Competitive Analysis:

Ascend GRF Vs Cisco GSR

October 1997

Table of Contents

Cisco Positioning of GSR

When Ascend successfully introduced the GRF 400 in 1996, Cisco was faced with an instant credibility problem with the ISP community: a vendor other than Cisco had developed a product that would effectively scale their networks while providing a reliable service – something that Cisco had yet to provide. Until that time, Cisco had been content to provide nominal feature improvements to their traditional routing solutions rather than recognizing the need to develop a next generation High Performance Switching router.

In an attempt to maintain mind share Cisco began a propaganda campaign promoting their next generation router. From the time it was first positioned in customer nondisclosures until its official announcement on September 8, 1997, the name of this next generation router changed often. An early internal acronym was the BFR (the politically correct name being Big *Fast* Router). This was later changed to the GSR (Gigabit Switch Router). Although the GSR acronym is still widely used, the official product naming is the *Cisco 12000 Series Router*. The 12000 family includes the 12004 (4-slot version) and the 12012 (12-slot version).

Cisco has positioned the 12000 toward high-speed ISP backbone requirements and claims that it complements the 7500 series - the 7500 platform is used for dedicated Internet access and aggregation and for Internet backbone applications at OC3 speeds and below. Cisco claims that the 12000 will support IOS software, albeit a ‘stripped-down’ version. Cisco is also claiming that Tag Switching (MLPS) will be supported on the 12000 series in the future.

Quick Comparison: GRF 400 Vs Cisco 12004

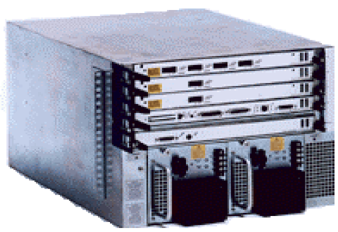

Cisco 12004 14" Height

GRF 400 5.25" Height

	Cisco 12004	**Ascend GRF 400**	**Ascend Advantage**
FCS Date	Q198 claimed	10/96	Market tested and proven
Actual Bandwidth	2.5 Gbps Useable	4 Gbps	60% higher bandwidth: Provides superior scaling and higher performance
Number of Media Slots	3	4	Superior media flexibility
Chassis Height	14" High	5.25" High	Greater bandwidth and density with 1/3 the space requirement: Compact design is critical at space restricted POPs
Media Types Currently Available	OC3/SONET O12/SONET OC12/STM4 ATM	OC3/SONET OC3/STM1 ATM O12/SONET OC12/STM4 ATM HSSI OC3/PPP;Frame Relay 10/100 BaseT FDDI; CDDI HPPI	Extensive WAN and LAN media card offering provides flexible interface options and eliminate the need to purchase separate products.
Maximum PPS Tested	Not Available	1.2 Million PPS (OC12)	Unmatched industry leading performance
NMS	CiscoView (routers only)	NAVIS	NAVIS provides NMS integration with entire Ascend product line (access devices, routers, switches)

Quick Comparison: GRF 1600 Vs Cisco 12012

Cisco 12012 56" Height

GRF 1600 21" Height

	Cisco 12012	**Ascend GRF 1600**	**Ascend Advantage**
FCS Date	End of 1997 claimed	10/97	Market tested and proven
Actual Bandwidth	7.5 Gbps 30 Gbps option 2H98	16 Gbps	2x bandwidth provides superior scaling and higher performance
Number of Media Slots	11	16	Superior media flexibility
Chassis Height	56" High	21" High	Greater bandwidth and density with 1/3 the space requirement: Compact design is critical at space restricted POPs
Media Types Currently Available	OC3/SONET O12/SONET OC12/STM4 ATM	OC3/SONET OC3/STM1 ATM O12/SONET OC12/STM4 ATM HSSI OC3/PPP;Frame Relay 10/100 BaseT FDDI; CDDI HPPI	Extensive WAN and LAN media card offering provides flexible interface options and eliminate the need to purchase separate products.
Maximum PPS Tested	Not Available	4.8 Million PPS (OC12)	Unmatched industry leading performance
NMS	CiscoView (routers only)	NAVIS	NAVIS provides NMS integration with entire Ascend product line (access devices, routers, switches)

