Virtual Circuits	1024 active VCs
ATM Multicast	Implements IP multicast extensions in RFC 1112
ANSI Standards	Complies with ANSI T1.105-1991 and ANSI T1E1.2/93-020R1
SDH (STM-1, STM-4c)	Complies with Synchronous Digital Hierarchy standards
Cables	Single-mode and multimode options
Hot-swap Capability	Yes

RELIABILITY AND AVAILABILITY

The GRF IP Switch, built for continuous carrier-class operation, offers a robust suite of availability features, including hot-swappable media cards and redundant and hot-swappable power supplies.

Hot-swappable Media Cards

Users can swap media cards without powering down the system. After inserting a replacement media card, the media card automatically performs configuration and self-testing. Upon completion of these functions, the GRF downloads current routing information to the media card route table, thus enabling the card to begin to route IP datagrams.

Hot-swappable Power

The GRF can be configured with two identical power supplies. They work on a current-sharing basis: each one provides half the power load. If a supply fails, the other unit immediately assumes the entire load without any interruption to the system. The GRF forwards appropriate status messages to the maintenance interface. The network administrator can replace the defective power supply without powering down the system. After that, the GRF resumes normal current-sharing operation without any interruption in service.

Redundant, Load-balancing, Hot-swappable Fans

The GRF 1600 features redundant, load-balancing cooling fans. Under normal conditions, each fan operates at 50 percent of its capacity. If the temperature goes beyond a predetermined threshold, the fans automatically speed up. If a fan fails, the other fan automatially increases its speed and maintains proper temperature. Additionally, the cooling drawer can be replaced while the system remains on-line.

Thermal Protection

The GRF IP Switch has extra cooling capacity for protection. In the event of a fan failure, the control board detects the failure and sends an alert to the logging facility. The system continues running as long as the temperature remains below the maximum operating temperature. If the temperature rises above the maximum operating temperature, the system sends a warning alert and then smoothly shuts itself down.

GRF System Processes

DATA PACKET PROCESSING

Figure 10—Data Packet Processing

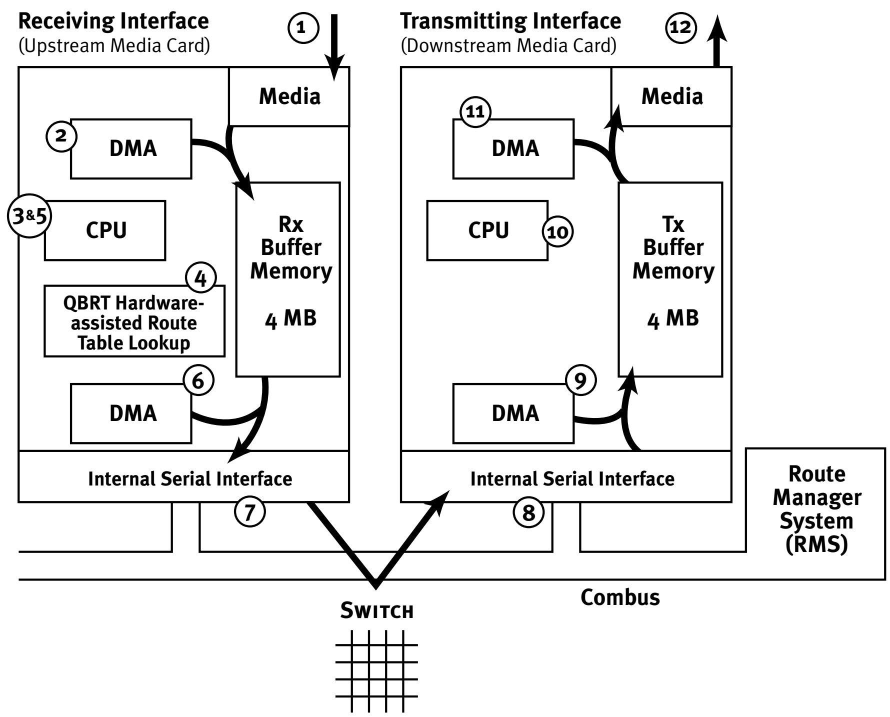

The process below describes how the GRF IP Switch processes dynamic routing products.

