ipinfusion™

Product Bulletin

ZebOS®-XP 1.2

The networking industry is moving away from the hardware dominated solutions and adopting networks defined by software platforms. The industry demands of breakneck speeds requirements of billions of new mobile users, the advent of Big Data, and the emergence of IT as a servicehas necessitated this industry change.

The widespread availability of commercial networking chipsets changes the economics of the network, driving commoditization of hardware and the emergence of software-defined networking (SDN). IP Infusion is at the core of this networking revolution. The emergence of cloud computing and network virtualization introduces new requirements to innovate at software speeds. In cloud computing and network virtualization, the challenge is to remove the network as a bottleneck to ubiquitous computing.

IP Infusion addresses these challenges with an SDN architecture that features:

- Split Control Plane design
- Modularity and programmability
- High degree of scalability
- Hardware abstraction to operate on a variety of silicon switching chipsets
- Embedded operating system (OS) abstraction to run on a variety of control processors

Modern next generation terabit scale aggregation, transport and data center networks power businesses through a new generation of business critical, time sensitive high bandwidth applications and services.

Software defined networking, data center virtualization, Cloud computing, High-performance computing, Data warehousing, and Disaster recovery strategies, among others prevalent in the current environment, are prompting a whole new set of requirements for network infrastructure.

ZebOS®-XP (ZebOS Extended Performance) network platform is designed to meet the needs of all of today's high performance terabit scale networks and devices.

- **Resilient:** Provides critical business-class availability
- **Modular:** Capable of extending network capabilities to evolve with business needs and provide extended lifecycles
- **Portable:** For consistency across platforms
- **Secure:** Protect and preserve data and preserve data and operations
- **Flexible:** Integrate and enable new technologies
- **Scalable:** Accommodates and grows with the business and its requirements
- **Easy to use:** Reduces the amount of learning required, simplify deployment, and ease of manageability

The ZebOS-XP (extended performance) software platform provides superior scaling and performance to support current and future multiprocessor hardware platforms, helping ensure scalability well into the future.

Compared to current shipping versions for ZebOS, the ZebOS-XP has the following attributes:

- Up to 4x times route scaling advantage when subjected to rapid increases in IP sessions and total bandwidth requirements for large network topologies.
- Up to 8x ECMP route scaling for Layer 3 unicast routing protocols offering substantial increases in bandwidth by load-balancing traffic over multiple paths.
- Significantly improved convergence time for Layer 2 switching and Layer 3 unicast routing designed to provide stability and to perform well in environments with a huge number of routes with low CPU utilization.

This product bulletin introduces IP Infusion's Software Release 1.2 for ZebOS-XP Network Platform and summarizes the new features it offers.

Building on a Proven Foundation

ZebOS-XP is a highly-evolved modular operating system that builds on more than 13+ years of innovation and experience in high-performance switching and routing. ZebOS-XP has its roots in the original ZebOS operating system used worldwide in business-critical networks.

As a direct result of having been deployed and evolving from nearly a decade in the extremely critical storage area networking space, ZebOS-XP can deliver the performance, reliability, and lifecycle expected in terabit scale aggregation, transport and data center networks.

ZebOS-XP Offers Significant Performance Improvements

It features:

- System infrastructure built from ground up for scale and performance
- Support for multi-slotted chassis designs
- Hooks for supporting software-defined networking

New Feature Overview

ZebOS-XP 1.2 provides a rich and comprehensive feature set to address the high demands of mission critical networks in mobile backhaul, carrier transport, and data center networks. ZebOS-XP 1.2 supports all hardware and software supported in legacy ZebOS Software Release versions 7.10.x or earlier. This ZebOS-XP 1.2 release supports the generic combinations of operating systems, TCP/IP stacks, and processors as outlined in Table 1:

Table 1. General Requirements for ZebOS-XP®

Operating System • WRL 5.0, Kernel version 3.4.10 • Fedora 17, Kernel version 3.4.44 • VxWorks – 6.9 (Basic L2, L3 protocols only)	**Processor** • Intel x86 or x86_64, PowerPC
Compilation environment; Fedora 17: • autoconf version 2.68 • gcc version 4.7.0 • binutils version 2.20.x • glibc version 2.15 **Wind River Linux 5:** • autoconf version 2.63 • gcc version 4.6.3 • binutils version 2.19 • glibc version 2.11	**Typical Server Test Platform Configuration** • Disk: 2GB • Memory: 512MB • Processor: x86 or x86_64 • WRL 5.0 on kernel 3.4.10 with 32-bit configuration • NIC: Any VM/OS supported NIC
Special Software Requirements Linux Kernel Patches IP Infusion ships the kernel patches for MPLS Forwarder, Layer-2 Forwarder, etc. Refer to the ZebOS-XP® Network Platform Installation Guide for information on how to apply these kernel patches.	