GSR Component Pricing and Availability

Components	List Price	Availability
Chassis Options		
"5" Gbps 12004	$14,900	1H98
"15" Gbps 12012 (non redundant)	$24,900	FCS (Q497)
"15" Gbps 12012 (redundant fabric)	$33,000*	1H98
"60" Gbps 12012 (non redundant)	$50,400*	2H98
"60" Gbps 12012 (redundant fabric)	$58,400*	2H98
Software		
System Software	$7,500	FCS (Q497)
Processor Mem Upgd (64 MB Default)		
128 MB Upgrade	$2,900*	FCS (Q497)
256 MB Upgrade	$6,700*	FCS (Q497)
Redundant Processor Option		
Redundant GRP	TBD	2H98
Switch fabric "Slices"		
Switch Fabric Card (SFC)	$8,500*	2H98
Redundant Switch Fabric Card (CSC)	$8,000	2H98
Media Cards		
SONET		
4-port OC3/SONET-MM (MM Mode)	$37,000	FCS (Q497)
4-port OC3/ SONET-SM (Single Mode)	$45,000	FCS (Q497)
1-port OC12/ SONET-MM (Multi Mode)	$25,000	FCS (Q497)
1-port OC12/ SONET-MM (Multi Mode)	$25,000	FCS (Q497)
4-port OC12/SONET- MM (Multi Mode)	TBD	2H98 or later
4-port OC12/ SONET-SM (Single Mode)	TBD	2H98 or later
1-port OC48/ SONET-MM (Multi Mode)	TBD	2H98
1-port OC48/ SONET-SM (Single Mode)	TBD	2H98
ATM		
4-port OC3/STM1-MM (Multi Mode)	$37,000	1H98
4-port OC3/STM1-SM (Single Mode)	$45,000	1H98
1-port OC12/STM4-MM (Multi Mode)	$25,000	FCS (Q497)
1-port OC12/STM4-SM (Single Mode)	$29,000	FCS (Q497)
1-port OC48/ STM4-MM (Multi Mode)	TBD	2H98
1-port OC48/ STM4- SM (Single Mode)	TBD	2H98
LAN		
Gigabit Ethernet (Port Density unknown)	TBD	2H98
Media Card Memory Upgrade Options		
Buffer Memory (32 MB Default)		
64 MB Upgrade (per card)	$1,500	FCS (Q497)
128 MB Upgrade (per card)	$3,500	FCS (Q497)
Route Table Memory (32 MB Default)		
64 MB Upgrade (per card)	$1,000	FCS (Q497)
128 MB Upgrade (per card)	$2,900	FCS (Q497)
256 MB Upgrade (per card)	$6,700	FCS (Q497)
Other		
Cache Engine	$30,000	2H98

*pricing is estimated: Cisco is reluctant to quote actual cost

Price Comparison

Following is a sample price comparison between the Ascend GRF and the Cisco GSR. Notice that the narrow range of media cards for the GSR will both introduce additional product requirements and limit the connection type to these external products. The media limitations of the GSR create a tremendous cost disadvantage when comparing it to the Ascend GRF.

Configuration Example: 1 OC12/SONET; 4 OC3/SONET; 8 100BaseT

Products	Quantity	List Price
Cisco 12004 Components		
12004 chassis (AC power)	1	$14,900
System Software	1	$7,500
1-port OC12/ SONET-MM	1	$25,000
4-port OC3/SONET-MM	2*	$74,000
256 MB Route Table Memory Upgrade	4^	$26,800
7507 Router (with RSP4 and IP software)	1#	$25,900
2-port 100BaseT cards for 7500	4	$73,600
1-port OC3/SONET-MM card for 7500	1	$22,000
	Total	**269,700**
Ascend GRF Components		
GRF 400 chassis (AC power)	1	$15,650
System Software		No charge
1-port OC12/ SONET-MM	1	$24,000
2-port OC3/SONET-MM	2	$40,000
8-port 10/100BaseT	1	$20,000
	Total	**$99,650**

*2nd OC3/SONET card required for connection to 7507
#7507 required for 100BaseT support
^Memory upgrade to ensure full route table support

How Cisco Attacks the Ascend GRF

Cisco Attack: The GRF is not in the "Same League" as the GSR
This is strictly a bandwidth claim by Cisco. Cisco will compare the "60 Gbps" bandwidth of the 12012 to the 16 Gbps of the GRF 1600 and 4 Gbps of the GRF 400.
Ascend Response:
Cisco artificially inflates the actual bandwidth of the GSR by a factor of two – this is explained in detail within the GSR *Architecture Review* section (page 9). The GSR will only provide 7.5 Gbps of actual bandwidth at FCS with a future scaling option to 30 Gbps in 2H98. Today***, the GRF 400 provides 60 percent more bandwidth than the Cisco 12004 and the GRF 1600 provides more than twice the bandwidth of the Cisco 12012.*** Additionally, it is important to remember that both the Cisco 12012 and 12004 require more than three times the shelf space of the GRF 1600 and GRF 400 respectively.