Receive Interface (Upstream media card)

1. Data packet arrives on the media.
2. A DMA engine automatically transfers the packet to receive buffer memory.
3. The CPU looks at the packet header and programs the hardware route lookup FPGA with the destination IP address.
   (Steps 4 and 5 are done in parallel.)
4. QBRT hardware finds the next hop in the route table. If no match is found, hardware will automatically return the default route if it exists.
5. The CPU prepends a GRF Internal Encapsulation Format (GRIEF) header to the IP packet. A GRIEF header contains the information needed by the downstream card to place the packet on its specific media. The result of the hardware search, the next-hop, is copied into the next-hop field of the GRIEF header.
6. A DMA engine transfers the packet with GRIEF header to the internal serial interface.
7. The internal serial interface sets up a switch connection to the downstream card. The gallium arsenide parallel-to-serial converter streams the packet across the switch core at 1 Gb/s.

Transmit Interface (Downstream media card)

8. The downstream card converts the serial data stream back into a parallel format.
9. A DMA engine automatically transfers the packet into transmit buffer memory.
10. The CPU examines the next-hop in the GRIEF header and uses that information to build a media header that will deliver the packet to the next-hop router or host. (For example, an ARP lookup.)
11. A DMA engine transfers the packet to the media interface.
12. The packet is transmitted on the media.

Serial processes are decoupled with the use of independent FIFOs maintained by the DMA engines.

DYNAMIC ROUTING PACKET PROCESSING

Figure 11—Dynamic Routing Packet Processing

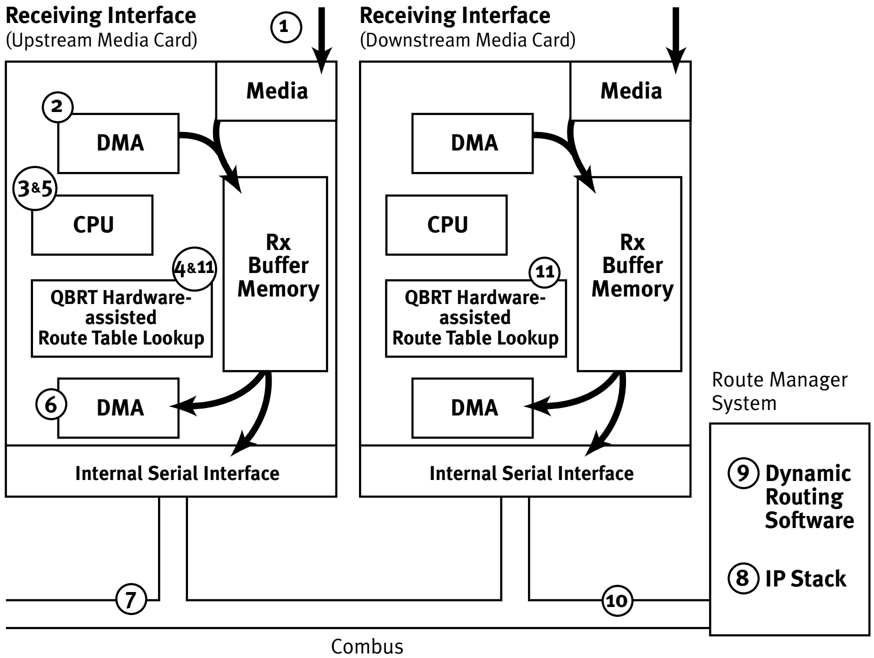

The process below describes how the GRF IP Switch processes dynamic routing packets.

