# ExtremeWare Operating System

## Introduction

Extreme Networks® understands that the e-commerce marketplace and 24/7 Internet and Enterprise application access are serious business, especially for the service providers and enterprise network managers tasked with keeping these web sites and application servers running in real time. That's why we developed our ExtremeWare® Operating System, the reliable, resilient software that runs on all Extreme Networks "i" series, and Summit® 200/300/400 series switches. Common code on these Extreme Networks switches means Plug-and-Play compatibility as well as consistent behavior and stable performance right out of the box.

ExtremeWare delivers the uncompromising management, control and security needed on today's demanding service provider, enterprise and co-located networks. Its standardsbased multi-layer switching and Policy-Based Quality of Service (QoS) give service providers and corporate network managers alike the tools they need to make the most of their capacity.

Also key is the flexibility of making your network design decisions for Layer 2, 3 and 4-7 switching independent from your QoS and security policies. For example, even if configured for simple Layer 2 switching, the switch can implement bi-directional bandwidth management and security policies at Layer 3, 4 or even based on user name.

## Multilayer Switching: Layer 2 Features

ExtremeWare can serve as a safety valve, starting at Layer 2. The software can set up and oversee thousands of Virtual LANs (VLANs) per switch. This gives service providers plenty of room to scale their infrastructure to meet customer demand while identifying and facilitating secure traffic flows across their networks. For example, users can be assigned to VLAN groups based on their need to access specific servers or applications. Segmentation of this sort eliminates congestion and allows traffic to be isolated end-to-end, preserving security and using bandwidth efficiently.

ExtremeWare also accommodates all commonly implemented VLAN types, ensuring maximum flexibility. These include schemes based on protocol, port, and Media Access Control (MAC) address, as well as 802.1Q tagging.

## Layer-Independent Quality of Service and Access Control

Even when working at Layer 2, ExtremeWare takes advantage of information available at other layers to establish and enforce QoS and access control policies. Layer-independence means that when an Extreme Networks product is routing or switching packets, QoS and access control decisions can still be made using criteria pulled from Layers 1-4 or from other sources.

That is a critical advantage, since layer-independent intelligence is the key to Policy-Based QoS and access control, which allows service providers to keep tighter control over traffic and offer end-to-end Service Level Agreements (SLAs) that meet bandwidth guarantees while still limiting rates independent of the network design decision to switch or route.

There is a lot more to ExtremeWare's Policy-Based QoS. Filters can be set to scan and classify incoming packets according to user name (Layer 7) or any possible IEEE 802.1p and DiffServ tags values. For instance, minimum and maximum throughputs (expressed as percentages of total available bandwidth) can be assigned to applications, users or user groups. Fine-tuning the lower limit protects delay-sensitive traffic and helps ensure that applications do not slow to a crawl. Setting the upper range prevents oversubscription, thus keeping any one application from hogging resources or overwhelming the network. ExtremeWare additionally allows the mapping and overwriting of IEEE 802.1p tags and DiffServ code points. This is crucial to the successful deployment of an internally coherent Policy-Based QoS network strategy that contains trusted and not-sotrusted sources of user-classified information. The result: an unparalleled level of control over each packet that traverses a service provider's infrastructure or enterprise network.

## Multilayer Switching: Layer 3 Features

ExtremeWare offers an equally extensive set of Layer 3 switching features all geared to increasing control and management on very large networks. Here again, this translates into zippier response times and happier end-users (not to mention shop-till-they-drop e-commerce customers).

The switching software implements the most popular and powerful versions of Border Gateway Protocol (BGP): BGP4, External BGP (EBGP) and Internal BGP (IBGP).

Implementing these protocols gives ExtremeWare the ability to establish route policies and route maps; perform route aggregation; and oversee route redistribution between BGP and other protocols, including Open Shortest Path First (OSPF). BGP also makes it possible to create route communities, which simplifies management on very large networks. Similarly, route confederation and reflection allows IBGP networks to scale almost infinitely without incurring huge management overhead, and MD-5 authentication allows for secure BGP peering. Extreme Networks BGP implementation, sophisticated design and support for BGP Route Flap Damping provides rapid convergence time and scales easily, even in environments with large routing tables.

BGP is only the beginning. ExtremeWare also handles all common router-to-router protocols, including versions 1 and 2 of the Routing Information Protocol (RIP) and IS-IS. Its feature-rich implementation of OSPF has been field-proven on some of largest public and private networks now in service.

