



The Next Generation Intelligent Network Infrastructure

How the Enterasys S-Series provides critical capabilities for a
future-proofed Local Area Network

The Next Generation Intelligent Network Infrastructure

How the Enterasys S-Series provides critical capabilities for a future-proofed Local Area Network

Executive Summary

The past decade has brought many changes and advancements to the traditional local area network (LAN). The LAN's evolution has accelerated from simple connectivity for a group of computing devices in a building or campus to a multifaceted service delivery system that connects larger numbers of users and increasingly diverse devices and applications, including IP wired and wireless communications devices, video surveillance, building access security, virtualized computing and applications, and enterprise cloud computing. In addition to providing greater bandwidth, LANs of the future must be more intelligent and possess the ability to automatically adapt and reconfigure connectivity when users and services move throughout the enterprise. The LAN must provide context for who or what is connecting to it, what the user or device's role is and be able to provision role-based connectivity, bandwidth and security. This paper introduces unique and industry-leading Enterasys solutions that simplify day-to-day administration and reduce operational costs while providing a robust and future-proofed network infrastructure.

Introduction

This paper discusses three components of the Enterasys S-Series family of next generation intelligent switch routers:

- Intelligent traffic classification and bandwidth management
- Network survivability during times of oversubscription
- Identity and context-based connectivity provisioning

Ethernet-based LANs are generally designed using tiered layers of devices that form an infrastructure to connect users to applications. LANs are most commonly designed using either two or three tiers of connectivity devices. In general three tier networks are larger and span an area such as an educational or industrial campus. The more common two tier network is found in business environments that exist within a single building. General design principles and economics dictate that the network is more likely to be subject to oversubscription closer to the network edge where users and peripheral devices are connected. This design is acceptable as the average user's computing and communications devices only generate small volumes of traffic. However, as more and more IP-enabled devices connect to the network and real time applications such as VoIP, video and SaaS are deployed, traffic volume exponentially increases to the point where intelligent bandwidth management, application and service prioritization and control are becoming an imperative.

A general design approach to dealing with increased demands on the network is to add additional links or bandwidth between devices to create higher performance paths – known as *“throwing bandwidth at the problem”* – and hope that it will cure any potential bottleneck or choke point between the network edge and core. The alternative is to manage the data utilizing the available links and bandwidth, but the issue is that the increase in administration needed to make adjustments and fine tune the network to reduce the possibility of oversubscription becomes unacceptable to network managers.

Enterasys believes the service oriented network of the future must be intelligent, and use switches that have advanced prioritization mechanisms, granular control over users and devices, automated connectivity provisioning and user/device mobility. Switches used in the network of the future must efficiently manage higher volumes of traffic to ensure that mission critical applications and services will run smoothly and provide the best possible user experience and productivity.

Intelligent Traffic Classification and Bandwidth Management

All S-Series switches include the Flex-Edge feature, which provides the unique capability to classify traffic as it enters the switch. The benefit of this is that the switch is significantly less vulnerable to network congestion issues at peak traffic times. Traffic critical to ensuring the “always up” operational state of the network and to maintain application continuity is identified and prioritized at ingress, prior to being passed to the packet processor. This feature is only available from Enterasys. Network high availability is assured, and important users and applications are guaranteed bandwidth and priority.

With the Flex-Edge feature, traffic can be prioritized in the MAC chip as it enters the switch prior to being handled by the packet processor. Some traffic types such as control traffic are automatically assigned high priority and passed to the packet processor. There are three packet prioritization modes that are configurable, which are described below. The benefit of the Flex-Edge feature is that it allows the switch to intelligently discard low priority traffic in times of network congestion while maintaining continuity for business critical applications and services.

The S-Series Flex-Edge feature provides the following pre-processing functions:

- Ensures that high priority traffic is forwarded to the internal packet processors during periods of congestion
- Asserts flow-control (MAC PAUSE) on front panel ports as needed to avoid dropping packets
- Ensures that lower priority traffic is dropped first in situations where packets must be dropped.