1. Routing packet arrives on the media.
2. A DMA engine automatically transfers the packet to receive buffer memory.
3. The CPU looks at the packet header and programs the hardware route lookup with the destination IP address.
   (Steps 4 and 5 are done in parallel.)
4. QBRT hardware finds the next hop in the route table.
5. The CPU prepends a GRF Internal Encapsulation Format (GRIEF) header to the IP packet. A GRIEF header contains the information needed by the downstream card to forward the packet on the media. The result of the hardware search indicates to the CPU that this packet should be delivered to the IP Switch Control Board's Route Manager. The CPU places a next-hop field in the GRIEF header that points to the IP Switch Control Board's Route Manager.
6. A DMA engine transfers the packet with GRIEF header, to the Combus Interface.
7. The Combus Interface transfers the packet to the IP Switch Control Board's Route Manager across the Combus, the out-of-band (i.e. not across the high-speed serial switch) data path used for internal messages.
8. The IP Switch Control Board's Route Manager receives the packet off the Combus and passes it up the stack to the dynamic routing software.
9. The dynamic routing software processes the packet and determines what adjustments should be made to the global routing table.
10. Route update messages are simultaneously broadcast to each of the media cards over the combus.
11. Each media card receives the route update packet from the IP Switch Control Board's Route Manager and makes the appropriate changes to the hardware table lookup's route table. Note that very few CPU cycles on the media cards were consumed processing the dynamic route update packets.

Configuration Information

Configuring the GRF system requires a base system, up to four (GRF 400) or 16 (GRF 1600) media cards and the appropriate memory, power and cable options. The following is a description of the configuration options available for the GRF IP switch.

BASE SYSTEM CONFIGURATION

All GRF base systems are configured with the listed components. Minimum configuration includes at least one IP Forwarding Media Card and GRF software. Additional components may be added as required.

GRF 400 Base System

GRF 400 IP Switch Base System with AC Power and four-slot chassis (GRF-4-AC) or GRF 400 IP Switch Base System with DC (48 volt) Power and four slot chassis (GRF-4-DC):

- IP Switch Control Board
- 64 MB DRAM (Upgradable to 256 MB)
- One power supply
- Face plate covers for empty media-card slots

GRF 1600 Base System

GRF 1600 IP Switch Base System with AC Power and 16-slot chassis (GRF-16-AC):

- Control Board
- IP Switch Board
- 64 MB DRAM (Upgradable to 256 MB)
- One power supply
- Face plate covers for empty media-card slots

IP FORWARDING MEDIA CARDS

The GRF 400 supports up to four IP Forwarding Media Cards. The GRF 1600 supports up to 16 IP Forwarding Media Cards. The following cards are available for all GRF systems:

• GRF-MC-HSSI	HSSI IP Forwarding Media Card; two ports
• GRF-MC-AOC3M	ATM OC-3c IP Forwarding Media Card; multimode (ATM/Q); two ports
• GRF-MC-AOC3S	ATM/Q OC-3c IP Forwarding Media Card; single mode (ATM/Q); two ports
• GRF-MC-FDDI	FDDI/Q IP Forwarding Media Card (FDDI/Q); four ports
• GRF-MC-CDDI	CDDI/Q IP Forwarding Media Card (CDDI/Q); four ports
• GRF-MC-EN8	10/100Base-T IP Forwarding Media Card (10/100Base-T/Q); eight ports
• GRF-MC-EN4	10/100Base-T IP Forwarding Media Card (10/100Base-T/Q); four ports
• GRF-MC-HIPPI	HIPPI IP Forwarding Media Card; one port
• GRF-MC-IPSM	IP/SONET OC-3c; multimode (IP/SONET/Q); one port
• GRF-MC-IPSS	IP/SONET OC-3c; single mode (IP/SONET/Q); one port
• GRF-MC-AOC12S	ATM OC-12c IP Forwarding Media Card; single mode (ATM/Q); one port

Note: Cards labeled with “/Q” feature the QBRT hardware-assisted route table lookup.

GRF SOFTWARE

The following software is required for all base systems:

• GRF-SR-STD	GRF Software; includes manual and power cord

MEMORY

The GRF base system comes with 64 MB DRAM. In all GRF memory configurations, 32 MB is used for the file system and 18 MB is used for system operation. For example, in the base system there is 64 MB of total memory, 32 MB of available memory and 14 MB of usable memory. (Refer to the following table for detailed memory configuration information.) Up to six additional 32 MB DRAM SIMMs may be added to support larger dynamic routing tables and larger numbers of peers (for a total of 256 MB, 204 MB usable). Additional memory must be purchased in 32 MB pairs (64 MB).