OSPF functions include:

- ABR/ASBR
- Equal Cost Multi Path Routing (ECMP)
- Stub/NSSA
- Access Policies
- Automatic link--speed metrics
- Route filtering on ASBR,ABR external and ABA interarea route advertisements
- Loopback passive and direct attach interfaces
- Originate default route
- Password//MD-5 authentication
- Per-function debug tracing tools
- Route redistribution with external route summarization
- Virtual links
- Opaque LSAs

Given that the ExtremeWare Operating System is engineered to meet the needs of large networks, it is no surprise that it supports all common router services, including Dynamic Host Configuration Protocol/User Datagram Protocol (DHCP/UDP) relay, IRDP, and controls on IP Option and Internet Control Message Protocol (ICMP) responses.

ExtremeWare fields a range of IP multicast routing protocols, including PIM Dense Mode (PIM/DM), PIM Sparse Mode (PIM/SM) and Distance Vector Multicast Routing Protocol (DVMRP). Since Extreme Networks switches handle multicasts in hardware, forwarding is done at line speed.

In smaller networks, unicast and multicast routing protocols can be a overkill for the task. Static unicast and multicast routes provide Layer 3 simplicity.

Even though IP advocates argue that their protocol is everywhere, the switching software can deal with legacy IPX traffic using RIP and SAP.

## Network Address Translation

Network Address Translation (NAT) is a feature that allows one set of IP addresses (generally private IP addresses) to be converted to another set of IP addresses (generally public Internet IP addresses). The conversion is done transparently for the hosts by having a NAT device (in this case an Extreme Networks switch) re-write the source IP address and L4 port of the packets. NAT is generally used to conserve IP address space by mapping a large number of inside (private) addresses to a much smaller number of outside (public) addresses. Extreme Networks implementation supports port-based translation as well.

## Policy-Based Routing

In its simplest form, ExtremeWare's policy-based routing services enable an Extreme Networks product to override its routing table and redirect traffic based on the Layer 3-4 information within a packet. This is no-risk routing. If the next hop router is down, ExtremeWare offers immediate failback to the routing tables.

Routing redirection can prove very handy in a number of applications. For instance, it can be used to balance traffic across two separate links, even if only one link is listed in the routing tables. Redirection can also perform a load-balancing function. If multiple next hop routers are defined in the route policy, the switch will load-balance traffic across the routers and perform health checks to allow forwarding is done only to routers that are functioning.

## Server Load Balancing and Web Cache Redirection

Keeping web traffic clipping along is the key to successful sites and satisfied customers. ExtremeWare makes sure that speedy response is the norm, thanks to server load balancing and intelligent redirects to web caches.

Redirection of traffic flows to cache devices for web or other applications get an equally comprehensive treatment. Multiple caches as well as multiple groups of caches are allowed, using multiple redirection rules and groups at wire speed. Cache redirections also takes so-called affinity into account, increasing efficiency by directing users with common destinations to the same cache device. ExtremeWare's built-in SLB capabilities are an attractive choice for deployment scenarios that do not require the performance of a dedicated appliance.

## Multilayer Resiliency

When web sites are unreachable, e-commerce customers go somewhere else and corporate Intranet users can not get their jobs done.

ExtremeWare deals with this basic business truth by offering redundant physical connections with subsecond failover/failback. ExtremeWare implements a mechanism called software controlled redundant port that allows to configure a port to backup a specified primary port. In this capacity, the redundant port will track the link state of the associated primary ports link state. Should the link go down on the primary port, the redundant port will establish link and become active. Due to its simplicity, the feature is very popular to design multihomed redundancy without the complexity of a protocol.

Load sharing (IEEE 802.3ad trunking) also features subsecond failover/failback. ExtremeWare's implementation includes static configuration as well as dynamic via LACP. What's more, redundant physical connections can be teamed with load sharing, so a load-shared trunk fails over as a group.

ExtremeWare implements the 802.1D Spanning Tree Protocol (STP) as well as Rapid Spanning Tree Protocol (RSTP, 802.1w). In Extreme Multiple Instance Spanning Tree Protocol mode, it allows a port or VLAN to belong to multiple STP domains and therefore adds significant flexibility to STP network design, further increasing resiliency. The implementation is PVST+ compatible and 802.1Q interoperable.

