Building a Distributed Fabric with Raptor Adaptive Switch Technology (RAST™)

Edwin Hoffman

WHAT IS RAPTOR ADAPTIVE SWITCH TECHNOLOGY?

Raptor Adaptive Switch Technology (RAST) is a multichannel, bit-serial, nibble-parallel, distributed fabric interconnect operating at a raw 12.5 Gigabits per second (Gbps) allowing for an Ethernet packet forwarding rate of 10 Gbps. In November 2004, the raw data rate will rise to support 12.5 Gbps of Ethernet packet forwarding. The special copper line driver technology from Raptor Networks Technology, Inc. allows these channels to communicate over a 45-foot cable that carries all 4 channels to a RAST-compatible system. A fiber version of this RAST technology is also available, which initially drives the RAST signaling from 300 meters on multimode fiber (MMF) to 10 kilometers on single-mode fiber (SMF) using Raptor's standard optical driver modules. Future developments of this Raptor Small Form Factor 10 Gbps (XFP) technology will extend the range to 40 km and beyond as these optics, or other technologies, become available.

WHAT DOES RAST DO?

RAST provides an in-band signaling system that carries Ethernet Packets and special signaling information over a physical connection. The RAST protocol allows RAST-compatible switches to exchange Media Access Control (MAC) address and route information between RAST-connected systems. Allowing these systems to "know" the MAC and route tables of all switch elements in every other switch unit, effectively creates a single-switch system.

Each switch element within the fabric obtains a unique address making each switch unit a separate addressable element within the entire fabric; a switch unit may contain multiple RAST-compatible switch elements. This addressing methodology is but one of the features responsible for the advanced versatility of Raptor's switches for the RAST fabric. Another important feature of RAST technology is the ability to deploy a form of Fiber Distributed Data Interface (FDDI) ring topology.

In a RAST fabric topology, messages containing Ethernet Packets up to Jumbo Packet size of 9216 bytes, take the Shortest Path First (SPF) to the egress element. The RAST fabric form of SPF significantly reduces backbone bandwidth requirements and eliminates the Carrier Sense Multiple Access/Collision Detection (CSMA/CD) and Spanning Tree requirements for any message within the RAST fabric.

Raptor Networks Technology, Inc. achieves this because the RAST fabric protocol directly addresses each switch element. This is accomplished because in the fabric, every switch element is aware of all MAC addresses in the RAST fabric. Thus, an edge element connected via 10 Gbps to another network topology is aware of all MAC addresses needed to access devices on that network via its Address Resolution Protocol (ARP) table including MAC addresses connected to its edge system. The RAST fabric achieves this by listening to Link State Messages from RAST fabric ingress switch systems. If a MAC address is known in an ingress switch, it is also known to the entire RAST fabric.