- Two 32 MB DRAM SIMMs (GRF-MEM-64)

The following table provides guidelines for memory configuration. All media cards can hold up 150K route entries. The Control Board, depending on memory configuration, can hold 35,805 to 521,730 route prefixes. Select the amount of memory according to your routing environment. Additional memory may be required for higher average numbers of routes per BGP peer.

Table 9 – Memory Configuration

Customer Profile	Total Memory	Available	Usable	No. of Route Entries on Media Card	No. of Route Prefixes[1] in Dynamic Routing Database	No. of Peer[2] Sessions
Static Routing: (High-Performance Computing Environment)	64 MB	32 MB	14 MB	150K (limited to 35,805 on base system)	35,805	0
Small POP:	128 MB	96 MB	78 MB	150K	199,485	3
Medium POP/ISP Backbone:	192 MB	160 MB	142 MB	150K	362,165	7
Large POP/Exchange Point/ Route Reflection Server:	256 MB	224 MB	204 MB	150K	521,730	10

[1] Total number of route addresses and associated policies received from peers including directly connected addresses.
[2] Assumes BGP peer with 50K route entries.

Figure 12 – System RAM

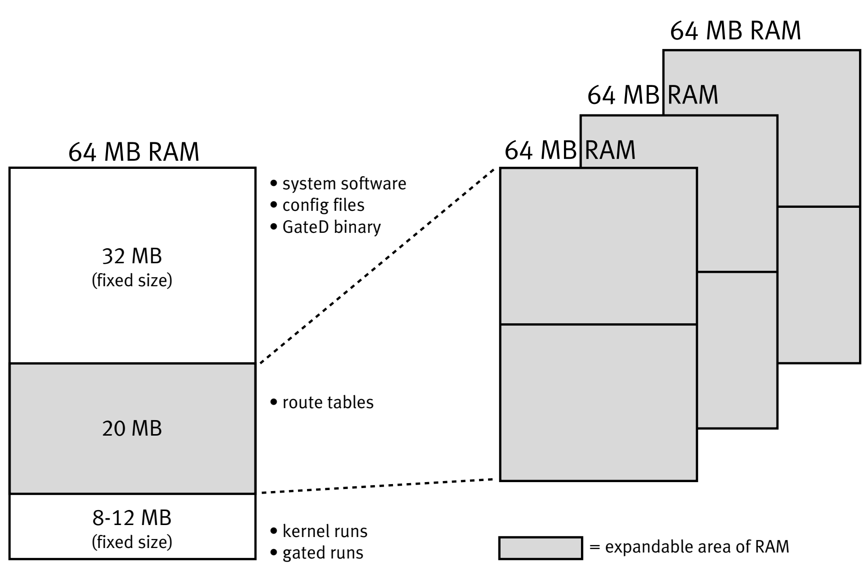

REDUNDANT POWER SUPPLY

To maximize uptime, the GRF supports a redundant power supply. With two power supplies, the GRF automatically distributes the power load to both power supplies. If one of the power supplies fails, it can be swapped out while the system remains on-line. Ascend recommends the second power supply to customers with mission-critical operations.

- Second AC Power Supply for Dual Supply Option on GRF 400 (GRF-AC-PS)
- Second DC (48 Volt) Power Supply for Dual Power Supply Option on GRF 400 (GRF-AC-DC4)
- Second AC Power Supply for Dual Supply Option on GRF 1600 (GRF-AC-AC16)

IP SWITCH CONTROL BOARD – SPARE

A spare IP Switch Control Board (GRF 400) or Control Board and IP Switch Board (GRF 1600) may be ordered (along with the required GRF software). The boards are not hot-swappable.

