


# SERVERIRON

- ▶ INTEGRATED LAYER 2 THROUGH LAYER 7 SWITCHING
- ▶ SUPPORT FOR UP TO 16,000,000 CONCURRENT SESSIONS, AND 56 GBPS OF THROUGHPUT
- ▶ HIGH-AVAILABILITY SERVER LOAD BALANCING WITH ACTIVE/ACTIVE CONFIGURATION AND STATEFUL FAIL-OVER
- ▶ INDUSTRY'S MOST POWERFUL CONTENT SWITCHING CAPABILITIES, INCLUDING URL, COOKIE AND SSL SESSION ID BASED SWITCHING
- ▶ CONTENT-AWARE CACHE SWITCHING
- ▶ HIGH PERFORMANCE VPN/FIREWALL LOAD BALANCING
- ▶ ROBUST PROTECTION AGAINST DENIAL OF SERVICE (DoS) ATTACKS
- ▶ MOST COMPREHENSIVE GLOBAL SERVER LOAD BALANCING WITH DNS PROXY AND CLIENT PROXIMITY MEASUREMENTS
- ▶ SUPERIOR SUPPORT FOR ALL MAJOR STREAMING MEDIA PROTOCOLS
- ▶ ENHANCED NAT CAPABILITIES



## INTERNET WEB SWITCHES



### SERVERIRON™ INTERNET TRAFFIC MANAGEMENT SYSTEM WITH INTERNET IRONWARE® LAYER 2 THROUGH 7 SWITCHING

Foundry Networks® award-winning ServerIron Family of Internet web switches provide high performance, Layer 2 through 7 switching, enabling network managers to control and manage today's exploding web transaction, web application and e-commerce traffic flows. Internet IronWare, Foundry's unique software suite of Internet traffic management capabilities – powers the ServerIron web switches to direct requests to the right server and application based on the information that resides beyond the traditional Layer 2 and 3 packet headers. ServerIron eases escalating Internet traffic overload, dramatically increases service availability, reduces the burden of server farm management, and allows the entire web facility to scale to its fullest potential.

Built on Foundry's proven IronCore™ architecture, ServerIron enables Internet traffic switching based on Layer 2 through

Layer 7 definitions. ServerIron delivers industry-leading performance for Internet traffic management functions, including local and global server load balancing, firewall load balancing, and transparent cache switching. ServerIron's superior content-switching capabilities include support for up to 256 URL-based content rules, along with cookie and SSL session ID-based switching. Furthermore, ServerIron provides the foundation for high service availability, disaster recovery, location and server transparency, backbone cost control, and a consistent user experience.

ServerIron supports four major Internet traffic management applications:

1. **Reliable Server Load Balancing (SLB)** – Distribute IP-based services and transparently balance web traffic across multiple servers while continuously monitoring server, application and content health. This enhances overall reliability and availability of the services while simultaneously ensuring server farm accessibility.



2. **Global Server Load Balancing (GSLB)** – Distribute services transparently across multiple web sites and server farm locations and balance the traffic across those sites/servers on a global basis while monitoring web site/server and application health. By directing the client to the best site for the fastest content delivery, ServerIron enhances overall web experience and reduces bandwidth costs.
3. **Firewall Load Balancing (FWLB)** – Increase the network’s overall firewall performance by distributing Internet traffic load across multiple firewalls. Overcome firewall scalability limitations, increase firewall throughput and performance, and improve firewall resiliency by eliminating the firewall as a “single point of failure”.
4. **Transparent Cache Switching (TCS)** – Eliminate the need to configure each client browser, improve Internet response time, decrease WAN access costs, and increase overall web caching solution resiliency by balancing web traffic across multiple caches. ServerIron improves service availability by implementing cache health checking and redirects client requests to the next available cache server or directly to the origin server in the event of a cache or server farm failure.

## Key Benefits

### SERVER AND APPLICATION AVAILABILITY

ServerIron ensures service availability by offering switch, server, link, and session level redundancy.

In the event of a server or application outage, ServerIron provides detection and sub-second fail-over to the next server in a logical group that supports a like service. ServerIron detects application error conditions such as “404 – Object not found” before the client sees them and transparently redirects the requests to other servers without any manual intervention.