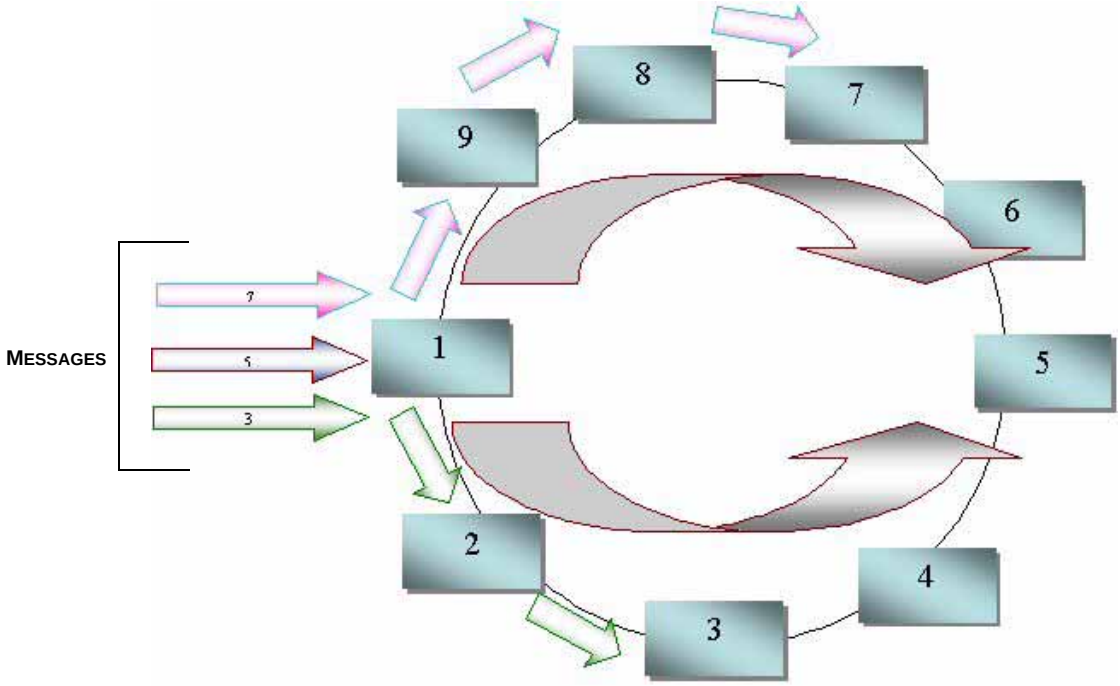

Figure 1: RAST Fabric Message Paths Ring

Figure 1 illustrates three message paths in a RAST fabric. Message 3 requires two hops between RAST 1 and RAST 3. Message 5 requires three hops between RAST 1 and RAST 5, but could also take a four-hop direction if traffic loading is required. Message 7, requires two hops via RAST 9, RAST 8, to RAST 7. In the event of a failure of the ring, RAST automatically directs all traffic in the alternate direction.

RAST fabric provides full support for ring, star, and meshed topologies without need for spanning tree protection or extra protocols, such as Resilient Packet Ring (RPR). Topologies can be used that contain all three of these major topology types.

RAST fabric supports link aggregation. When higher bandwidths are needed and redundant links are required, RAST “binds” together up to four links to create a 40 Gbps pipe (50 Gbps in the future). See Figure 2. These bound links do not need to be connected to the same hardware unit. In certain meshed topologies, this method is used to provide a fully-redundant, meshed core system.

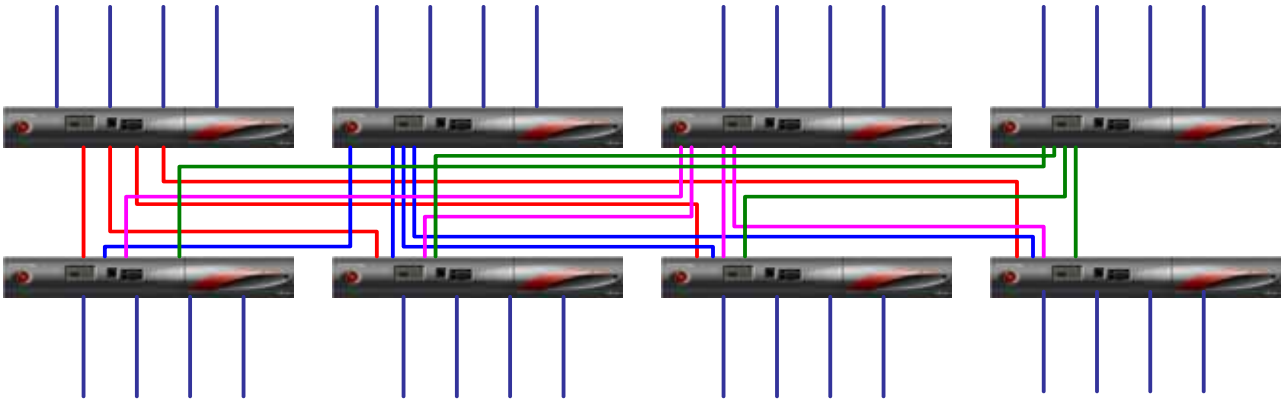

Figure 2: 40 Gbps Pipe using Four 10 Gbps RAST Connections

Of great importance is the capability of RAST fabric in a “Rings of Stars,” which is unavailable to standard Ethernet switching topologies! In standard Ethernet switching topologies, this configuration could only be accomplished with very complicated RPR technology and management. The capability of RAST fabric to automatically add switch elements into a cohesive, single-distributed fabric makes the Raptor Adaptive Switch Technology a unique and very desirable switching topology.