- GRF 400 IP Switch Control Board – spare (GRF-CB-4)
- GRF 1600 IP Switch Control Board – spare (GRF-AC-CB16)
- GRF 1600 IP Switch Board – spare (GRF-AC-SWB16)

COOLING FAN DRAWER – SPARE

The GRF 1600 supports hot-swappable cooling fans. To support this a spare cooling fan drawer may be ordered:

- GRF 1600 Fan Drawer – spare (GRF-AC-FAN16)

CABLES

All of the GRF IP Forwarding Media Cards support industry-standard cables. Since Ascend has certified the viability of GRF-to-GRF HIPPI connections using 51-meter cables, it offers HIPPI cables directly:

- HIPPI Cable; 5 Meters (GRF-CBL-HP5)
- HIPPI Cable; 25 Meters (GRF-CBL-HP25)
- HIPPI Cable; 51 Meters (GRF-CBL-HP51)

Table 3 – Cable Support

IP Forwarding Media Card	Cable Support
HSSI	Two 25 twisted-pair shielded coax cables
ATM OC-3c Multimode	Two 62.5/125 micron fiber optic cables
ATM OC-3c Single Mode	Two 9/125 micron fiber optic cables
10/100Base-T	Four/Eight STP or UTP CAT5 cable
FDDI	Four multimode cables (62.5/125 micron optical fiber)
CDDI	Four Category 5 twisted pair unshielded copper cables
IP/SONET OC-3c Multimode	One 62.5/125 micron fiber optic cable
IP/SONET OC-3c Single Mode	One 9/125 micron fiber optic cables
HIPPI	Two twisted-pair copper cables
ATM OC-12c Single Mode	One 9/125 micron fiber optic cable

Product Specifications

Table 11 — GRF 400 System Chassis Specifications

GRF 400 System Chassis Specifications	
IP Forwarding Media Cards	1 - 4 per chassis, multiple media types supported simultaneously
Dimensions	5.25 in x 19 in. x 19 in. [13.34 cm. 48.26 x cm. 48.26 x cm.] (with rack mount ears)
Weight	Chassis with one power supply 26.5 lbs. [11.9 Kg] Single power supply 6 lbs. [2.7 Kg] Single media card 2 lbs. [0.9 Kg]
AC Voltage	100 to 240 VAC, 60=50 Hz., 6.0-3.0 Amp.
Power Cord	IEC320 to NEMA 5-15P, international options available
Power Consumption	500 W
Altitude	0 - 10,000 feet (0 to 3,048 meters)
Relative Humidity	10% to 90% non-condensing
Temperature	32 - 104°F [0 - 40°C]

Table 12 — GRF 1600 System Chassis Specifications

GRF 1600 System Chassis Specifications	
IP Forwarding Media Cards	1 - 4 per chassis, multiple media types supported simultaneously
Dimensions	17.5 in x 19 in. x 19 in. [48.26 x cm. 48.26 x cm.]
Weight	Chassis with one power supply 100 lbs. [45 Kg] Single power supply 21.5 lbs. [9.8 Kg] Single media card 2 lbs. [0.9 Kg]
AC Voltage	85 to 264 VAC, 60/50 Hz. 10 Amp. maximum
Wall Receptacle	Standard US 3-prong plug, 115 V grounded receptacle, international options available
Power Consumption	1200 W
Altitude	0 - 10,000 feet (0 to 3,048 meters)
Relative Humidity	10% to 90% non-condensing
Temperature	32 - 104°F [0 - 40°C]

Table 13 — GRF 400 IP Switch Control Board Specifications

GRF 400 Control Board Specifications	
Processor	Intel Pentium 166 MHz
Processor L2 Cache	512 KB Fast SRAM
DRAM	64 MB - 256 MB SECDED memory (8 SIMM sockets for 32 MB SIMMS)
FLASH Memory	85 MB
Switch	16 Gb TriQuint TQ8017
PCMCIA Slots	2 (1 PCMCIA I/II, 1 PCMCIA I/II/III)
Serial Adapters	1
10/100Base-T	1 autosensing
Software	Ascend Embedded OS

Table 14 — GRF 1600 Control Board Specifications

GRF 1600 Control Board Specifications	
Processor	Intel Pentium 166 MHz
Processor L2 Cache	512 KB Fast SRAM
DRAM	64 MB - 256 MB SECDED memory (8 SIMM sockets for 32 MB SIMMS)
FLASH Memory	85 MB
PCMCIA Slots	2 (1 PCMCIA 1/11, 1 PCMCIA 1/11/111)
Serial Adapters	1
10/100Base-T	1 autosensing
Software	Ascend Embedded OS

Table 15 — GRF 1600 IP Switch Board Specifications

GRF 1600 Switch Specifications	
Switch	16 Gb TriQuint TQ8017

Table 16 — Internet Standards

The list below summarizes the Internet standards supported by the GRF.