Note: The Flex-Edge feature is completely separate from the S-Series port priority (IEEE 802.1D) configuration. 802.1D support is provided by the S-Series fabric subsystem.

Traffic received on front panel ports is categorized from highest value to lowest value as follows:

Control traffic

- Protocol packets necessary for maintaining network topology and high availability
 - Specific packet types:
 - L2 protocol packets (BPDU, GVRP, LACP, IEEE 802.1X)
 - L3 protocol packets (VRRP, OSPF, BGP, LDP, RSVP, RIPv2, PIM, IS-IS)
 - Hard-coded (not user-configurable)

Favored port

- Packets received on a particular front-panel port
 - Configurable by an administrator

Favored user

- Packets containing a particular SA or SA/SIP
 - Configurable by an administrator

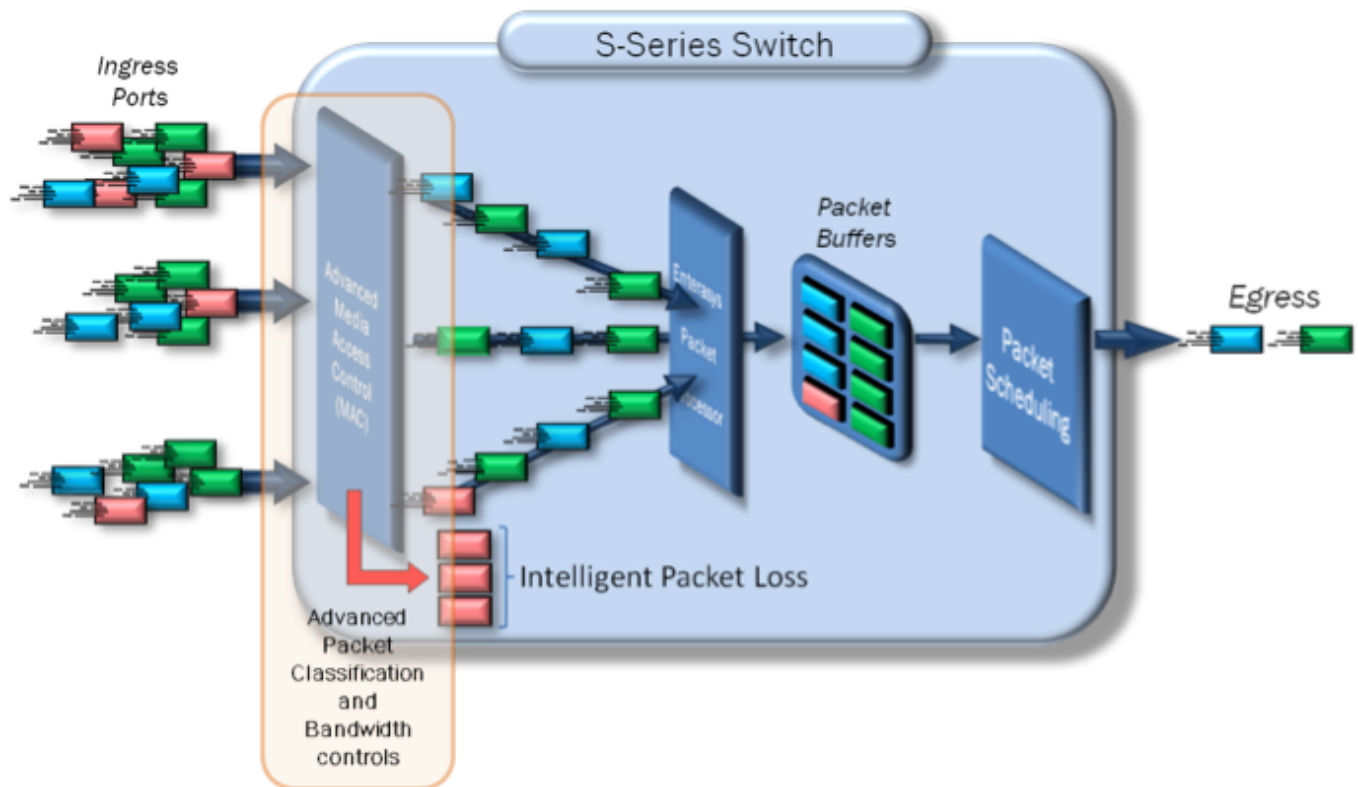
Normal

- All traffic that doesn't fall into any other category

Disfavored

- Configurable by an administrator

The S-Series MAC chip contains an ingress content aware processor which is used to classify packets into the control and favored user categories. The MAC's default port priority settings are used to classify packets into the favored port category.



All packets in each category are mapped to a corresponding internal priority level by the MAC's content aware processor. Each priority level is in turn mapped to a distinct CoS queue. In this context, an egress port is one that is connected to an internal packet processor. The packet processor will then identify users and devices and apply roles and rules (policies) based on the user role within the business or organization. Devices like printers, access points, video cameras, and other devices will also have their role specific set of rules applied.

Network Survivability during Times of Oversubscription

A critical capability for an enterprise switch is the ability to continue operation in times of peak demand. The following section describes the importance of switch port buffering and how it streamlines data flow and improves network performance.

What is port buffering?

The term "buffer" in computing and networking pertains to a memory space where data is stored while waiting to be sent to its destination. An Ethernet switch uses buffering to temporarily hold data content of packets until the switch can send the packets to their destination. The buffering process happens when network congestion occurs at or beyond the egress port of the switch. A buffer is analogous to a