Ethernet Automatic Protection Switching (EAPS) is used by ExtremeWare as a means of providing carrier-class availability in an enterprise environment. EAPS is a protocol designed to prevent loops in a Ring topology running Layer-2 traffic. Its role is similar to what Spanning Tree Protocol (STP) accomplishes, however it has the advantage of being able to converge in far less than 1 second when a link breaks as opposed to STP's 30+ seconds. Failover timing only depends on the number of VLANs, not the number of switches; timing will be sub 50ms in most deployments. In its latest implementation, EAPSv2, it supports multiple interconnection points between rings and superloops. To take advantage of spatial reuse and broaden the use of a ring's bandwidth, EAPS supports multiple domains running on a ring at the same time.

ExtremeWare's implementation of the Virtual Router Redundancy Protocol (VRRP) allows a group of switches to function as a single virtual default gateway, with multiple switches providing redundant routing services to users. For instance, if a default gateway fails, a backup router can assume forwarding responsibilities.

Another of ExtremeWare's key resiliency features is the Extreme Standby Routing Protocol™ (ESRP™), which can be implemented at both Layers 2 and 3. These layered redundancy features can be used in combination or independently. ESRP tracks link connectivity, VLANs, learned routes, and ping responses. Additionally, it typically supports failover in 6 to 8 seconds. Basically, ESRP extends the VRRP's capabilities, adding Layer 2 resiliency and loop prevention and Layer 3 default router redundancy. It can be used as a STP substitute and can be scaled to protect thousands of VLANs. In fact, multiple instances of ESRP in the same VLAN allow direct host attachment to standby switches.

Loopback-detect protects from accidental or malicious loops in the networks, for example happening at the edge. Extreme Loop Recognition Protocol (ELRP) is similar as it detects loops, but is geared towards diagnostics usage instead of protection and allows one-time and periodic checks with feedback through CLI and trap/syslog mechanisms.

Snooping Internet Group Multicast Protocol (IGMP) data guarantees that multicast traffic is shunted only to those recipients interested in a particular stream—another way to get the most out of every packet and every pipe. ExtremeWare also supports PIM Snooping. ExtremeWare supports static IGMP membership on a per virtual port basis to force multicast traffic down a particular link without a host having to be present. It also offers IGMP filter capabilities to block traffic from unwanted groups.

## Reducing Edge Complexity Using Unistack™ Stacking

Stacking for the Summit 200-24, Summit 200-48, Summit 300-24, Summit 400-24 and Summit 400-48 is the latest addition to Unified Access Architecture. This allows users to physically connect up to eight individual Summit switches together as a single logical unit. This logical unit behaves as a single switch with a single IP address and a single point of authentication. Network complexity is greatly reduced as a result of configuring multiple chassis as a single switch that can be managed as one entity by EPICenter™ network management software from Extreme Networks. This is a substantial consolidation of managed resources. Stacking in this manner allows for redundant uplinks on two different switches, which is considerably more resilient than having both redundant ports on a single nonstacked switch. Even in the very unlikely event of multiple failures, stacking will continue to operate.

## Scaling up to Accommodate Growth

ExtremeWare gives service providers and enterprise network managers a number of ways to tame the troubles associated with rapidly growing user populations and traffic loads. VLAN aggregation, for instance, is a simple, intelligent way to scale resources when confronted with many isolated subnets, each comprising relatively few devices. This scheme assigns individual customers to sub-VLANs whose members can be grouped together to ease IP address administration and make more efficient use of unique IP IDs. Customers on sub-VLANs share a single router and subnet mask; subnet-like isolation between customers can be maintained. In fact, customers may optionally be prevented from talking to other customers.

VLAN aggregation makes perfect sense in co-location or "multi-tenant " building environments, since it allows an ISP or ASP to create separate sub-VLAN for each customer. This isolates traffic while letting customers connect many small sites or individual users.

The switching software is equally adept with virtual MANs (vMANs). Essentially, this technology makes it possible to deliver transparent LAN services over a metropolitan-area network using Layer 2 point-to-multipoint tunnels across a common physical backbone. These logical pipes carry customer-defined VLANs and their associated tagging schemes, without requiring service providers to get involved. Equally impressive, this scheme enables Extreme Networks switches to handle more than the IEEE 802.1Q limitation of 4,096 VLANs, allowing for unparalleled economies of scale.

Another method to overcome the standards defined 4K limitation for VLAN IDs is a method called "VLAN Translation". This technique allows the remapping of VLAN IDs on the aggregation path to say a metro core and hence combine VLAN IDs that might be managed locally into common VLAN IDs or into the VLAN ID scheme of the core.