Supported Standards for GRF Chassis with IP Switch Control Board	
RFC 768	User Datagram Protocol
RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 793	Transmission Control Protocol
RFC 854	Telnet Protocol Specification
RFC 855	Telnet Option Specifications
RFC 857	Telnet Echo Option
RFC 858	Telnet Suppress Go Ahead Option
RFC 860	Telnet Timing Mark Option
RFC 862	Echo Protocol
RFC 863	Discard Protocol
RFC 864	Character Generator Protocol
RFC 867	Daytime Protocol
RFC 868	Time Protocol
RFC 894	Standard for the Transmission of IP Datagrams Over Ethernet Networks
RFC 959	File Transfer Protocol
RFC 1034	Domain Names - Concepts and Facilities
RFC 1035	Domain Names - Implementation and Specification
RFC 1058	Routing Information Protocol
RFC 1073	Telnet Window Size Option
RFC 1079	Telnet Terminal Speed Option
RFC 1112	Host Extensions for IP Multicasting
RFC 1155	Structure and Identification of Management Information for TCP/IP-based Internets
RFC 1157	A Simple Network Management Protocol (SNMP)
RFC 1184	Telnet Linemode Option
RFC 1191	Path MTU Discovery
RFC 1212	Concise MIB Definitions
RFC 1213	Management Information Base for Network Management of TCP/IP-based Internets: MIB-II
RFC 1215	A Convention for Defining Traps for Use with the SNMP
RFC 1227	SNMP MUX Protocol and MIB
RFC 1267	A Border Gateway Protocol 3 (BGP-3)
RFC 1305	Network Time Protocol (v3)
RFC 1350	The TFTP Protocol (Revision 2)
RFC 1397	Default Route Advertisement in BGP2 and BGP3 Versions of the Border Gateway Protocol
RFC 1403	BGP OSPF Interaction
RFC 1471	The Definitions of Managed Objects for the Link Control Protocol of the Point-to-Point Protocol

Supported Standards for GRF Chassis with IP Switch Control Board (continued)

RFC 1473	The Definitions of Managed Objects for the IP Network Control Protocol of the Point-to-Point Protocol
RFC 1512	FDDI Management Information Base
RFC 1583	OSPF v 2
RFC 1643	Definitions of Managed Objects for the Ethernet-like Interface Types
RFC 1722	RIP Version 2 Protocol Applicability Statement
RFC 1723	RIP Version 2 Carrying Additional Information
RFC 1771	A Border Gateway Protocol 4 (BGP-4)
RFC 1901	Introduction to Community-Based SNMP v2
RFC 1902	Structure of Management Information for Version 2 of the Simple Network Management Protocol (SNMP v2)
RFC 1903	Textual Conventions for Version 2 of the Simple Network Management Protocol (SNMP v2)
RFC 1904	Conformance Statements for Version 2 of the Simple Network Management Protocol (SNMP v2)
RFC 1905	Protocol Operations for Version 2 of the Simple Network Management Protocol (SNMP v2)
RFC 1906	Transport Mappings for Version 2 of the Simple Network Management Protocol (SNMP v2)
RFC 1907	Management Information Base for Version 2 of the Simple Network Management Protocol (SNMP v2)

Table 17—HSSI

HSSI Specifications

Max Transmission Unit	Frame relay - 4,352 bytes PPP - 1,496 bytes HDLC - 4,352 bytes
Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 kg)
DC Power Supplies	+5.0 V = 10 amps maximum, total watts = 50
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104°F (0 to 40°C)
Card Connectors	2-row, 50-pin shielded tab connectors
Cables per Card	Two, 25 twisted-pair shielded coax cables
Cable Receptacles	2-row, 50-pin receptacle heads