funnel; when a funnel is filled it temporarily holds its contents while also emptying its contents at a slower rate than it is filled. Ultimately the funnel is emptied and none of its content is lost. This analogy assumes that the funnel is not continuously being filled at a higher rate than it is being emptied. Buffering is the key mechanism to reduce the contention and packet loss in an Ethernet switching device. Using switches with large buffers reduces packet loss.

Line rate switches and buffering

While it is true that a line rate switch doesn't need buffering to switch traffic in theory, this often fails in practice for a number of reasons. A line rate switch works well in low-latency networks with devices that are all attached at the same speed, and are able to accept data from the switch at line rate. Problems arise with this architecture when the device attached to the line rate switch can't handle any more traffic and sends a pause frame. If this signal is not given the all clear before the buffer fills up, packet loss can occur. This type of situation can cause a chain reaction in a line rate network resulting in poor performance. Another situation that can create this issue is if the uplink is oversubscribed and results in resource contention. For example, with several Gigabit workstations all transmitting at bursts to 500MB over a 1 Gbps uplink, the switch will drop packets as the uplink is oversubscribed and the buffer is saturated. It doesn't take long to overwhelm a typical 256Kb buffer at Gigabit speeds.

Larger packet buffers improve network performance

The larger the buffer, the less likely a packet will be dropped. This is especially important when going from a high speed link (for example, a server connected at 1 Gbps or 10 Gbps) to a lower speed such as a client connected via wireless, a VPN over the Internet, or a slow edge switch (10 Mbps). The reason for this is that a loss in TCP traffic will force a retransmit of the entire TCP window. In other words, the amount of traffic on the network will increase by two to three times for every dropped packet. In a UDP stream such as voice or video, dropped packets result in “clipping” on a voice call or a video stream momentarily freezing. Also, implementing Quality of Service (QoS) can actually increase your total network traffic as rate limited or shaped traffic forces the sending device to retransmit after the switch buffer fills up and packets are dropped.

