

## Ethernet Protocols Primer for AdvancedTCA

[www.diversifiedtechnology.com/atca](http://www.diversifiedtechnology.com/atca)

Ethernet has proven to be the early leader in AdvancedTCA. This fact is not a surprise considering the widespread acceptance of Ethernet and its continuous penetration into new markets. Even when not used on the extended fabric, Ethernet is present on the base fabric, cementing its place in AdvancedTCA. Having been around for more than 25 years, the Ethernet specification has evolved, and related specifications have been written, to meet the changing needs of networks. The many specifications can make it difficult track which ones are important for a particular application.

With the advent of CompactPCI 2.16 and now AdvancedTCA, system designers must become more familiar with networking to select a switch for their system. Simply specifying “Ethernet switch” as a requirement probably is not enough for most systems. However, requiring the system's switch to support every Ethernet specification greatly increases the cost of the switch by requiring potentially unused features. Some of these Ethernet specifications define important features for AdvancedTCA, while others can be desirable, and still others are unnecessary for most AdvancedTCA systems. (Note that Table 1 lists acronyms for the terms discussed here.)

ACL	Access Control List
BGP	Border Gateway Patrol
CLI	Commnad Line Interface
DiffServ	Differentiated Services
GARP	Generic Attribute Registration Protocol
GMRP	GARP Multicast Registration Protocol
GVRP	GARP VLAN Registration Protocol
IGMP	Internet Gateway Management Protocol
LAG	Link Aggregation Group
MSTP	Multiple Spanning Tree Protocol
NAT	Network Address Translation
OSPFv2	Open Shortest Path First
QoS	Quality of Service
RADIUS	Remote Authentication Dial-In User Service
RIPv2	Routing Information Protocol
RSTP	Rapid STP
SNMP	Simple Network Management Protocol
STP	Spanning Tree Protocol
ToS	Type of Service
VLAN	Virtual Local Area Network
VRRP	Virtual Router Redundancy Protocol

### Important Ethernet features in AdvancedTCA

Low end, unmanaged Ethernet switches might work for a small minority of AdvancedTCA systems. However, the majority of systems will need to make use of at least a few advanced features. Here are some important features that should probably be used in every AdvancedTCA Ethernet switch.

Virtual Local Area Networks are a simple but powerful tool. By segmenting a network with VLANs, security can be increased and broadcast traffic decreased. While nearly every managed switch supports basic VLANs, there are several important extensions to VLANs that are not always supported including GARP VLAN Registration Protocol and protocol-based VLANs. GVRP can greatly ease the configuration of larger networks by propagating VLANs automatically though the network. Protocol-based VLANs allow network segmentation by protocols rather than by physical ports. This is useful for networks that need to support diverse protocols.

Redundancy is a key theme of the AdvancedTCA specification. The cabling to AdvancedTCA systems should also feature redundancy to further increase the availability of the system. To support redundancy in an Ethernet network there are two key protocols, Multiple Spanning Tree Protocol and Link Aggregation Group. Spanning Tree Protocol STP allows for loops and redundant links to exist in networks. It does this by allowing only one active path between any two network devices. Rapid STP is an extension of STP to make the network adapt faster when changes occur. Multiple STP is an extension to RSTP to make STP work with VLANs. Because only one path can be active at a time, the STP protocols provide 1+M redundancy (1 active, M standbys). With a LAG designers can group multiple links between network devices as a single link to increase throughput and provide redundancy. Combining MSTP and LAGs allows for N-way active (all N active) redundancy for higher throughput and higher availability between systems and networks.

In every network some data is more important than other data. Under high network usage the important data should take precedence over the less important data. Quality of Service is a generic term used to describe services that prioritize traffic. There are many different ways to implement QoS with two of the more useful being Access Control Lists and Differentiated Services (DiffServ). ACLs are easy to setup rules that define how to prioritize traffic based on various things such as an IP address, or a Layer 4 port number. DiffServ is a standard that replaces the rarely used Type of Service field of traditional IP and uses it to prioritize traffic. While not as easy to setup as ACLs, DiffServ allows for more powerful rules, which leads to better control of data flow.