To provide 100 percent availability, ServerIron includes an active-standby or active-active redundancy capability that protects against session loss. When enabled, this feature allows network administrators to establish primary and secondary load balancing switches to support identical configuration parameters. In active-standby mode, one unit operates while the other unit sits as a backup. In active-active mode, both units symmetrically operate. In either situation, each switch continuously monitors the health of the other. In the event that one device fails, the other switch takes over without losing sessions or connectivity. ServerIron also supports link-level redundancy that ensures server connectivity by automatically switching sessions from a failed link to a redundant link.

ServerIron’s extensive service health check capability monitors Layer 2, Layer 3, Layer 4, and Layer 7 connectivity and services and determines the servers’ ability to respond to user requests. This ensures fast detection of service problems and eliminates service outage.

### MAXIMUM SCALABILITY

ServerIron supports high port density on both the stackable and chassis platforms, allowing for support of massive server farms and network devices such as firewalls and cache.

Internet IronWare running on ServerIron simplifies network design by enabling network managers to create a server farm, represented by a single IP address known as a Virtual IP (VIP) address. ServerIron appears as a virtual server with a VIP address that controls, monitors, and directs client requests to the most appropriate real server in a server farm. By supporting a wide selection of intelligent load balancing methods, network administrators can transparently and easily scale server capacity regardless of the server platform. ServerIron delivers these benefits without using expensive hardware add-ons or server-side agents.

ServerIron allows ISPs and enterprises to deploy GSLB to transparently expand server capacity on a worldwide basis by redirecting web service requests across multiple data centers located around the world and scale Internet capabilities to global proportions.

For enhanced security and performance, ServerIron’s FWLB eliminates firewall bottlenecks and scales firewall implementations by balancing and distributing load across multiple firewalls. With load balancing support for synchronous, non-synchronous, Network Address Translation (NAT), Layer 2, and Layer 3 firewalls, network administrators can deploy firewalls in the most flexible and reliable manner. ServerIron supports active-standby as well as active-active FWLB configurations. ServerIron’s FWLB supports environments built on static or dynamic routing protocols including RIPV2 and OSPF.

FIGURE 1





FIGURE 2



ServerIron aids in the widespread deployment of Internet caching, used to improve web-service response time and user experience. Network administrators can now use multiple caches to increase redundancy because ServerIron supports the creation of up to four caching server groups per device for added cache reliability. ServerIron distributes load across multiple caching servers for scalability and sends traffic directly to the origin servers if all the caches fail.

#### EASY TO SET UP AND MANAGE

ServerIron is simple to configure and manage using the Foundry Command Line Interface (CLI) or built-in, web-browser based interface. The CLI uses well known Cisco-like commands allowing network administrative staff to easily configure all Foundry products.

In addition, ServerIron's support for Simple Network Management Protocol (SNMP) allows device management using applications such as HP OpenView, available on major server platforms including Sun Solaris, HP-UX, and Windows NT.

The IronView NMS application can be used to monitor and chart SLB and TCS data polled at regular intervals. Formats include bar graph, line graph, and pie chart allowing network managers to easily collect and display detailed information about network traffic destined to server farms. Network administrators can also gauge the amount of traffic between servers and clients, as well as which application is dominating network traffic.

Extensive accounting and statistics allow network managers to easily collect and display detailed information about network traffic destined to server farms. Network administrators can also gauge the amount of traffic between servers and clients, as well as which application is dominating network traffic. In addition, ServerIron tracks the number of active and open sessions per server. These statistics can also be used to track traffic loads on servers that support multiple applications.

#### IRONCLAD WEB PERFORMANCE

ServerIron delivers unmatched Layer 4 through Layer 7 switching performance. All ServerIron platforms utilize a unique distributed switching fabric and powerful processors to deliver load-balancing capacity of over 600,000 connections per second with no session loss. As well, ServerIron scales to maintain 16,000,000 concurrent sessions. This ensures web site availability during peak Internet traffic load.