Buffering and QoS

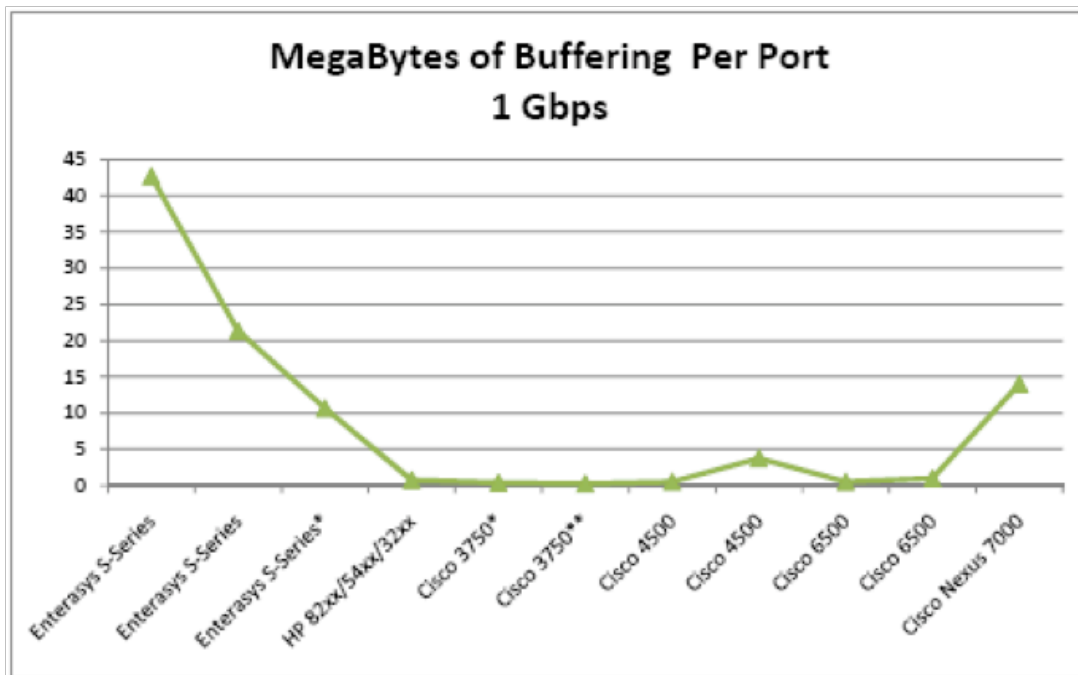
Regardless of architecture or device manufacturer, all routers and switches prioritize traffic by allowing the high priority traffic through first. Mechanisms such as weighted fair queuing are used to provide appropriate handling for lower priority traffic and prevent bandwidth starvation, but they can only achieve so much. The other traffic must be buffered or dropped in order to make room for the higher priority traffic. On a very lightly used link a device with small buffers will work adequately, and often these are the links that really don't need QoS to begin with as they rarely face bandwidth contention. The buffers come into play when a highly utilized link needs to allow priority traffic. If the switch starts dropping packets, this compounds the problem because the devices connected to it are forced to retransmit the packets that were dropped. This process of drop and retransmit increases the demand for bandwidth and lengthens the time required for the transmission.

The LAN's response to insufficient buffering

The relationship between the probability of packet loss and buffering is exponential. If the switch has two times the buffering, it is eight times more likely to NOT drop any packets. This allows the network to avoid packet loss at times of high traffic load. It also means the switch is more efficient at queuing for QoS and for prioritizing without causing additional traffic due to dropped packets. For example, if a slow FTP transfer of 10 Gigabytes of data is rate shaped and is dropping 1% of the packets, a worst case scenario results in a 300% retransmission rate. To look at it another way, the 10 GB transfer just became 30 GB. On the other hand, if the buffers are deep enough to avoid packet loss, the transfer will remain the original size, effectively reducing network utilization by 20 GB.

