September 1996

Ascend

FAQ

GRF 400

1. *What is an IP switch?*

   An IP switch is a Layer-3 network device able to route using a rich Layer-3 header, while forwarding IP traffic over a high-speed switch. An IP switch must maintain the high bandwidth and low latency of a switch while handling all the routing functions, including route table lookup and packet forwarding.

   The GRF IP switch is a full-fledged router that uses a fast internal switching engine combined with equally fast route table lookup to achieve wire-speed performance. The GRF is unique because it performs Layer-3 switching over all standard media.

2. *What is driving the need for IP switching?*

   The main factors driving the need for IP switching are:
   - Limited scalability and performance of conventional routers
   - Exponential increase in IP traffic on the Internet
   - New network applications like web browser-based services
   - Increased access requirements for ISPs, carriers and on-line service providers

3. *What is the difference between IP switching and routing?*

   A conventional router's architecture is based on a shared-bus backplane, cache-based route table and centralized processing. These shared-resource routers maintain performance in simple, stable network environments. In a network with a large address space, however, cache hits become less frequent. This leads to bus contention and CPU saturation, causing performance to degrade fast.

   The GRF 400 IP Switch uses a 4 Gb switch and a distributed processing architecture with up to four IP Forwarding Media Cards that perform Layer-3 decisions. Since each card has 1 Gb/s bandwidth dedicated on a per card basis, the GRF scales linearly. GRF media cards implement a hardware-assisted full route table lookup that locates the next hop 100 times faster than software-driven route table lookups. These features enable the GRF to perform Layer-3 decisions at switching speeds, with no degradation in performance in dynamic routing environments.

4. *Why is distributed processing better than centralized processing?*

   A centralized routing architecture consolidates all of the functions of the router in one or a few processors: routing packets, maintaining the route table and protocols, and performing maintenance and management of the device. When the central processor(s) become overloaded, it affects the performance of the entire system.

   The GRF 400's routing architecture distributes routing functions onto its IP Forwarding Media Cards. This eliminates central CPU saturation and bus bottlenecks.

ASCEND

5. *Who needs IP switching today?*

Businesses that provide IP networking services to customers who are demanding access to the Internet. These organizations fall into three groups:

- Internet access providers
- On-line service providers
- Carriers that carry IP traffic on their telecommunications infrastructures

6. *What is the difference between route cache lookup and full route table lookup?*

A route cache is a portion of memory in which a list of most recently used routes are stored. This cache only contains a few routes. If a route is required and it is not in the cache, a centralized processor must perform a full route table lookup to determine the route.

Using route cache becomes less effective as network address space increases. A router's overall performance is impacted by the number of cache misses that occur. For example, a 50 percent decrease in performance is typical with only a 10 percent cache miss. The diversity of the Internet suggests that cache misses in the backbone and at major aggregation points are much greater than 10 percent.

The GRF performs a hardware-assisted full route table lookup faster than a conventional router's cache lookup. The GRF's cacheless architecture is not dependent upon unrealistic cache hit rates for its line-speed performance.

7. *What is NetFlow? Does the GRF use something similar?*

NetFlow is a modification to Cisco's IOS that extends the caching concept to processes such as access lists and filtering. In environments where these processes apply (typically not the Internet), NetFlow improves performance by as much as 10 percent. However, Cisco achieves this performance gain at the expense of extra load on the CPU. Since Cisco's architecture is based on a shared-resource paradigm, any added processing on a per-packet basis can be very costly and lead to CPU overload and router failure.

With NetFlow, Cisco's router caches the results of any additional packet processing in a special cache. All subsequent packets are compared to the contents of the NetFlow cache. If the packet in question requires the same additional processing before being routed, rather than perform the actual processing, Cisco makes use of the information stored in the cache, thereby avoiding the large performance degradation associated with performing the analysis on every packet.

NetFlow sounds like a powerful enhancement. But in fact:

- No switching is ever involved
- Most packets do not need extra processing and never get the benefit of the extra caching
- Only 10 percent improvement is gained when the benefit of caching is available
- NetFlow caching is also available on the new "Versatile Interface Processor" (VIP) cards, but can only be applied to "flows" on networks attached to a given VIP card, other "flows" must use the centralized RSP (Route Switch Processor) facility

The GRF does not require any special flow-based techniques because of its hardware-assisted lookup and low switch latency.