Cisco Attack: The GRF does not match the density provided by the GSR
Cisco is trying to use the artificially inflated bandwidth claims of the GSR, rather than actual port density, as justification that it provides superior density over the Ascend GRF.
Ascend Response:
The GSR provides superior line card density for OC3 media only. The GRF offers more slots (within a much smaller footprint) and in most configurations a single 21" high, 16 Gbps GRF 1600 provides superior density than a 56" high, 30 Gbps Cisco 12012. Additionally, Ascend will deliver higher port OC3 cards in the near future.

Cisco Attack: The GRF is a "First Generation" Crossbar Switch Architecture
Cisco claims that the GRF has no QoS support and that the single FIFO queue per line card suffers from head-of-the-line (HoL) blocking, causing fabric congestion and dropped packets. Thus an ISP cannot charge for differentiated service based on QoS.
Ascend Response:
The GRF currently supports QoS by ensuring that high priority packets will not be dropped from an output queue. Additionally, HoL blocking should not be appended to a statement about QoS – it is an entirely separate issue.

Although statistically possible, encountering HoL blocking on the GRF is mathematically improbable. The GRF provides the industry's highest speed IP forwarding techniques, and combines this with sufficient buffer memory to help alleviate HoL. Additionally, Ascend will be increasing buffer memory and adding hardware queues to future media cards. To date, Ascend has not lost a single sale to HoL blocking.

Cisco Attack: The GRF routing table is limited to 150K entries vs. 250K for the GSR
Ascends Response:
With the exception of the HSSI and HPPI cards (media options which the Cisco GSR does not even support), all other GRF media cards support a routing table of 250k. If HSSI or HIPPI cards are configured in a GRF, then the routing table will be "limited" to 150k per line card within the GRF. However, this is simply not an issue. Use of the word "limited" here implies a shortage, which simply doesn't exist. Today, the size of a full routing table is around 45,000 entries. Additionally, each media card within an Ascend GRF supports the entire route table –the Cisco GSR requires very expensive memory upgrades on each card to support the entire route table.

Cisco Attack: The GRF does not provide the value of Tag Switching or IOS
Ascend Response:
There is no need to, or value in, supporting IOS for Internet applications. In fact, even Cisco is not supporting traditional IOS software on the GSR. This fact is explained in the GSR *Architecture Review* section (page 9).

As for Tag Switching - no vendor supports Tag Switching today. Although Cisco has been touting Tag Switching it is still not ratified as a standard and, in its current form, will not be adopted by the IETF. The evolving industry standard for IP switching is MLPS. Ascend is helping to define this standard. The GRF will support the evolving MLPS standard after it becomes ratified.

Cisco GSR Detailed Overview

GSR Architecture Review

Chassis Options

The 12000 family supports two chassis options: a 4-slot 12004 and a 12-slot 12012. One slot within the 12004 contains the centralized Gigabit Route Processor (GRP), leaving 3 slots for media cards. In the 12012, 11 slots are available for media cards and one slot is reserved for the GRP.

Power Supplies

The Cisco 12000 family supports non-redundant or redundant AC and DC power configurations. Power supplies are load-sharing and hot swappable. The 12000s ship standard with One DC power supply. Cisco claims that a single DC supply is sufficient to handle a fully loaded 120012. A second DC power supply can be added for 1:1 redundancy. For AC configurations, two power supplies are required while a third or fourth AC power supply can be added to provide for 1:1 or 1:n redundancy.

Switch Fabric

The Cisco 12004 provides a static switch fabric that Cisco claims provides bandwidth of 5 Gbps. According to Cisco, the Cisco 12012 provides a switch fabric that can scale from 15 Gbps to 60 Gbps. However, ***Cisco's bandwidth claims are artificially inflated by a factor of two.*** The switch fabric of the 12004 supports a single fabric that provides 622 Mbps of unidirectional bandwidth to each of the 4 media card slots in the system. In reality the 5 Gbps fabric is actually a "2.5x2 Gbps" – meaning that bandwidth is measured by counting unidirectional traffic as it enters the fabric and then again as it exits the fabric.