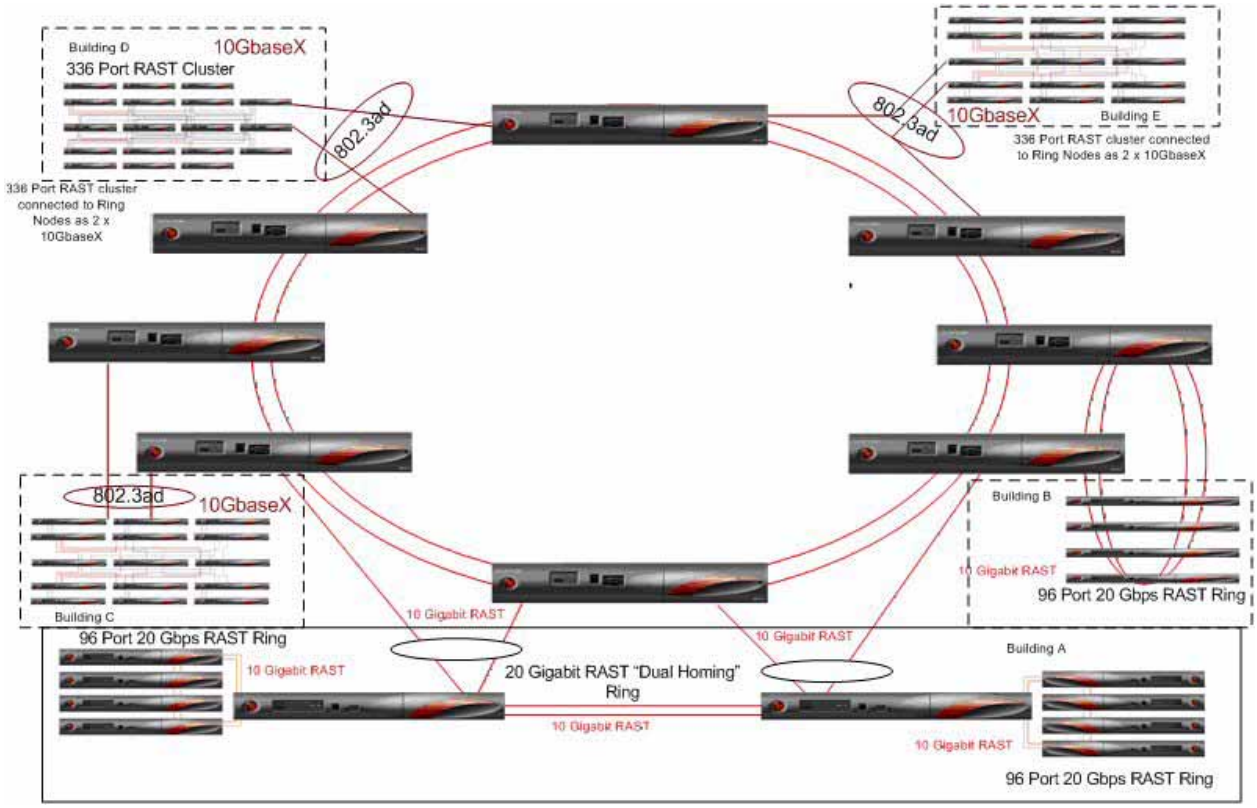

Figure 3: 20-Gigabit RAST Ring

Figure 3 illustrates a combination of 10-Gigabit Ethernet and RAST tree topology networks connected via a Ring comprising a 20 Gbps backbone. This use of RAST ports enables the connection of many multiples of RAST-ported switch elements into a high-speed backbone system. The ultimate speed of this backbone is limited only by the number of RAST ports used. For example, with the Ether-Raptor 2800, a four RAST link ring can be constructed through several ER-1808s resulting in a 40 Gbps ring. Another example is the deployment of two ER-1010 (24 ports, 1000 BaseX, six ports RAST) switch elements connected to each ER-1808 (on the tree side). This provides 48 ports of 1000 BaseX from each Ring access onto a backbone.

The RAST fabric’s multilink, switch-switch connectivity, autoconfiguration, and reconfiguration provide the ability to mirror RAST ports from any switch element to any other RAST port. The RAST fabric can also transport 802.1Q and 802.1p characteristics between RAST ports (switch elements) with the unique ability of transporting 802.3AD Aggregated Link (Trunk) as well.