HSSI Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1293	Inverse Address Resolution Protocol
RFC 1332	The PPP Internet Protocol Control Protocol (IPCP)
RFC 1333	PPP Link Quality Monitoring
RFC 1490	Multiprotocol Interconnect over Frame Relay
RFC 1661	The Point-to-Point Protocol (PPP)
RFC 1662	PPP in HDLC-like Framing

Table 18 – FDDI/CDDI

FDDI/CDDI Specifications

Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 kg)
DC Power Supplies	+5.0 V = 9 amps maximum, total watts = 45
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104˚F (0 to 40˚C)
Cable Support	Multimode cables: 62.5/125 micron optical fiber (FDDI) UTP-5 - twisted pair (CDDI)
Physical Connectors	MIC transceivers RJ-45 (CDDI)
Transceiver Requirements	Must meet ANSI X3.166 standards
Optic Components	Internally-produced 1,300 nm surface-emitting LEDs and photodiodes
Optical Output (avg.) (62.5/125 um)	-19 dBm minimum, -14 dBm maximum
Optical Input Sensitivity	-33.5 dBm minimum, -14 dBm maximum
Optical Wavelength	1,270 nm minimum, 1,330 nm typical, 1,380 nm maximum
Spectral Width	130 nm maximum

FDDI/CDDI Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1390	Transmission of IP and ARP over FDDI Networks
RFC 1512	FDDI Management Information Base

Table 19 – ATM OC-3c

ATM OC-3c Specifications

Cable Requirements	• Single-mode cables: 9/125 micron optical fiber • Multimode cables: 62.5/125 micron optical fiber
Physical Connectors	Duplex SC transceivers
Connector Requirements	Must meet JIS C 5973 standards
LASER Components	Class1, 1310 nm LASER diodes
Optical Input (avg.)	• -32.5 dBm minimum, -8 dBm maximum Single Mode (SM) • -32.5 dBm minimum, -14 dBm maximum Multimode (MM)
Optical Output (avg.)	• -14 dBm minimum, -8 dBm maximum Single Mode, 9/125 um, (SM) • -19 dBm minimum, -14 dBm maximum Multimode, 62.5/125 um, (MM)
Optical Wavelength	• 1,261 nm minimum, 1,310 nm typical, 1,360 nm maximum (SM) • 1,270 nm minimum, 1,310 nm typical, 1,380 nm maximum (MM)
Spectral Width	• 7.7 nm maximum (SM) • 120 nm typical, 200 nm maximum (MM)
Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 kg)
DC Power Supplies	+5.0 V = 9 amps maximum, total watts = 45
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104 °F (0 to 40° C)

Single-Mode (SM)
Multimode (MM)

ATM OC-3c Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1483	Multiprotocol Encapsulation over ATM Adaptation Layer 5
RFC 1577	Classical IP and ARP over ATM
RFC 1626	Default IP MTU for use over ATM AAL-5
RFC 1755	ATM Signaling Support for IP over ATM

Table 20 – 10/100Base-T

10/100Base-T Specifications

Cable Requirements	STP and UTP CAT5
Physical Connectors	RJ-45
Transmission Distance	100 meters
Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 cm)
DC Power Supplies	+5.0 V = 10 amps max, total watts = 45
Hot-Swap Capability	Yes
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104°F (0 to 40°C)

10/100Base-T Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 826	Ethernet Address Resolution Protocol: converting network protocol addresses to 48-bit Ethernet address for transmission on Ethernet hardware
RFC 894	Standard for the transmission of IP datagrams over Ethernet networks
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1643	Definitions of Managed Objects for the Ethernet-like interface types

Table 21 – HIPPI

HIPPI Specifications

Transmission Distance	Point-to-point and LAN: up to 25 meters (82 feet); Ascend certifies viability of GRF-to-GRF connections using 51-meter cables
Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 kg)
DC Power Supplies	+5.0 V = 8.8 amps maximum, total watts = 44
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104˚F (0 to 40˚C)
Cable Support	Twisted pair copper cables: 5, 25 and 51 meter cables available from Ascend
Physical Connectors	• Cable: 2-row 100-pin shielded tab connectors as specified in HIPPI-PH • Card: 2-row 100-pin panel-mount receptacle as specified in HIPPI-PH

HIPPI Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1374	IP and ARP on HIPPI