Remote management is a must for large networks, and (Simple Network Management Protocol) provides the framework and is definitely an important feature for AdvancedTCA switches. While a CLI provides a management interface, it is difficult to use it programmatically, even if it follows the “industry standard.” SNMP provides remote management and monitoring that is relatively easy to use in a programmatic way. Tools built around SNMP allow for constant monitoring, statistics collection, dynamic reconfiguration, firmware updates, configuration backups, and more on the entire network. SNMPv3 is of particular interest because of the broad security features it brings to SNMP. The end result of SNMP support is less downtime because problems can be identified and solved more quickly.

Specific to AdvancedTCA 3.0 hub boards, cross-connect can be a very important feature. Originally defined in PICMG ENC-3.0-2.0-001 and now an optional part of PICMG 3.0 R2.0, cross-connect is a technology that allows a single switch to connect to both shelf managers in an AdvancedTCA system. This is done by taking channel 1 of the base fabric, on both the backplane and the switch, and breaking it into two 10/100BASE-TX ports, one for each shelf manger, rather than the default 10/100/1000BASE-T port for a single shelf manager. This decouples shelf manager failures from switch failures. While cross-connect is important in certain AdvancedTCA shelves, in other shelves it is not implemented. So the switch should support running with cross connect enabled or disabled.

### **Weighing costs versus benefits**

Every AdvancedTCA system has different requirements of the Ethernet switches used. Often the cost versus benefit analysis of a feature will show that the feature is needed, even when the cost is high. (See Table 2) Here are examples of features that are very important to certain networks while not used in others.

Common Name	Name	Specification	Importance in AdvancedTCA					Typical added cost							
			1	2	3	4	5	1	2	3	4	5			
VLANs	VLANs	IEEE 802.1Q													
GVRP	GVRP	IEEE 802.1P													
Protocol-based VLANs	Protocol-based VLANs	IEEE 802.1v													
Spanning Tree Protocols	Spanning Tree Protocols														
STP	STP	IEEE 802.1D													
RSTP	RSTP	IEEE 802.1W													
MSTP	MSTP	IEEE 802.1S													
LAGs	LAGs	IEEE 802.3ad													
QoS	QoS														
ACLs	ACLs														
DiffServ	DiffServ	IETF RFCs 2474, 2475, & 3260													
SNMP v3	SNMPv3	IETF RFCs 2570-2580 or 3410-3418													
AdvancedTCA Cross-connect	AdvancedTCA Cross-connect	PICMG 3.0 R2.0													
Routing	Routing														
Static	Static														
RIP v2	RIPv2	IETF RFC 1723 & 2453													
OSPF v2	OSPFv2	IETF RFC 2328													
BGP4	BGP 4	IETF RFC 1771													
VRRP	VRRP	IETF RFC 2338													
IGMP Snooping	IGMP Snooping														
GMRP	GMRP	IEEE 802.1P													
Port mirroring	Port mirroring														
RADIUS	RADIUS	IETF RFC 2865 & 2866													
802.1X	802.1X	IEEE 802.1X													
NAT	NAT	IETF RFC1631													

Layer 3 routing support on the switch is not actually required in many AdvancedTCA usage models. Carefully consider a system's requirements before requiring Layer 3 routing because it will carry a large price premium over Layer 2 switching. Nodes within a shelf should be on the same subnet, that is, Layer 2 switching, to minimize latency and network complexity. Even when connecting multiple shelves, a Layer 2 switched network is usually the right choice for the same reasons. If a system connects to multiple external networks, Layer 3 routing must be used. If Layer 3 routing is required, the type of routing protocol must then be carefully considered, as the cost delta between different protocols can be large. For routing within a network, static routing, Routing Information Protocol, and Open Shortest Path First are often used. Designers use Border Gateway Protocol to route between autonomous networks, although this often carries a large price premium over the other routing protocols.