The Ether-Raptor family supports 802.3AD capability for 32 groups of up to 8 x 1000 BaseX ports in each group. Also, the Ether-Raptor family transports 802.3AD through the RAST fabric via RAST ports, and delivers to an egress port elsewhere in the network or fabric.

The Ether-Raptor family provides RAST support for the following switch elements:

- ER-1010–24 1000 BaseX, plus six RAST, or 10-Gigabit, ports
- ER-1808–8-port RAST core switch, or 8-port 10-Gigabit Ethernet switch

The establishment of a RAST fabric does not require the deployment of core switch elements. This is accomplished via the direct connection of RAST ports between ER-1010 switch elements. In this example, all six RAST ports of an ER-1010 (x 3) can use a 3 In-3 Out RAST Ring (similar to Figure 3) on each member of the RAST fabric. This produces a 30 (37.5) Gigabit Ring with 72 Gigabit Ports and provides for extension to over a 40 km/leg campus, acting as the campus backbone. This RAST fabric ring provides a far higher speed backbone than any Ethernet switch over three fibers, and provides the added features of autoconfiguration and reconfiguration, in the event of failure.

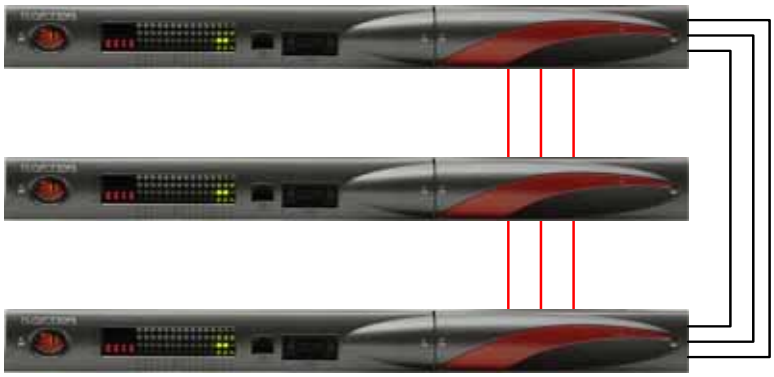

Figure 4: 3 In-3 Out RAST Ring

PORT MIRRORING OVER RAST FABRIC VIA RAST

Hardware inherent in all Ether-Raptor switches allows the user to designate a port on any RAST-connected switch and mirror any port on any RAST-connected switch to that mirror port. Presently Ether-Raptor switches can only mirror the ingress or egress traffic for any port to the mirror port. Management software enables this function.

VLAN 802.1Q OVER RAST FABRIC VIA RAST

Hardware based and software assigned to any Virtual Local Area Network (VLAN) can be transported to any or all switch elements in a RAST network.The user may select automatic VLAN movement or can “assign” the VLAN to particular switches. Ether-Raptor supports both port- and protocol-based VLAN and, in addition, supports GARP VLAN Registration Protocol/General Attributes Registration Protocol (GVRP/GARP) for automatic VLAN transport.

ROUTING OVER A RAST NETWORK

When routing (direct IP forwarding or VLAN routing) over a RAST network, a single “hop” can occur over a large geographical network. For example, in November 2004 up to 64 ER-1010s could be connected in a RAST network. If they were connected with 40 km between these devices, the RAST network would extend over 2560 km, making that single-routed hop very large!

ACCESS CONTROL LISTS (ACL) AND QUALITY OF SERVICE (QOS) OVER RAST FABRIC VIA RAST

Ether-Raptor hardware supports eight priority queues for Ethernet Priority Class of Service. Priority is automatically carried between switch elements over the RAST link, provided VLAN/Priority Tagging is enabled. In addition, a full implementation of per-hop DIFFserv behavior allows complete control over QoS. Any QoS can be applied to the entire RAST cluster as a single command. With access control lists (ACL), Ether-Raptor supports 200 ACL implementations per RAST element. Therefore an ER-1010 supports 400 ACLs with 10 rules each. In addition, an ACL can be applied for the entire RAST cluster with a single command.

LINK AGGREGATION (802.3AD) OVER RAST FABRIC VIA RAST

RAST also allows an 802.3AD link aggregation group to exist with elements from different RAST ported elements within the same RAST fabric.