Table 22 – IP/SONET OC-3c

IP/SONET OC-3c Specifications

Transfer Rate	155 Mbps
Transmission Distance	LAN; WAN distances with use of repeaters
Interfaces Per card	1 full duplex OC-3c SONET/STM-1 SDH (supports an APS 1+1 Architecture Switching)
Cable Requirements	• Single-mode cables: 9/125 micron optical fiber • Multimode cables: 62.5/125 micron optical fiber
Physical Connectors	Duplex SC transceivers
Connector Requirements	Must meet JIS C 5973 standards
LASER Components	Class1, 1310 nm LASER diodes
Optical Output (avg.) (9/125 um)	• -14 dBm minimum, -8 dBm maximum Single Mode (SM) • -19 dBm minimum, -14 dBm maximum Multimode (MM)
Optical Wavelength	• 1,261 nm minimum, 1,310 nm typical, 1,360 nm maximum (SM) • 1,270 nm minimum, 1,310 nm typical, 1,380 nm maximum (MM)
Spectral Width	• 7.7 nm maximum (SM) • 120 nm typical, 200 nm maximum (MM)
Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 kg)
DC Power Supplies	+5.0 V = 9 amps maximum, total watts = 45
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104 °F (0 to 40° C)
Single-Mode (SM) *Multimode (MM)*	

IP/SONET OC-3c Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1293	Inverse Address Resolution Protocol
RFC 1332	The PPP Internet Protocol Control Protocol (IPCP)
RFC 1333	PPP Link Quality Monitoring
RFC 1490	Multiprotocol Interconnect over Frame Relay
RFC 1619	PPP over SONET/SDH
RFC 1661	The Point-to-Point Protocol (PPP)
RFC 1662	PPP in HDLC-like Framing

Table 23 – ATM OC-12c

ATM OC-12c Specifications

Cable Requirements	• Single-mode cables: 9/125 micron optical fiber • Multimode cables: 62.5/125 micron optical fiber
Physical Connectors	Duplex SC transceivers
Connector Requirements	Must meet JIS C 5973 standards
LASER Components	Class1, 1310 nm LASER diodes
Optical Input (avg.)	• -28 dBm minimum, -7.0 dBm maximum Single Mode (SM) • -26 dBm minimum, -14.0 dBm maximum Multimode (MM)
Optical Output (avg.)	• -15 dBm minimum, -8 dBm maximum Single Mode, 9/125 um, (SM) • -19 dBm minimum, -14 dBm maximum Multimode, 62.5/125 um, (MM)
Optical Wavelength	• 1,274 nm minimum, 1,310 nm typical, 1,356 nm maximum (SM) • 1,270 nm minimum, 1,310 nm typical, 1,380 nm maximum (MM)
Spectral Width	• 2.5 nm maximum (SM) • 120 nm typical, 200 nm maximum (MM)
Card Dimensions	9 in. high x 16 in. long (22.9 cm x 40.6 cm)
Card Weight	~ 2 pounds (0.9 kg)
DC Power Supplies	+5.0 V = 9 amps maximum, total watts = 45
Operating Humidity	10% to 90%, non-condensing
Operating Temperature	32 to 104 °F (0 to 40° C)

Single-Mode (SM)
Multimode (MM)

ATM OC-12c Internet Standards

RFC 791	Internet Protocol
RFC 792	Internet Control Message Protocol
RFC 1112	Host extensions for IP multicasting
RFC 1191	Path MTU Discovery
RFC 1483	Multiprotocol Encapsulation over ATM Adaptation Layer 5
RFC 1577	Classical IP and ARP over ATM
RFC 1626	Default IP MTU for use over ATM AAL-5
RFC 1755	ATM Signaling Support for IP

ASCEND **Remote Networking Solutions That Work.™**

Ascend Communications, Inc.
One Ascend Plaza
1701 Harbor Bay Parkway
Alameda, CA 94502, USA
TEL: 510.769.6001

FAX: 510.747.2300
E-mail: info@ascend.com
Toll Free: 800.621.9578
FAX Server: 415.688.4343
Web site: http://www.ascend.com

2-97
GRF16-AG-02