## Extensibility

Network change and growth are happening more rapidly now than ever. Implementation of standards based technology allows extensibility to incorporate new applications and services on an as-needed basis. A best-of-breed approach can be used to implement Extreme Networks devices in the Intelligent core and Unified Access, while voice, video or other application specific equipment can be chosen independent of vendor brand.

Link Layer Discovery Protocol (LLDP) is now used in ExtremeWare to aid in the deployment and maintenance of certain network devices that supports the protocol. As a media independent protocol intended to run on all IEEE 802 devices, LLDP may be used to discover routers, bridges, repeaters, WLAN access points, IP telephones, network cameras or any LLDP enabled device, regardless of manufacturer. Extreme Networks employs LLDP not only as a means to simplify deployment of access devices, but also as a trouble-shooting and firmware management tool and eventually as a way to legitimize new services.

## Configuration and Management

Even the most advanced switching features lose their appeal if they are a difficult to set up and oversee. With ExtremeWare this is not an issue. All configuration and monitoring can be via telnet, command line interface, web interface, SNMP v1/v2c/v3 and Secure Shell 2 (SSH2).

To keep web-based management as effortless as possible, a Hypertext Transfer Protocol (HTTP) server is built into every Extreme Networks switch. Even seemingly minor management details have been addressed. For instance, ExtremeWare allows usernames/passwords to be shared, so a single sign-on is all that is needed when using telnet, command line interface, or web browser.

Similarly, all configuration files are in ASCII, so they can be readily edited. New config files can be downloaded and replicated to multiple switches. Switches can be programmed to reboot after downloads or to store new configurations without rebooting.

The same sort of attention has been paid to troubleshooting. ExtremeWare is loaded with debug-trace logging utilities for BGP4, OSPF, RIP, IPX, and Spanning Tree. Traceroutes are available, as is a tool that identifies and sorts the processes running in a switch's CPU.

ExtremeWare's SNMP agent supports the SNMPv1,v2c protocol versions, with multiple community string support as well as SNMPv3, adding security via encryption and authentication.

Remote monitoring and managing are also given the full treatment. ExtremeWare keeps tabs on RMON statistics, alarms, history, and events—all of which are collected at line speed without slowing the switch by a single microsecond, as well as SMON. The switching software also boasts port mirroring capabilities that can completely replicate bidirectional traffic streams, including VLAN traffic on 802.1Q port. Finally, ExtremeWare implements a broad range of SNMP MIBs (listed individually in the technical specifications).

NetFlow provides a way for a switch to capture and export traffic classification or precedence information as data traverses, or flows, across portions of a network. sFlow provides Layer 2 and Layer 3 traffic sampling and statistics. A collector receiving the sampled traffic and statistics provides IT managers a good overview of their network usage on application as well as user IP address level. It allows identifying bandwidth consuming applications and users.

When performing Policy Management in order to manage multiple switches as a cohesive group, ExtremeWare in the switches is coupled with the Policy Manager in ExtremeWare Enterprise Manager, our powerful Web and Java-enabled management tool. This allows end-to-end scoping and provisioning of QoS policies throughout the network infrastructure as well as VLAN provisioning, statistics tools, configuration scripting tools, and many other features.

## Security and Control

Corporate web sites and e-commerce servers are favorite targets of hackers and other cybercreeps. That is why ExtremeWare goes to such lengths to deliver bulletproof security and foolproof access control.

User authentication is nothing short of exhaustive. ExtremeWare implements both Terminal Access Controller Access Control System (TACACS+) and an extended version of the Remote Authentication Dial-In User Server (RADIUS) that can verify user capabilities on a per-command basis. What's more, a different access profile can be used for each remote access method (web, SNMP read, SNMP read/write, Telnet, SSH2). And if management data needs encryption protection, ExtremeWare offers encryption and key exchange with SSH2 (client and server), secure file transfer with SCP and SNMPv3 authentication and encryption.

Network Login is a network access security feature introduced by Extreme Networks. In today's world network security has become necessity for all networks and the network login process is a new feature offered to help solve security problems by giving addresses only to those users who properly authenticate. Network Login does this by putting any port in a VLAN with network login enabled, into a state where it will not forward any packets. All packets sent by a client on this port will not get beyond this port to the rest of the network until authentication using RADIUS servers takes place. In many cases, the RADIUS server will interact with central data repositories for user authentication such as an LDAP directory without putting the burden of the LDAP protocol into the network infrastructure. Network Login supports 802.1x, web-based and MAC-based mechanisms. These methods can be enabled individually or together to provide smooth implementation of a secured network.