The same "Cisco math" is applied to the switch fabric measurement of the 12012. The 12012 fabric is implemented by using from 1 to 4 fabric "slices" that enable the system to scale its bandwidth. Cisco claims that a single fabric slice for the 12012 provides 15 Gbps of bandwidth. However these fabric slices only provide 7.5 Gbps of useable bandwidth.

The switch fabric of the Cisco 12012 is physically implemented with two different types of cards – the Clock Scheduler Controller (CSC), and up to three Switch Fabric Card (SFC). The CSC provides the system clock and contains the scheduler algorithm, including a 7.5x2 Gbps slice of bandwidth. The SFC provides incremental 7.5x2 Gbps slices of bandwidth.

The Cisco 12012 ships standard with a single 7.5x2 Gbps fabric. This single fabric slice provides 1.25 Gbps of bandwidth to each of the 12 slots. More specifically, there are two 622 Mbps rails connecting each slot to the fabric slice – one rail supporting inbound traffic and the other rail supporting outbound traffic. The “15 Gbps” 12012 can effectively only support a single OC12 interface on each media slot.

Adding three additional 7.5x2 Gbps fabric slices will increase the overall useable bandwidth to 30 Gbps (not 60 Gbps as Cisco claims). With four fabric slices configured in a 12012, each media slot is provided a total of 5 Gbps of unidirectional bandwidth (four rails of 622 Gbps supporting inbound traffic and four rails of 622 Gbps supporting outbound traffic).

Future OC48 media cards will need to connect to each of the four fabric slices in order to operate. ***Therefore, a 12012 must be configured with four fabric slices simply to support OC48 media cards.***

Cisco will not provide a 12012 option greater than 15 Gbps (7.5 Gbps actual) until 2H98 – providing that OC48 interfaces, or multiport OC12 interfaces are available at that time. In the initial release of the 12012, no combination of available media cards will saturate the bandwidth of a single 7.5x2 Gbps fabric slice.

It is unclear how service is effected when additional bandwidth slices are added to the 12012. Additionally, based on the fabric design, when the standard 7.5x2 Gbps configuration is exceeded, customers are forced to automatically upgrade their 12012 to the 60 Gbps model (30 Gbps actual) – there are no incremental bandwidth increase options as Cisco would lead one to believe.

Fabric Redundancy

There is no fabric redundancy available on the 12004. Sometime in 1H98 the 12012 will support a fifth fabric slice (CSC) that can be added to provide 1:1 fabric redundancy for OC12 and slower media card configurations – this is because only one “active” fabric slice is required for these configurations. In the future, if OC48 media is configured, the redundant fabric slice will only provide 1:4 fabric redundancy – this is because four “active” fabric slices must be configured to support each OC48 media card. Again, it is unclear how service is effected when an additional bandwidth slice is added to the GSR. Additionally, it is unknown how long the re-booting time will be for the redundant fabric slice to come on line.

Route Processor: GRP

Today the 12000 family is controlled by a single centralized Gigabit Route Processor (GRP). The GRP is a R5000, 200MHz processor, providing 64 to 256 MB EDO memory and 20 MB Flash memory. The GRP is not responsible for switching packets. Instead, the GRP's primary function is to run the routing algorithms (BGP, EIGRP, OSPF, and so on) and create the Express Forwarding table.

The Express Forwarding table is a technology that is fundamentally different from what Cisco customers came to expect from their earlier Cisco routers (AGS, AGS+, 7000 and 7500). This technology allows the 12000 family to forward switching decisions to each media card. Express Forwarding sends information including: protocol classification, queuing and route look up which is stored on each media card.

Essentially, the Filtering Engine of the GRP allows the GSR to distribute the routing table on each media card, which will eliminate some of the route caching problems associated with Cisco's previous routers. However, unlike the GRF, which can support the entire routing table on each media card, ***only portions of the routing table are kept on each GSR media card*** (when default memory if used). There are very expensive memory upgrade options for the GSR that will most likely be required in order to support large route table sizes.

Route Processor Redundancy

Cisco is claiming that in a ***future*** release, the Cisco 12000 will support dual GRP configurations. ***The future redundant GRP option will require the use of a media slot***, further reducing the overall port density of the 12000.