Buffering in an Enterasys S-Series switch

Each packet processor in an Enterasys S-Series switch has 512 Megabytes of buffer memory. Compare this to commodity switches that often have a shared 8MB buffer for 24 ports – which is 333 Kilobytes of buffers per port. Below are some visual representations of S-Series buffering when compared to common competitors. Note that the only switch that comes close is the Cisco Nexus 7000 with a full line rate card at 10Gbps -- a switch that costs two to three times as much as the Enterasys S-Series.



* - 2 to 1 oversubscribed

** - 4 to 1 oversubscribed

Note: The S-Series dynamically allocates buffer memory to its ports based on traffic pattern and load. The table above indicates the per port buffer memory when the total memory per packet processor is divided equally between the it's ports.

Role and Rules-based Policy Overlay

The S-Series has built-in intelligence that allows it to apply an extensive set of user or device-based roles and rules (policies) that align their usage of the network and its related services with the business and its operational processes. Leveraging a centralized visibility and control point to deliver enterprise-wide policy administration, an IT organization can quickly and effectively adapt to any security or business policy change with a simple action. With a single “click” of a mouse, a policy change can be administered throughout the enterprise. From that point, the S-Series switches that have been configured with policy profiles will enforce them whenever and wherever a user or device connects to the network. The Enterasys Policy Framework can dynamically provision network communications for any user or connected device based on specific rules defining the users’ organizational business role and security requirements. A user will be able to use the network to do their job and the network will be protected from misuse and out-of-compliance activities. No matter where or when a user-based or machine-centric device connects to an Enterasys network, they will be governed by the business and security rules of the organization.

There are three attributes to the Enterasys role and rules-based policy capabilities:

Central role-based administration

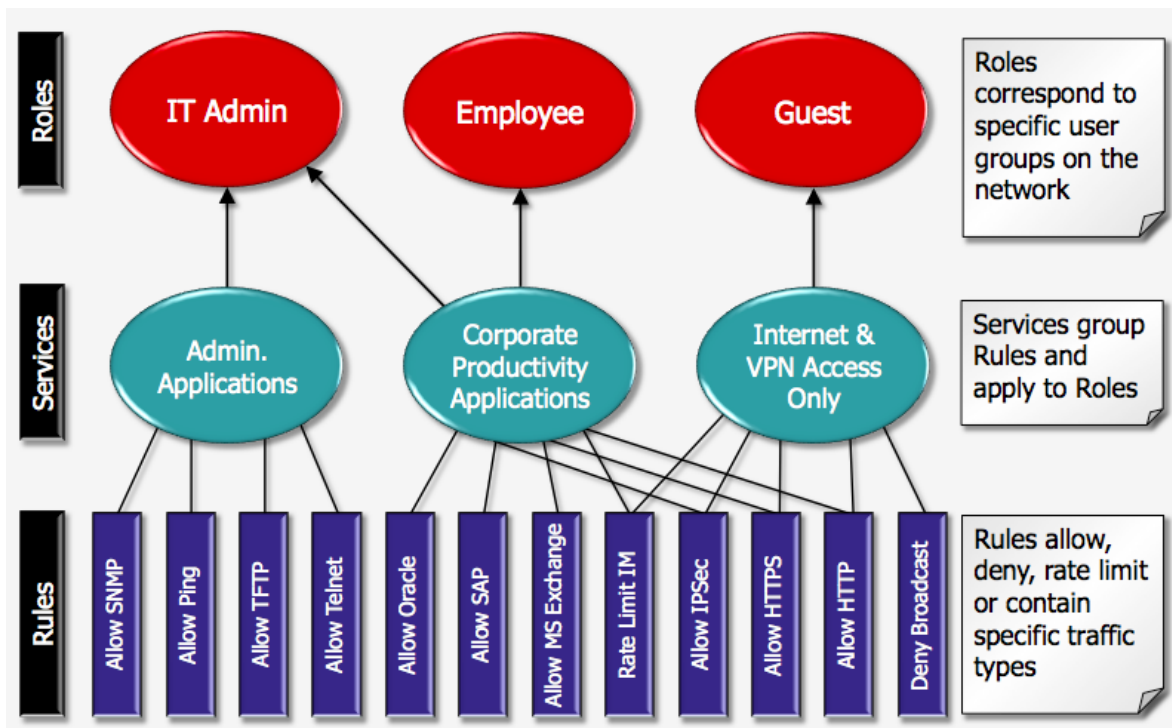
For a set of user/device roles and rules to be valuable they must tightly align with the business and operational processes of a company. In addition, because there are many different organizations within a typical company, all having different business and operational responsibilities, network communication roles and rules must be defined in accordance with the role of the user or device in the business. Role-based administration is the modeling and alignment of network technology, information services, and business processes. The Enterasys NMS Policy Manager application provides the operational interface for role and rules-based administration of company policy.