## Key Features

#### EXCEPTIONAL PERFORMANCE AND CAPACITY

- ▶ **Industry Leader in Concurrent Session Capacity** – ServerIron effectively handles over 16,000,000 concurrent connections to accommodate more client traffic as the web site experiences growing traffic demands.
- ▶ **SwitchBack** – Also known as direct server return, SwitchBack takes advantage of the inherently asymmetrical nature of web traffic. The client-to-server traffic flows through the load balancing device but the return (server-to-client) traffic, which typically consumes more bandwidth because it contains the information that the client has asked for, switches directly to the client on the return path. By avoiding the load-balancing device, SwitchBack provides wire-speed throughput servicing the clients.
- ▶ **Throughput** – The various ServerIron platforms provide differentiated system performance and throughput levels from 2 Gbps with ServerIronXL, right up to 56 Gbps with ServerIron 800 depending on policies enabled on the ServerIron.
- ▶ **Session Processing** – Foundry leads the industry with real-world session processing capabilities of up to 600,000 connections per second.
- ▶ **Symmetric Load Balancing** – Multiple switches can be deployed to increase load-balancing capabilities in parallel and multiply the total connection capacity and overall performance. Also known as an active-active configuration, Internet IronWare's symmetric load balancing provides cross-platform fault tolerance, picking up the full load where the failed switch left off without losing any state information.
- ▶ **Switching Capacity** – Built on Foundry's custom ASIC designs, ServerIronXL, ServerIronXL/G, ServerIron 400 and ServerIron 800 respectively deliver 4.2 Gbps, 32 Gbps, 128 Gbps and 256 Gbps of total switching capacity.
- ▶ **Trunking for Performance** – You can configure trunk groups between ServerIron switches or between the ServerIron switch and the server to increase overall server farm bandwidth, throughput, performance, and redundancy.

ServerIronXL supports up to five trunk groups containing from two to four 10/100Base TX ports, and ServerIron 800 supports up to 22 trunk groups.

- ▶ **IP Filters** – Network managers can define up to 1024 IP filters to selectively control SLB and TCS traffic. These dynamic filters, which take effect immediately without requiring a reboot, match on source and destination IP address, network mask, and TCP/UDP port information.
- ▶ **Overflow Protection** – In the event that the local servers exceed their maximum capacity, Foundry's Internet traffic management systems can load balance the subsequent requests to remote server farms.
- ▶ **Massive Server Farms** – ServerIronXL supports up to 24 10/100 ports and 2 Gigabit ports, ServerIronXL/G supports up to 8 Gigabit Ethernet fiber ports, and ServerIron 800 supports up to 168 10/100BaseTX ports or 56 Gigabit ports. With support for unlimited Virtual IP addresses, and up to 2048 real servers, ServerIron switches provide the highest connectivity to server farms.
- ▶ **High Performance Web Hosting** – ServerIron's many-to-one load-balancing features enable network managers to define multiple VIPs and track service usage by VIP. With this capability, a single server and port number can load balance multiple web sites across a shared set of real servers.

#### SERVER AND APPLICATION HEALTH CHECKS

- ▶ **Layer 3** – Upon configuration, ServerIron immediately and automatically health checks the server via ARP and ping to determine whether the server is ready for operation. If problems arise, ServerIron can automatically take corrective action.
- ▶ **Layer 4** – When the network manager binds an application (TCP/UDP port) on a real server to an application on a virtual server, ServerIron performs a Layer 4 health check, which the server must pass. This guarantees that clients do not hit "dead" servers.
- ▶ **Layer 7** – With health checking enabled for a service (TCP/UDP port), ServerIron can perform a Layer 7 application-specific health check immediately after successful completion of the Layer 4 health check. These can include the following: HTTP, DNS, SMTP, POP3, LDAP, NNTP, IMAP4, FTP, Telnet and RADIUS.

#### LOAD BALANCING METHODS

- ▶ **Round Robin** – Assigns connections sequentially among servers in a logical community. Round robin treats all servers as equal regardless of the number of connections or response time.
- ▶ **Least Connections** – Assigns a connection to the server with the least number of open connections. This option works well in web sites in which there is a group of servers with dissimilar performance capabilities. Least connections ensures adequate distribution and avoids server overload.
- ▶ **Weighted Percentage** – Allows managers to assign a performance weight to each server. Use weighted percentage to ensure that those servers capable of processing connections faster receive the largest number of connections.
- ▶ **Response Time** – Directs requests to the server providing the best response time. Response time measurements include the network latency and the response time for the application-level health checks.
- ▶ **Combined Response Time and Least Connections** – Directs requests to a server based on the server and application response time, and the number of active connections on the server.
- ▶ **Slow Start** – To protect the server from a surging flow of traffic at startup, ServerIron implements a unique slow-start service that allows real servers to gradually accept connections when the server comes up. This is especially useful for SLB implementations using least connections as the load balancing method. Since most servers today cannot handle more than 2000 new connections per second, this feature helps ensure stability when bringing new servers online.
- ▶ **Maximum Connections** – Used to protect servers from bogging down due to high amounts of Internet traffic, this feature allows administrators to limit the number of concurrent connections handled by a particular server and ensures that the traffic does not outpace the performance of the server.
- ▶ **Connection Rate Limiting** – In addition to tracking the total number of connections being sent to a given server, ServerIron can also control both the TCP and UDP rate for a given firewall, cache server or real server. The feature is useful for environments where many short lived connections can be generated beyond the servers capabilities. Connection rate limiting removes the burden from the server, helping maintain a quick response time.
- ▶ **Cookie Switching** – This feature directs HTTP requests to a server group based on information embedded in a cookie in the HTTP header. The next time the client requests information from the server, the cookie specifies which server group should handle the request. Cookie switching ensures that a particular server group always handles requests from a particular client, even across sessions, thereby guaranteeing client persistence and a more satisfying end-user experience.