8. *What is the aggregate bandwidth of the GRF switch?*

The GRF 400's switch operates at 4 Gb/s and dedicates 1 Gb/s to each IP Forwarding Media Card.

9. *What is the packet throughput rate of the GRF 400?*

The GRF 400 forwards packets at 2.8 million pps using 256-byte packets.

10. *What are the media supported by GRF?*

The GRF 400 supports the following media:

- 10/100Base-T (8 autosensing ports per card)
- FDDI (4 ports per card)
- HSSI (2 ports per card)
- ATM OC-3c (2 ports per card)
- HIPPI (1 port per card)
- IP/SONET OC-3c (PPP and Frame Relay; 1 port per card; available 12/96)
- ATM OC-12 (1 port per card; available 1Q97)

11. *What routing protocols are supported by the GRF?*

The GRF can support the following popular routing protocols:

- OSPF
- IS-IS
- RIP v1/v2
- BGP3/4
- MTU Discovery
- ICMP
- DVMRP
- CIDR

12. *How many route entries does the GRF support?*

The GRF supports route tables with up to 150K entries. The current Internet has between 35K and 45K routes but the number is expected to grow over the next year or two to about 100K to 125K.

13. *How long does it take the GRF to perform a full route table lookup?*

The GRF uses a hardware-assisted full route table lookup that can be performed in less than 3 microseconds, even when the route table contains 150,000 routes. For most networks, the next hop is found in less than 1 microsecond. This is 100 times faster than software-driven route table lookups.

14. *What management is available in the GRF?*

The GRF supports standard and proprietary MIBs for reads, writes and traps and can be remotely monitored by SNMP network management platforms. The system is configured through a command-line interface similar to Ascend's MAX TNT product.

15. *What are the security features of the GRF?*

The GRF will filter on a single or combination of the following:

- The Protocol (ICMP, TCP, UDP, etc.)
- Source Address
- Destination Address
- Protocol Port Number (Single Number or Range of Numbers) for TCP and UDP
- Established TCP Connections

RADIUS support gives the GRF the ability to authenticate administrators connecting to the GRF.

16. *Does GRF support logical interfaces on the ATM and HSSI cards? How many logical interfaces are supported?*

Yes, the GRF supports up to 512 VCs on the ATM card and 1,812 DLCIs on the HSSI (frame relay) Card.

17. *Is it possible to configure different MTUs for different logical interfaces?*

Yes. Additionally, GRF software supports MTU Discovery, which dynamically sets the MTU size per TCP connection (Path MTU Discovery, RFC 1191).

18. *Is traffic shaping over ATM supported? How many different speed classes can you define?*

Yes, traffic shaping is supported on the ATM cards. A user can define their own classes based on:

- SCR (In 1 Kb/s increments)
- PCR (In 1 Kb/s increments)
- MBR (In 1 Kb/s increments)
- QoS field of (HI, MED, or LOW) for Priorities

19. *Is traffic shaping supported also for incoming ATM SVCs?*

Yes, the GRF can support traffic shaping on our ATM cards.

20. *Is IP Multicasting supported and, if so, what protocols are used?*

We provide full support for routing IP multicast traffic using DVMRP. The following is a feature list of the multicast support:

- The GRF uses IGMP to add and delete interfaces as members in multicast groups to enable data traffic flow
- Mrouted (version 3.8) runs on the GRF and passes information to the kernel and the media cards to update multicast routing tables
- Logical interfaces are defined as multicast capable as part of their configurations
- Traffic is routed directly from one port to another across the switch
- The GRF will not be allowed to be an endpoint of multicast data

Ascend Communications, Inc.
1275 Harbor Bay Parkway
Alameda, CA 94502, USA
TEL: 510.769.6001
FAX: 510.814.2300

E-mail: info@ascend.com
Toll Free: 800.621.9578
FAX Server: 415.688.4343
Internet Home Page:
http://www.ascend.com