When Layer 3 routing is used, Virtual Router Redundancy Protocol works nicely if the switch is acting as a gateway. With VRRP, you can set multiple routers to appear as a single router to nodes and other routers. This creates highly available gateways for Layer 3 networks. If either switch goes down, the nodes will dynamically failover to use the other switch as the network gateway. VRRP often does not add much cost to a switch. Consider VRRP a requirement if Layer 3 routing is used, because it can be so useful.

Internet Gateway Management Protocol Snooping and GARP Multicast Registration Protocol bring the power of multicasting to Layer 2 networks. Though multicast is not often used in most networks, it is critical to some networks. Without multicast, data is either sent to a single port or all ports of a switch. Multicasting is a Layer 3/Layer 4 technology that allows a router to easily send one set of data to multiple recipients. Both IGMP Snooping and GMRP extend this technology to Layer 2 networks. IGMP Snooping does it by "snooping" the higher level multicast packets and forwarding the packets only to ports that have members of the multicast group in that packet. GMRP works by propagating multicast settings through a network. Either technology reduces broadcast traffic on networks that use multicast.

Often unused, port mirroring can be a valuable tool for addressing network problems. With port mirroring all the data going out of a port, or ports, can be copied to another port for monitoring. Port mirroring can be used in traffic analysis, problem debugging, and even security monitoring.

PICMG 2.16 and now AdvancedTCA are pioneers in backplane Ethernet technology as the specification defining Ethernet over backplanes has yet to be finalized. Each Ethernet specification and feature has its place, but many of them do not fit well into backplane environments and more specifically AdvancedTCA. Here are two examples of features that are not often needed in AdvancedTCA.

Most switches used in AdvancedTCA systems do not need user level authentication because they do not have users in the edge/access sense. Remote Authentication Dial-In User Service and IEEE 802.1X are common technologies that provide user security at the port level. If an AdvancedTCA system does not have any direct users, RADIUS and 802.1X are not required at all. Even if an AdvancedTCA system has direct users, it is more likely the users are connected in via the node boards or AdvancedMCs rather than the switches themselves and thus the switches do not need RADIUS or 802.1X.

Network Address Translation has found a place in many homes and businesses, it does not belong in AdvancedTCA systems. As its name suggests, NAT allows one or more IP addresses on one network to be translated into one or more IP addresses on a different network. This is most often used to hide several nodes behind one IP address. This has the side effect of breaking the end-to-end connectivity of networks. Several popular protocols will not work though a NAT connection and even more protocols will not work through two NAT connections. Because of this NAT is never used deep in the network and thus does not fit into most AdvancedTCA systems.

### **Million-dollar Overkill?**

Carefully selecting which specifications to support is important to keep down the cost of a switch in AdvancedTCA systems. This is even more important, as Ethernet continues to expand into AdvancedTCA, and AdvancedMC Ethernet carriers become more common. Table 2 shows just a fraction of the hundreds of specifications related to Ethernet. As shown, certain technologies can be much more useful than others.

Ethernet switch/routers range in price from 20-dollar consumer switch/routers to million-dollar core routers. In the same way that a moped and a Formula 1 racecar provide the same basic functionality, the consumer switch/router and core router provide the same basic function of creating a network. There are obviously differences between the two devices in features, performance, availability, and manageability among other things though. AdvancedTCA system designers need to weigh features to reach a comfortable balance between cost and functionality. While most AdvancedTCA systems cannot get by with the 20-dollar consumer switch/router type switch, the million-dollar core router is overkill.

*JP Landry has a BS in Computer Engineering from Louisiana State University and works as a design engineer for Diversified Technology, Inc.*

To learn more, contact JP at:

### **Diversified Technology, Inc.**

476 Highland Colony Pkwy

Ridgeland, MS 39157

Tel: 800-443-2667

E-mail: [jp.landry@dtims.com](mailto:jp.landry@dtims.com)

Website: [www.diversifiedtechnology.com/atca](http://www.diversifiedtechnology.com/atca)