System Software

Although Cisco claims that the GSR will "support all the advantages of IOS", the system software offered is not a typical IOS option. Instead, the GSR system software will be a "stripped-down" version of IOS. It will support a minimum of protocols and "Cisco-specific" features. Even this stripped-down IOS offering will prevent the GSR from fully interoperating with 3rd party products – this will reduce the customer's ability to customize their networks or leverage their existing investment in non-Cisco equipment.

For marketing reasons, Cisco purposely "blends" together traditional IOS software options, Tag Switching software, and the separate GSR-specific system software. Viewed by some as their biggest strength in the multi-protocol enterprise world, IOS has become Cisco's biggest weakness in the service provider world. Cisco is forced to continue positioning IOS as their significant differentiator – even if the product does not support it or can not derive any value from it.

Media Cards

Each media card is based on a common architecture that includes:

Layer 3 Switch Processor: Each media card contains a 200Mhz MIPS R5000 RISC processor. It is responsible for making the Layer 3 switching (forwarding) decisions that were initially created and then forwarded from the Express Forwarding table by the centralized GRP.

Route Table: Each media card has configurable route table memory to store the Express Forwarding table information that is distributed to it. Memory is configurable from 32 to 256 MB. Even with the expensive 256 MB memory option, the media cards may not be capable of supporting the entire route table – and the GSR will be subject to ongoing route table lookups to the central GRP creating a bottleneck.

Buffering: Each media card provides 32 MB of buffer memory, expandable to 128 MB. (16MB transmit/16MB receive to 64MB transmit/64MB receive). Cisco claims that this level of buffering will eliminate any round-trip delay problem.

Virtual Output Queue(VOQ): Cisco states that their VOQ function will effectively manage fabric utilization, multicast traffic, HoL performance, and (in the future) class of service provisioning. Up to 17 VOQs will be supported per media card, allowing each card to map a virtual queue to each output on the switch fabric.

Silicon Queuing Engine (SQE): User configurable software that, according to Cisco, helps optimize network utilization. The SQE is simply used to configure VOQs, and set packet discard parameters such as Random Early Detection (RED) and Weighted Random Early Detection (WRED).

Media Card Redundancy

Similar to the Ascend GRF solution, the SONET-based media cards of the GSR will support automatic protection switching (APS). APS is an inherent feature in SONET solutions: protecting against fiber failure, port failure, line card failure or even router failure. Fore each port the SONET cards provide one active port and a redundant port. No other media card options for the GSR provide redundancy options.

Cache Engine

Although each media card stores the information forwarded to it by the Express Forwarding function of the GRP, eliminating the cache miss problems associated with previous Cisco routers, Cisco's look up may still not be robust enough to meet the demands of increased Internet use. Designed specifically for Internet service providers providing large-scale WEB hosting, Cisco has announced a Cache Engine product. The ***Cache Engine is basically a computer workstation that contains special software and requires a local connection to a Cisco router***. It is believed that the local connection will only be via a Gigabit Ethernet connection – which will not be available on the GSR family until 2H98.

Cisco is hiding the fact that Internet service providers may need this external Cache Engine solution to make up for the continued deficiencies in route look up process of the GSR. Instead, Cisco is delivering a "cost savings" marketing spin. Cisco claims that the Cache Engine will reduce WAN usage costs while providing accelerated Web access by storing Web pages in a local network cache. Each Cache Engine, costing $30,000, can support about 500,000 users at a single point of presence and store approximately 25 million Web pages. Up to 32 Cache Engine workstations can be combined to form a cache farm.

GSR Limitations Explained

Many GSR limitations have been highlighted throughout the *Architecture Review* and within the various comparison matrices. In the following paragraphs some of these limitations have been expanded upon, while additional limitations are introduced.

Excessive Size

The 12012 is massive, measuring 56" high and can weigh as much as 380 pounds. The 4-slot 12004 is rather large also – measuring 14” high. The Ascend GRF provides a very compact design – the GRF 400 measures only 5.25” high and the GRF 1600 measures only 21” high. The massive size of the Cisco design will prevent a standard size equipment rack from supporting more than one 12012 – this is significant consideration for service providers as space can be severely restricted and expensive at a POP.