- ▶ **URL Switching** – ServerIron directs HTTP requests to a server or group of servers, using information contained inside the text of a URL string. URL switching grants the network administrator greater control over the web site deployment to place different web content on different servers, thereby eliminating the constraint of duplicating all content across all load-balanced real servers.
- ▶ **URL Hashing** – Using this feature enables ServerIron to examine information in the HTTP request (either the Cookie header or the URL string) and internally map this information to one of the real servers bound to the virtual server. This HTTP request and all future HTTP requests that contain this information then always go to the same real server.
- ▶ **URL Parsing** – Selects a real server or a group of servers by looking at the prefix or suffix of the URL or by applying a pattern-matching expression to the entire URL. ServerIron supports up to 256 URL rules and imposes no limit on the URL length examined.
- ▶ **SSL Session ID Switching** – SSL (Secure Socket Layer) is a protocol for secure World Wide Web connections used to protect confidential information with server authentication, data encryption and message integrity. In order for SSL to work, all the SSL connections between a client and server must reach the same host. ServerIron ensures that all the traffic for an SSL transaction with a given SSL ID always goes to the same server. It is a “must have” feature for commerce, financial, and shopping-cart based web sites.

#### IRONSHIELD™ SECURITY

- ▶ **Intelligent TCP and UDP Rate Limiting** – ServerIron further protects server farms, firewalls and cache servers by controlling the rate of TCP and UDP connection on an application port basis. Controlling the rate of new connections that each device receives ensures availability despite increased levels of activity due to DDoS attacks or surges in user access.
- ▶ **Network Address Translation (NAT)** – Using ServerIron’s bi-directional NAT, network administrators can assign real servers internal non-routable private addresses to enhance security and conserve address space. ServerIron also supports NAT functionality for environments where the real servers reside in different subnets. This translates to complete network design flexibility when laying out web server farms.
- ▶ **Enhanced NAT** – Unique to Foundry Networks, ServerIron’s generic NAT supports generic UDP, TCP,

and ICMP applications as well as active and passive FTP, Windows™ Media, RealAudio, Real Video, RealMedia, and QuickTime protocols.

- ▶ **Transaction Rate Limiting™** – ServerIron protects real servers against malicious attacks by allowing administrators to configure a threshold for the rate of incoming packets for applications. ServerIron automatically resets servers if clients do not respond within the user-definable timeframe.
- ▶ **SYN-Guard™** – ServerIron protects server farms against multiple forms of Denial of Service (DoS) attacks, such as TCP SYN attacks, by monitoring and tracking session flows. Only completed connection requests are sent to the server. This capability combined with the ability to support over 16 million concurrent connections protects against the highjacking of web server resources by malicious users and ensures connectivity for genuine clients.
- ▶ **High Performance Access Control Lists (ACLs) and Extended ACLs** – By using ACLs, network administrators can restrict access to specific applications from a given address or subnet. Filters can be easily set to deny access to servers by any particular port or VIP address. For example, a network administrator can deny FTP traffic to a particular address. Conversely, filters can be set to allow access by a subset of users or subnets.
- ▶ **Cisco-syntax ACLs** – ServerIron supports Cisco-syntax ACLs, which network administrators can easily cut/copy/paste from their existing Cisco products to drop them right into the ServerIron configuration for the ultimate in portability and security.