ZebOS-XP 1.2 software adds enhancements in the following key areas:

Hardware Integration – ZebIC

- Broadcom Trident, Triumph, Katana family of Devices
- Marvell Prestera CX Family of Devices
- Support for PMC-Sierra Winpath family
- Hardware and software supported in ZebOS® Software Release ZebOS-XP 1.1

Infrastructure Enhancements

- Generic Multicore and Multi-threading Support Framework Prototype for one protocol. BGP Protocol is supported in this release.

Routing Protocol Enhancements

- PIM SMDM multicast routing protocol support
- Support for OSPF Demand Circuits
- OSPF LFA (IP FRR) Loop Free Alternate mechanism RFC 5286
- PIM Group-to-Rendezvous-Point Mapping (RFC6226)
- ECMP Support for PIM

Carrier Ethernet 2.0 Enhancements

- Comprehensive support for Carrier Ethernet 2.0 specification

MPLS Enhancements

- LDP-ECMP Support for up to 32 LDP (Label Distribution Protocol) dynamic ECMP paths.
- Explicit Null label support for MPLS VPN
- RFC 6074 based provisioning, auto-discovery and signaling in layer2 VPNs

High Availability Enhancements

- Statefule Switch Over for BGP protocol

ZebOS-XP® 1.2 has a new Network Service Module architecture and enhanced support for routing and switching protocols. The Tables 2 through 9, provide an overview of the main software features and benefits introduced in ZebOS-XP® 1.2 release.

OS Support

- WRL 5.0, Kernel version 3.4.10
- Fedora 17, Kernel version 3.4.44
- VXWorks 6.9

Table 2. Infrastructure Performance Improvement

Feature	Description
Multicore and Multi-threading Support – Prototype Implementation	As multithreaded processors and multi-core chips become the norm, architects and designers of digital systems use these approaches exploit the concurrency in a computational workload. Multi-processor, or multi-core, systems exploit concurrency to spread work around a system. As many software tasks can run at the same time as there are processors in the system. This tractability can be used to improve absolute performance, cost or power/performance. Starting with release of ZebOS-XP the protocol modules will be refactored to exploit the underlying hardware platform capabilities. In this release BGP protocol is refactored to improve scalability and performance.

Table 3. High Availability Enhancements

Feature	Description
High Availability	**BGP Support for NSR with Stateful Switchover (SSO)** The BGP Support for Nonstop Routing (NSR) with Stateful Switchover (SSO) feature enables provider edge (PE) routers to maintain Border Gateway Protocol (BGP) state with customer edge (CE) routers and ensure continuous packet forwarding during a Route Processor (RP) switchover or during a planned In-Service Software Upgrade (ISSU) for a PE router. CE routers do not need to be Nonstop Forwarding (NSF)-capable or NSF-aware to benefit from BGP NSR capabilities on PE routers. Only PE routers need to be upgraded to support BGP NSR—no CE router upgrades are required. BGP NSR with SSO, thus, enables service providers to provide the benefits NSF with the additional benefits of NSR without requiring CE routers to be upgraded to support BGP graceful restart. **BGP SSO Caveats** • This release adds SSO support for IPv4 only • Only core BGP states are supported • VPN services are not supported at this time

Table 4. MPLS Enhancements

Feature	Description
MPLS Enhancements	**Explicit Null label support for MPLS VPN** The MPLS VPN—Explicit Null Label Support with BGP IPv4 Label Session feature provides a method to advertise explicit null in a Border Gateway Protocol (BGP) label session for a carrier supporting carrier (CSC) customer edge (CE) router, as per Section 9 from RFC 4364. **Provisioning, auto-discovery and signaling in layer2 VPNs (RFC 6074)** The Virtual Private LAN Switching (VPLS) control plane is used for auto-discovery and signaling. Auto-discovery involves locating all provider edge (PE) devices that participate in a particular VPLS instance. Signaling is accomplished by configuring pseudowires for a VPLS instance. Prior to the introduction of the VPLS BGP Signaling L2VPN Inter-AS Option B feature, Label Distribution Protocol (LDP) was used for signaling and Border Gateway Protocol (BGP) was used for auto-discovery, as specified in RFC 6074. With the introduction of the VPLS BGP Signaling L2VPN Inter-AS Option B feature, the VPLS BGP Signaling L2VPN feature supports RFC 4761 by simplifying the auto-discovery and signaling of all known PE devices in a VPLS instance by using BGP for both functions. Auto-discovery is defined per VPLS instance. Internal BGP (IBGP) peers exchange update messages of the L2VPN Address Family Identifier (AFI) and the Subsequent Address Family Identifier (SAFI) numbers with L2VPN information to perform both auto-discovery and signaling, which includes the Network Layer Reachability Information (NLRI).