Limited Bandwidth

As explained previously, Cisco artificially inflates the actual bandwidth of the GSR by a factor of two. Actual bandwidth is 2.5 Gbps for the 12004 and 7.5 Gbps or 30 Gbps for the 12012. ***Furthermore, it will be 2H98 before The GSR provides bandwidth beyond 7.5 Gbps.***

Poor Density

The maximum densities of the GSR platforms are extremely poor. The 12004 supports a maximum of 3 line cards, while the 12012 supports a maximum of 11 line cards. Other than OC3 media, the per-card density will be minimal (see line card descriptions on page 6).

Limited Media Card Offering

Today, and for the foreseeable future, the GSR will be severely limited in its media card offering. Restricted to ATM and SONET WAN interfaces prevent the GSR from being a viable option for the mutiservice WAN and LAN requirements prevalent in most POP, NAP and Backbone needs. Additional equipment will be required to provide solutions for mixed media requirements, thus significantly increasing the cost and management complexity.

Expensive Connection to 7500 Router

Cisco claims that the GSR provides an effective connection point for 7500 routers. However, connecting a 7500 router to a GSR is a very expensive proposition. ***Because of the limited media card options on the initial GSR release, the only connection option to the 7500 is via an OC3/POS interface*** - a line card which costs $22k-$25k on the 7500. Furthermore, most customers are utilizing the 7500 for ATM and frame relay speeds of DS3c and lower – adding OC3/POS on a 7500 in these environments will add unnecessary complexity and tremendous expense.

Inefficient Route Table Lookup with Default Memory

As explained in the GSR *Architecture Review* the Filtering Engine of the GRP allows the GSR to distribute the routing table on each media card. This will eliminate some of the route caching problems associated with the 7500 router. However, unlike the GRF, which can support the entire routing table on each media card, only portions of the routing table are supported on each GSR media card – unless the most expensive route table memory options are purchased. This means that the GSR is subject to ongoing route table lookups to the central GRP, which will create a bottleneck.

QoS Story

Cisco is delivering an effective QoS "story", but the solutions are geared toward Enterprise networks – their effectiveness will be minimized for service provider environments. Cisco is putting significant emphasis on queue management to eliminate HoL blocking and provide QoS services. In fact, Cisco is making a push about delivering QoS for the Internet and will try to win the routing business by describing their QoS story. In addition to the VOQ and SQE media card features listed above, Cisco will often present to the customer a white paper regarding their QoS.

The paper addresses all the criteria needed for QoS (input policing, congestion control, prioritized queuing). However, it describes a muddled set of Cisco technologies that will implement this capability (Tag Switching, Netflow monitoring, Express Forwarding, RSVP, and an apparently new capability called Committed Access Reservation CAR). Many of these technologies were developed for Cisco products designed for the Enterprise and furthermore, they will not be available on the GSR in the initial releases.

Disparate NMS Offerings

With numerous acquisitions, Cisco now offers a disparate group NMS packages. Cisco will often promote a WAN solution that combines their 7500 and GSR routers with switches acquired from LightStream and Stratacom. They are trying to plug all these boxes together, pitching end-to-end solutions with robust service level agreements, but they have no integrated NMS solution. Each product family has a different NMS, and each NMS offers a different level of functionality. As service providers put additional emphasis on effectively managing their increasingly complex services, the Cisco NMS weakness will become more apparent.

NAVIS is an integrated family of comprehensive, scalable applications that provides Ascend with a huge competitive advantage. NAVIS supports all Ascend products, and consists of applications that provide intuitive tools for simplified provisioning, accounting, statistics and reporting, fault management, management redundancy, and even customer-based network management. These applications are key for service providers to streamline their existing infrastructure, as well as to leverage them into new value-added service offerings.

Detailed Matrix Comparison

Features	Ascend GRF	Cisco GSR
Architecture		
Bandwidth	**GRF400: 4 Gbps** **GRF1600: 16 Gbps**	12004: 2.5 Gbps (5 claimed) 12012: 7.5 Gbps Today (15 claimed) 30 Gbps Future (60 claimed)
# of Media Card Slots	**GRF400: 4** **GRF1600: 16**	12004: 3 12012: 11
Maximum Performance	**300k pps per media card (via OC12)**	Unknown
NEBS Compliant	**Yes**	Yes
Maximum Port Density:		
ATM		
OC3/STM-1	**GRF400: 8** **GRF1600: 32**	12004: 12 12012: 44
OC12/STM-4	**GRF400: 4** **GRF1600: 16**	12004: 3 12012: 11
OC48/STM-16	**N/A**	Future - 2H98
OC192/STM-64	**N/A**	12004: N/A 12012: Future - 1999
SONET		
OC3 SONET	**GRF400: 4 (increase in '98)** **GRF1600: 16 (increase in '98)**	12004: 12 12012: 44
OC12 SONET	**GRF400: 4** **GRF1600: 16**	12004: 3 12012: 11 (44 Future – 2H98)
OC48 SONET	**N/A**	12004: N/A 12012: Future – 2H98
OC192 SONET	**N/A**	12004: N/A 12012: Future - 1999
PPP/Frame Relay		
OC3	**GRF400: 4 (increase in '98)** **GRF1600: 16 (increase in '98)**	N/A