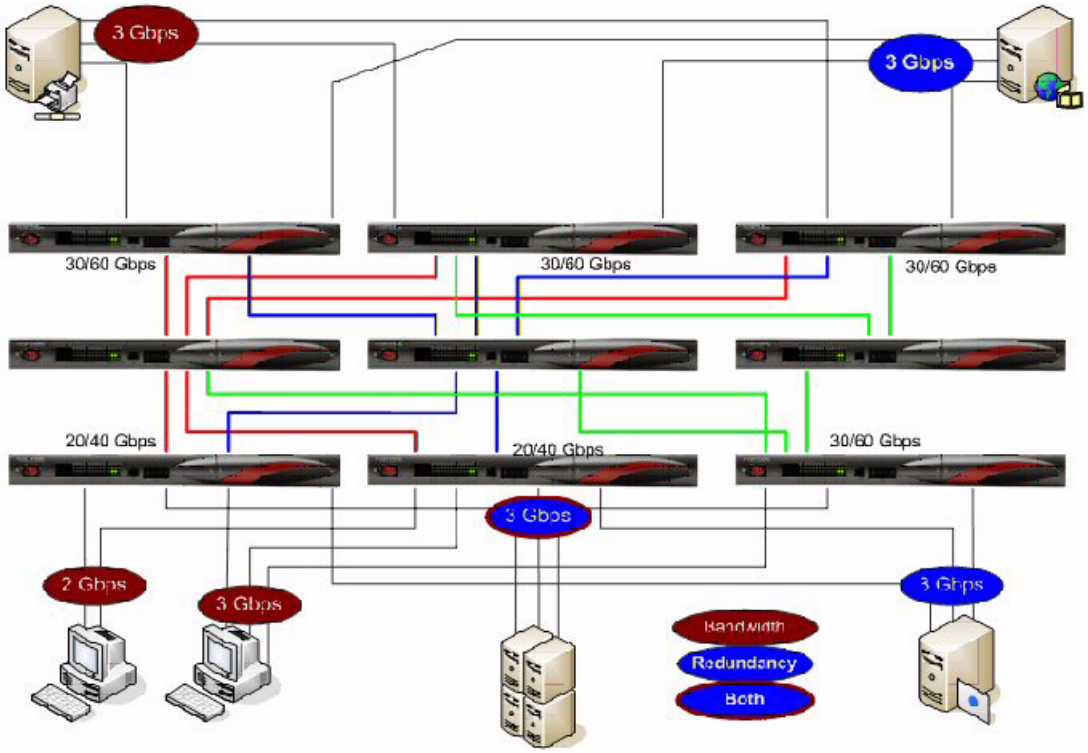

Figure 5: 802.3AD Link Aggregation

MULTIPLE LINK RAST

Whether deploying a RAST fabric ring topology or RAST fabric tree/star topology, RAST ports “automatically discover” any “added” switch. All pertinent information (MAC Table, ARP Table, etc.) is discovered and properties are sent in-band to all RAST-connected switches. Any newly-discovered switch is given a unique number in the RAST fabric and that number is used thereafter to pass messages to/from ports connected to other RAST-connected elements. This in-band messaging does not impinge upon existing messages containing Ethernet Packets moving between RAST elements, thereby providing maximum throughput. With the addition of multiple links between RAST switches, these RAST switch elements automatically load-balance traffic between the two RAST elements and move traffic over to the surviving link if one should fail.

When using a multiple-link RAST as a cascade ring, the same principle applies. All switch elements are “aware” of all possible routes between RAST elements. When a RAST message initiates with the address of another RAST element embedded in the RAST Header, the message takes the shortest path possible to the egress element. When connected over multiple RAST links, the shortest path will also include the alternate link connection.

SUMMARY

Raptor Adaptive Switch Technology (RAST) is an evolutionary jump for network connectivity much as the introduction of FDDI was.

Raptor Networks Technology, Inc. expects RAST to be an open *de facto* standard in the very near future. This will provide other vendors with the ability to integrate into their switch system(s) the availability for Raptor Networks Technology, Inc. and RAST should they wish to license this technology.

1421 E. Dyer Road, Suite 150
Santa Ana, CA 92705
Phone: 949-623-9300
Fax: 949-623-9400