#### SESSION PERSISTENCE FOR E-COMMERCE TRANSACTIONS

Internet IronWare maintains up to seven different methods of persistence to ensure that shopping-cart based applications and long-running web transactions proceed accurately. When combined with the ability to handle over 16 million concurrent sessions, ServerIron provides the industry’s best persistence methods to build high-performance networks for e-commerce.

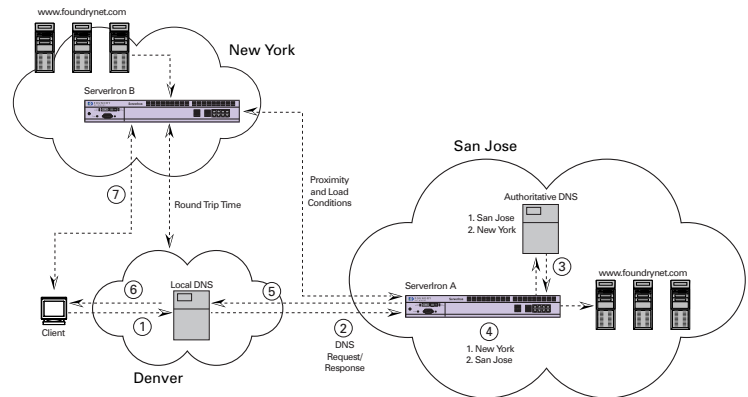
- ▶ **Port Tracking** – Some web applications define a lead port and follower ports. ServerIron ensures that all connections to the follower ports arrive at the same server as the lead port connection. For example, if SSL traffic follows HTTP traffic, the network administrator often defines the SSL port as the follower port to the HTTP port. This ensures that both types of traffic arrive at the same server for a given customer.

- ▶ **Sticky Ports** – ServerIron supports a wide variety of “sticky” connections, including those where web applications require that the client request for additional TCP/UDP ports always go to the same real server, to the server using arbitrary port numbers, or to the server using sequential TCP/UDP ports.
- ▶ **Additional Persistence Methodologies** – ServerIron supports many other types of persistence based on a large range of user programmable options, including Source IP/VIP/Port, Source IP/VIP, and SSL Session ID.
- ▶ **Mega Proxy Server Persistence** – Network managers can configure ServerIron to treat a range of source IP addresses as a single source to solve the persistence problem caused by certain mega proxy sites in the Internet.
- ▶ **Comprehensive Session Persistence** – ServerIron expands upon simple cookie-based switching by including support for Source IP based persistence. When configured for cookie-based persistence, ServerIron uses Source IP-based session persistence if there is no cookie present. This provides a superior level of session persistence over other vendor implementations.

#### HIGH AVAILABILITY SERVICES

- ▶ **Remote Backup Servers** – If no local servers or applications are available, ServerIron sends client requests to remote servers. The remote server can be another server farm managed by another ServerIron or just another group of real servers. By defining multiple remote servers, ServerIron will load balance the traffic among them using the selected load balancing method.
- ▶ **HTTP Redirect** – ServerIron can also use HTTP redirect to send traffic to remote servers if the requested application is not available on the local server farm.
- ▶ **Active/Standby** – When deployed in active-standby mode, the standby ServerIron will assume control and preserve the state of existing sessions in the event the primary load-balancing device fails.
- ▶ **Active/Active** – When deployed in active-active mode, both ServerIron switches work simultaneously and provide a backup for each other while supporting stateful fail-over.
- ▶ **Quality of Service** – Network administrators can prioritize traffic based on ports, MAC, VLAN, and 802.1p attributes, as well as by service port and application type; for example, ServerIron can grant priority to HTTP traffic over FTP.
- ▶ **Redundant Power Supplies** – ServerIronXL supports an optional built-in redundant power supply and

FIGURE 3



ServerIron 800 supports multiple (up to three additional) hot-swappable power supply options.

#### EASE WEBSITE MANAGEMENT

- ▶ **Mix and Match Servers** – ServerIron increases network design flexibility and investment protection for existing resources by allowing for application distribution on different servers based on performance requirements, with dynamic adjustment and without interruption of service.
- ▶ **Graceful Shutdown** – ServerIron allows you to gracefully remove a server from the load balancing rotation without disrupting any existing session. This increases web site stability when upgrading or repairing servers.
- ▶ **SNMP Reporting** – ServerIron reports SNMP traps and events to the specified trap receiver and can send those management messages to up to six different SyslogD servers.