Features	Ascend GRF	Cisco GSR
Maximum Port Density Continued:		
HSSI	**GRF400: 8** **GRF1600: 32**	N/A
LAN		
10/100 BaseT	**GRF400: 32** **GRF1600: 128**	N/A
FDDI/CDDI	**GRF400: 16** **GRF1600: 64**	N/A
Gigabit Enet	**TBD**	Future – 2H98
HIPPI	**GRF400: 4** **GRF1600: 16**	N/A
Protocol Support		
Routable Protocols	**IP**	IP
Interior Routing Protocols	**RIP, OSPF, IS-IS, EGP**	RIP, OSPF, IS-IS, IGRP, EIGRP, EGP
BGP4	Route reflections MED (Multi-Exit Discriminators) Communities DPA (Destination Preference Attribute) Flat/Weighted Route Dampening Confederations Next Hop-Self Static routing (IGP)	Route reflections MED (Multi-Exit Discriminators) Communities DPA (Destination Preference Attribute) Flat/Weighted Route Dampening Confederations Next Hop-Self BGP multipath Static routing (IGP)
Circuit/Route Features		
Route Table size	**Up to 250,000: Supported on <u>each</u> media card**	250,000: Requires memory upgrade t be supported on each card
SVC Performance	**Future**	Future
Multicast Support	**DVMRP**	DVMRP, PIM
Cell Buffering	**8 MB to 32 MB options per Car**	32 to 128 MB per Card
System Redundancy		
Switch Fabric	**No**	1H98: 1:1 (OC12 and below) 2H98: 1:4 (OC48)

Features	Ascend GRF	Cisco GSR
System Redundancy Continued:		
Processor	**N/A**	Future claim
Media Cards	**SONET Cards**	SONET Cards
Port Level APS	**SONET Cards**	SONET Cards
Redundant FLASH	**Future**	No
Distributed Forwarding	**Yes**	Yes
Distributed Queuing	**Future**	Yes; “Silicon Queuing Engine”
Traffic Control Features		
ATM Service Classes	**All classes**	All classes
Per VC Queuing	**Future**	Yes
Per VC Policing	**Future**	?

Conclusion

The GSR will not be as competitive as Cisco wants the industry to believe. Cisco began positioning the GSR over 12 months ago and has developed a formidable marketing campaign promoting the GSR to potential customers, analysts and the press. The industry trade publications often reported on the pending introduction of the GSR, and although Cisco was careful not to release detailed information on the product, the articles generally portrayed the GSR as a product aimed directly at the Ascend GRF and donned it the "Ascend Killer."

As of this writing the GSR is still not commercially available. For the foreseeable future, the GSR will be limited in regards to media card options, density and features. Additionally, the 12012 is only a 7.5 Gbps router, with a 30 Gbps option not available until 2H98 –drastically different from Cisco’s claims that the GSR will scale from 15 Gbps to 60 Gbps. The Ascend GRF offers superior bandwidth, flexible WAN and LAN media options and superior performance…all at a significantly lower cost than the Cisco GSR.

Furthermore, the GRF has been through exhaustive customer testing and has gained widespread acceptance. Over the next several months Cisco will undoubtedly go through significant problems during their customer implementations. By the time the GSR is ready to be widely deployed, the Ascend GRF will have more than an 18-month time-to-market advantage.

Comments and Questions

Additional competitive information about the subject matter in this document will be disseminated as it is obtained. For comments, questions or additional input, please contact:

Mark Gibson
Ascend Communications Corp.
1 Robbins Road
Westford, MA 01886 USA

Phone: (508) 952-1210
Fax: (508) 392-1484
E-mail: mgibson@ascend.com