Authentication and authorization

In order to effectively apply the appropriate network communication rules for the business role of the user or device connecting to the network, an identity must be established for that user or device. User-centric and machine-centric authentication technologies at the access layer of the network must be employed to identify the user and device such as a printer, phone, wireless access point, camera, storage device, etc. Authentication then allows the authorization of the appropriate network usage and associates the user or device to an organizational role. Various access control and authentication types are provided in Enterasys access-layer and distribution-layer switches, and also in Enterasys wireless access layer products. Authentication types include user-based 802.1x, device-based MAC, and guest-based web authentication. These authentication types work in conjunction with industry standard AAA services such as RADIUS servers and directory services (LDAP) to authenticate users and end-systems in the network environment.

Intelligent network infrastructure

Without a network infrastructure that can distinguish various traffic classes, applications, and threats, a meaningful network communications policy cannot be enforced. The network infrastructure where the policy rules will be enforced must be intelligent, meaning that the infrastructure products themselves must classify and manipulate traffic in a manner that aligns with the business policy. For example, if a policy rule says that MS-Blaster should not be introduced into the network, then the access layer products must be able to recognize the specific traffic associated with MS-Blaster and filter only that specific traffic in order to enforce the policy rule defined. Enterasys S-Series products provide this intelligent network infrastructure.



The diagram above shows the relationship between rules services and roles. Multiple rules can be created and grouped together for a specific service. Users and devices are identified and assigned services with their associated roles.

Enterasys ensures only the right users and devices have access to the right information from the right place at the right time through role-based access control policies. Policies are business-oriented based on users, devices and applications, whereas ACLs and VLANs are technology-oriented based on IP addresses. Policy profiles embrace secure user mobility, whereas ACLs and VLANs are relatively static and require a call to the help desk every time there is a move, add or change. With Enterasys, administrators define the roles for the users, applications and devices in the directory servers. Then they define the rules that permit, deny, rate-limit, contain or trigger automated responses to specific traffic types. Just like applications are becoming service-oriented, Enterasys is delivering service-oriented networks. With Enterasys, organizations typically need 1 network operator for every 2,000 to 5,000 users. With other vendors, the need is closer to 1 network operator for every 200-500 employees.

Summary

The Enterasys S-Series provides organizations with the following real world benefits:

1. Flex-Edge provides dedicated network resources to important users, devices and applications to guarantee bandwidth and maintain reliable application delivery.
2. Advanced port buffering streamlines and reduces network congestion due to retransmits and increases network response times. S-Series large port buffers reduce the need for costly and time consuming network upgrades, saving time and money while increasing network response times.
3. S-Series intelligent switching with role and rules-based policy capabilities ensure that users and devices can access applications and network resources from wherever they connect. Operational costs associated with adds, moves and changes are greatly reduced. A single network administrator can manage larger numbers of users and devices than with competitors solutions.

The S-Series is ideal for the service oriented networks of today and the future. LAN infrastructures built using the S-Series intelligently integrate business changing technologies such as virtualization and cloud computing while also supporting the diverse requirements for the next generation of network attached communications devices. Enterasys customers have experienced a significant reduction in operational costs, gained visibility into their networks and implemented control over users, devices and bandwidth.