#### GLOBAL SERVER LOAD BALANCING

ServerIron redirects client traffic globally by site availability, site load, and site response time. ServerIron also measures client/server proximity as defined by round trip delay and geographic location. All these features can work in conjunction with the network's existing DNS servers and results in minimizing network disruption when implementing GSLB. ServerIron continually monitors the sites to detect any changes in servers or services due to varying health and traffic conditions. Configurable site load thresholds enable network administrators to fine tune the health checking parameters to best suit the site's web server and service capabilities. [See Figure 3]

FlashBack™, a unique Foundry Internet traffic management feature, measures response time between the authoritative DNS site and each data center. ServerIron actually computes proximity based on the round trip time between the different global sites and uses these measurements as a benchmark to determine site responsiveness and the site's range to the client.

In addition, ServerIron understands geographically-based site selection to keep the requests within continental domains. ServerIron constantly monitors web traffic to create a knowledge base that enables a more intelligent GSLB methodology, powering smarter site selection criteria. ServerIron handles server farm outages or overloads even after the DNS lookup by automatically triggering HTTP redirection or implementing remote server load balancing. ServerIron GSLB provides the following:

- ▶ Acts as a DNS proxy to transparently intercept and modify the DNS responses, thereby directing customers to the best site.
- ▶ Handles the server farm outage or overload after the DNS lookup by HTTP redirect or remote server load balancing.
- ▶ Leverages the existing DNS server and minimizes disruption to the existing DNS environment.
- ▶ Provides continuous site monitoring to detect changes in site health conditions.
- ▶ Allows the network administrator to tune individual site load thresholds through configurable settings.
- ▶ Monitors and selects sites based on FlashBack speed, which measures performance using site, server and application responsiveness.
- ▶ Adds an evolutionary knowledge based in the global server load balancer that enables smarter site selection as more clients access the site.
- ▶ Grants the network administrator the ability to tune ServerIron settings and ensure that minimal differences in various site metrics do not affect site selection.

Foundry offers another unique Internet traffic management feature called “Global IP”. Working alongside Foundry Layer 3 switches (BigIron® and NetIron®) and powered by unique VIP health checking algorithms, Global IP injects routes to client web services based on regional data center server-farm health and provides GSLB capability without relying on the DNS protocol. This enables ISPs and enterprises to instantaneously build networking infrastructures on a global basis and offer ready-to-go differentiated services to customers for multi-site deployments.

### TRANSPARENT CACHE SWITCHING

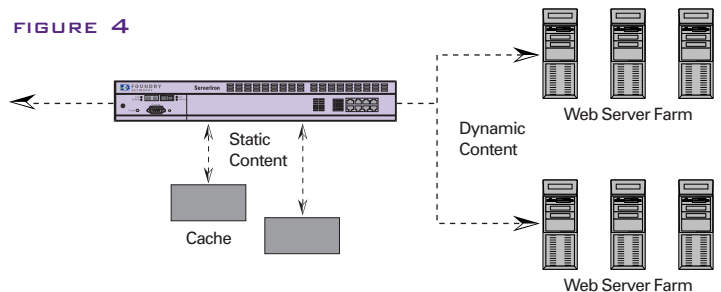
ServerIron switches improve Internet response time and decrease WAN access costs by redirecting web traffic destined for remote Internet hosts to local cache servers. Foundry’s Transparent

Cache Switching (TCS) can be used with any cache server that supports transparent redirection, including those from leading vendors such as CacheFlow, Cobalt, Compaq, Dell, Inktomi, Network Appliance, and Novell.

Foundry’s TCS offers network managers a resilient web caching solution that significantly simplifies administration. Unlike proxy server solutions that require manual configuration of each client’s browser, ServerIron transparently intercepts and switches HTTP client requests to an available cache server without reconfiguring the client’s browser. Network administrators can configure the device to switch traffic based on source and destination IP address. [See Figure 4]

ServerIron switches provide the industry’s most powerful content-aware cache switching features to build intelligent content networks that route traffic based on content rather than just IP addresses. ServerIronXL/TCS is exclusively optimized for content-aware TCS and is targeted to remove price/performance barriers to the mass deployment of content-aware cache switches. ServerIronXL/TCS includes the following content-switching features:

- ▶ Intelligent load balancing of caches to eliminate content duplication, increase cache-hit ratio, and improve the Internet response time.
- ▶ Accelerate delivery of dynamic content and optimize the cache utilization by bypassing the caches automatically for dynamic content.
- ▶ Specify content-based rules to determine what content should be cached.
- ▶ Organize caches into logical groups that serve different content to provide differentiated service offerings to content providers.