The web-based method does not require any specific client side software (a challenge for 802.1x), instead uses standard built-in technologies on clients (DHCP and a web browser), and hence is an easy to deploy security mechanism for all client devices supporting these. As long as the web browser requests traffic from any HTTP server on the network, Extreme Networks switches with Network Login enabled will redirect this traffic to the Network Login welcome page. The login welcome page is configurable to allow posting a custom greeting or for example guest login information for Internet access via a dedicated guest VLAN. A logout window and configurable timeout capability allows usage of Network Login in billed services environments. The web-based method is also an excellent way to deploy 802.1x client software and certificates in a secure fashion on any port without having to open up the network.

Another method that integrates into the existing Network Login infrastructure is using the MAC address on a port to authenticate via RADIUS. This adds an authentication scheme for a wider range of devices such as IP Telephones and security cameras.

Extreme Networks Secure Unified Access allows to dynamically provision user/group based policies—QoS and Access Control Lists (ACLs)—after authentication using Network Login/802.1x. This simplifies QoS provisioning and enables centralized security policies even in a roaming user environment logging in from different ports over time. The mechanism is provided using the Extreme Networks EPICenter Policy Manager.

MAC Address Security allows you to control the way the FDB is learned and populated. By managing entries in the FDB, you can block, assign priority (queues), and control packet flows on a per-address basis. It includes MAC lockdown capabilities on a per port basis and saving learned MAC addresses between reboot. This can be used to for example protect dedicated ports for VoIP phones or printers from abuse. Limiting the number of MAC addresses learned on a port also allows protection from rogue access points or switches or enforcement of service level agreements in tenant or service provider environments. Combined with traps and syslog messages sent out when exceeding the configured limits or MAC addresses or identifying other MAC addresses, MAC Address Security allows identifying port abuse such as rogue wireless access points or hubs/switches on edge ports.

IP Address Security provides two capabilities: DHCP Option 82 and Disable ARP Learning. DHCP Option 82 allows to map IP Address assignments to hosts by port, hence enabling logging for security reasons or to always assign the same IP address to a client on a given port ("virtual static IP addresses").

Disable ARP Learning functions as a "DHCP Enforcement". The switch will only learn through DHCP, not through ARP. This provides protection against accidental or malicious IP address hijacking by a user via static IP address configuration and hence is a security feature as well as a protection against duplicate addresses configured statically out of the DHCP servers address range.

Source IP Lockdown is a unique feature that Extreme Networks has developed to dynamically filter against invalidly sourced traffic. Many different types of network attacks use random source addresses for their traffic. Source IP Lockdown, places "source IP address" filters on all ports automatically. In other words, it will only allow traffic that is sourced from a valid DHCP-assigned address or authenticated user's IP static address to enter the network.

ACLs offer another level of protection against source-spoofing attacks by validating IP addresses. Validation is done at wire speed with no performance impact. Access lists also can be used when grouping traffic for Policy-Based QoS services. ExtremeWare also minimizes the effects of the Denial-of-Service (DoS) attacks used recently to bring several high-profile web sites to their knees. Safeguards include IP Option and ICMP response controls, as well as SYN rate limiting to balance the load on servers if an attack is detected. Besides extensive DoS testing using well-known attacks, ExtremeWare provides automatic protection mechanisms against traffic floods to the CPU such as broadcast, ICMP and Layer 2/Layer 3 learning, etc. When enabled, the feature is capable of automatically setting a temporary hardware ACL if a configured threshold is exceeded. It also includes rate limiting of switch CPU bound traffic such as ICMP. The feature is fully configurable and hence can be applied even for Day-Zero attacks.

Finally, ExtremeWare's routing access policies offer even more protection against denial-of-service attacks. Providers and enterprise networking managers can control service advertisements and establish trusted sources for BGP4, OSPF, IS-IS, RIP, DVMRP, and PIM/DM.

*Note: Differences between product lines do exist. Differences include licensing and functionality differences. Basic and Full licenses differ considerably from Edge and Advanced Edge licenses. Additionally, there are platform specific functionality and performance differences. Please see details at www.extremenetworks.com or contact your local Extreme Networks sales representative.*

**www.extremenetworks.com** **email: info@extremenetworks.**

**Corporate and North America**
Extreme Networks, Inc.
3585 Monroe Street
Santa Clara, CA 95051 USA
Phone +1 408 579 2800

**Europe, Middle East, Africa and South America**
Phone +31 30 800 5100

**Asia Pacific**
Phone +852 2517 1123

**Japan**
Phone +81 3 5842 4011