Glossary

AAA – is an Authorization, Authentication and Accounting architecture that incorporates a generic AAA server with an application interface to a set of application specific modules that perform authorization access control and accounting for devices connecting to a data network

ACL – Access Control List - an ACL specifies which users or system processes are granted access to objects, as well as what operations are allowed to be performed on given objects

BGP – Border Gateway Routing Protocol is the core routing protocol of the Internet

BPDU – Bridge Protocol Data Units - data frames used by the IEEE 802.1D Spanning Tree Protocol (STP) to exchange information about bridge IDs and root path costs between switches

DA – Destination Media Access Control (MAC) address of a device

DIP – Destination Internet Protocol (IP) address of a device

GVRP – Group VLAN Registration Protocol is used for registering VLAN trunking between multilayer switches, and by the GARP Multicast Registration Protocol (GMRP). The latter two are both mostly enhancements for VLAN aware switches, which requires IEEE 802.1Q.

IEEE 802.1X - is an IEEE Standard for port-based Network Access Control (PNAC)

IEEE 802.1D – Spanning Tree Protocol is the IEEE MAC Bridges standard which includes Bridging, Spanning Tree and others

IP – Internet Protocol is a protocol used for communicating data across a packet-switched internetwork using the Internet Protocol Suite, also referred to as TCP/IP (TCP – Transfer Control Protocol)

IS-IS – Intermediate System to Intermediate System Routing Protocol is a protocol used by network devices (routers) to determine the best way to forward datagram's through a packet-switched network. The IS-IS protocol is defined by the ISO/IEC 10589:2002 international standard.

LACP – IEEE 802.3ad Link Aggregation Control Protocol provides a method to control the bundling of several physical ports together to form a single logical channel. The 802.3ad standard was replaced by publication of IEEE 802.1AX-2008 on 3 November 2008.

LDAP - Lightweight Directory Access Protocol, or LDAP is an application protocol for querying and modifying directory services running over TCP/IP

LDP – Label Distribution Protocol is used to build and maintain Label Switch Router (LSR) databases that are used to forward traffic through Multi Protocol Label Switching (MPLS) networks

MAC – Media Access Control

MAC PAUSE – is is a pause frame used by Ethernet flow control which is a mechanism for temporarily stopping the transmission of data on an Ethernet computer network

OSPF – Open Shortest Path First Routing Protocol is a dynamic routing protocol for use in Internet Protocol (IP) networks

PIM – Protocol Independent Multicast is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN or the Internet

RIPv2 - Routing Information Protocol is a dynamic routing protocol used in local and wide area networks

RSVP - Resource ReSerVation Protocol, described in RFC 2205, is a Transport layer protocol designed to reserve resources across a network for an integrated services Internet

SA- Source Media Access Control (MAC) address of a device

SaaS – Software as a Service is the sharing of end-user licenses and on-demand use may also reduce investment in server hardware or the shift of server use to SaaS suppliers of applications file services

SIP – Source Internet Protocol (IP) address of a device

UDP – User Datagram Protocol is a minimal message-oriented transport layer protocol that is documented in IETF RFC 768

VPN – Virtual Private Network is a virtual computer network that exists over the top of an existing network

VRRP – Virtual Router Redundancy Protocol is a non-proprietary (but patented and licensed) redundancy protocol described in IETF RFC 3768 designed to increase the availability of the default gateway servicing hosts on the same subnet

Contact Us

For more information, call Enterasys Networks toll free at **1-877-801-7082**, or +1-978-684-1000 and visit us on the Web at enterasys.com



© 2010 Enterasys Networks, Inc. All rights reserved. Enterasys Networks reserves the right to change specifications without notice. Please contact your representative to confirm current specifications. Please visit <http://www.enterasys.com/company/trademarks.aspx> for trademark information.