# Technical and Physical Specifications

## LOAD BALANCING

### METHODS

Least connections  
Response time  
Response time + least connections  
Round robin  
Weighted distribution

## LAYER 2 SWITCHING

### CAPABILITIES

32,000 MAC addresses  
802.1d Spanning Tree Protocol  
802.1p prioritization  
Policy-based VLANs  
Port-based VLANs  
Layer 3 protocol VLANs  
Layer 3 protocol and subnet VLANs  
802.1q VLAN tagging

## PROTOCOL SUPPORT

TCP  
UDP  
SSL  
FTP  
Telnet  
SMTP  
HTTP  
IMAP4  
LDAP  
NNTP  
POP3  
DNS  
BootP  
TFTP  
SNMP  
VRRP/VRRPe

## STANDARDS

### COMPLIANCE

802.3, 10BaseT  
802.3u 100BaseTX,

100BaseFX  
802.3z 1000BaseSX  
802.3z 1000BaseLX  
802.3x Flow Control  
802.1q VLAN Tagging  
802.1d Bridging  
802.3 Ethernet Like MIB  
Repeater MIB  
Ethernet Interface MIB  
SNMP V1  
SNMP MIB II

## NETWORK MANAGEMENT

Integrated Command Line Interface  
Web-based GUI  
Telnet  
SNMP  
RMON  
HP OpenView for Sun Solaris, and Windows NT Standalone  
Windows NT

## WARRANTY

1 year hardware  
90 days software  
Upgrades to higher levels available

## MOUNTING OPTIONS

19" Universal EIA  
(telco) Rack  
Tabletop

Specifications subject to change without notice.

	SERVERIRONXL	SERVERIRONXL/G	SERVERIRON 400	SERVERIRON 800
Concurrent sessions	1,000,000	1,000,000	16,000,000	16,000,000
Throughput	2 Gbps	8 Gbps	24 Gbps	56 Gbps
Switching capacity	4.2 Gbps	32 Gbps	128 Gbps	256 Gbps
Packet forwarding rate (packets per second)	3,000,000 pps	12,000,000 pps	48,000,000 pps	84,000,000 pps
Number of Virtual IP addresses	Unlimited	Unlimited	Unlimited	Unlimited
Number of real servers	1,024	1,024	2,048	2,048
Number of 10/100 ports	24	N/A	72	168
Number of Gigabit ports	2	8	24	5 6
Total number of ports	26	8	72	168
Layer 3 switching capabilities	Supports servers on different subnets from that of Virtual IP address	Supports servers on different subnets from that of Virtual IP address	OSPF, RIP V2	OSPF, RIP V2
Physical dimensions	2.75" h x 17.5" w x 16.75" d (6.7cm x 44.5cm x 42.2cm)	2.75" h x 17.5" w x 16.75" d (6.7cm x 44.5cm x 42.2cm)	8.75" h x 17.5" w x 15" d (22.2 x 44.5 x 38.1 cm)	20.75" h x 17.5" w x 15.25" d (52.7 x 44.5 x 38.7 cm)
Weight	18 – 22 lbs (8 – 10 kg)	18 – 22 lbs (8 – 10 kg)	60 lbs fully loaded (29.9 kg)	117 lbs fully loaded (43.7 kg)
Power Requirements Specifications	110v/220v auto-sensing	110v/220v auto-sensing	<b>4-slot Chassis with Single (1) Power Supply:</b> <b>Input Voltage and Current</b> <b>Power Supply Rating</b> -70 to -40 VDC: 17A 100 to 120 VAC (auto-ranging): 8A 200 to 240 VAC (auto-ranging): 4A AC line frequency: 47-63 Hz	<b>8-slot Chassis with single (1) Power Supply:</b> <b>Input Voltage and Current</b> <b>Power Supply Rating</b> -70 to -40 VDC: 17A 100 to 120 VAC (auto-ranging): 8A 200 to 240 VAC (auto-ranging): 4A AC line frequency: 47-63 Hz



FOUNDRY NETWORKS

2100 Gold Street • P.O. Box 649100 • San Jose, CA 95164-9100 • Tel 408.586.1700 • Fax 408.586.1900 www.foundrynetworks.com