Table 4. MPLS Enhancements continued

Feature	Description
MPLS Enhancements	Both BGP standards (RFC 6074 and RFC 4761) for the auto-discovery protocol for VPLS use the same BGP AFI (25) and SAFI (65) but they have different Network Layer Reachability Information (NLRI) encoding, which makes them incompatible with each other. CLI configuration is needed to distinguish the two encoding types as they are mutually exclusive per neighbor. RFC 6074 provides guidelines for specifying length encoding as bits while RFC 4761 provides guidelines for specifying length encoding as bytes. **LDP-ECMP Support for up to 32 dynamic LDP (Label Distribution Protocol) ECMP paths** ECMP support feature for LDP performs load balancing for LDP based LSPs by having multiple outgoing next-hops for a given IP prefix on ingress and transit LSRs. An LSR that has multiple equal cost paths to a given IP prefix can receive an LDP label mapping for this prefix from each of the downstream next-hop peers. As the LDP implementation uses the liberal label retention mode, it retains all the labels for an IP prefix received from multiple next-hop peers. Without ECMP support for LDP, only one of these next-hop peers will be selected and installed in the forwarding plane. The algorithm used to determine the next-hop peer to be selected involves looking up the route information obtained from the RTM for this prefix and finding the first valid LDP next-hop peer (for example, the first neighbor in the RTM entry from which a label mapping was received). If, for some reason, the outgoing label to the installed next-hop is no longer valid, for example, the session to the peer is lost or the peer withdraws the label, a new valid LDP next-hop peer will be selected out of the existing next-hop peers and LDP will reprogram the forwarding plane to use the label sent by this peer. With ECMP support, all the valid LDP next-hop peers, those that sent a label mapping for a given IP prefix, will be installed in the forwarding plane. In both cases, ingress LER and transit LSR, an ingress label will be mapped to the next hops that are in the RTM and from which a valid mapping label has been received. The forwarding plane will then use an internal hashing algorithm to determine how the traffic will be distributed amongst these multiple next-hops, assigning each “flow” to a particular next-hop.

Table 5. Carrier Ethernet

Feature	Description
Carrier Ethernet 2.0 Enhancements	CE 2.0 is set of Metro Ethernet Forum Carrier Ethernet 2.0 Certified network elements that connect to transport Carrier Ethernet services for all users, locally and worldwide. Carrier Ethernet services are carried over physical Ethernet networks and other legacy transport technologies. CE 2.0 greatly expands from 3 services in CE 1.0 to 8 services, 2 of each respectively in E-Line, E-LAN, E-Tree, and E-Access (defined in MEF Standards MEF 6.1, 33) The industry’s first standardized Multi-CoS with application-oriented CoS Performance Objectives, new metrics (MEF 6.1, 10.2, 20, 23.1) Interconnect through the integrated delivery of MEF Service Attributes (MEF 10.2, 26.1, 33) allows ubiquitous deployment spanning multiple providers Manageability, (MEF 7.1, 16, 17, 30, 31) plus additional specifications In this release, the existing ZebOS-XP Carrier Ethernet support will be enhanced to comply with the latest CE 2.0 specifications with focus on the following area: • Carrier Ethernet 2.0 E-Line, E-LAN, E-Tree, and E-Access, as defined in MEF Service Specifications and Implementation Agreements (MEF 6.1, 6.1.1, 33 & 26.1)

Table 6. Routing Enhancements

Feature	Description
Multicast Routing Enhancements	**PIM sparse-dense mode** Protocol-Independent Multicast (PIM) 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. It is termed protocol-independent because PIM does not include its own topology discovery mechanism, but instead uses routing information supplied by other routing protocols. The main variants of PIM are PIM Sparse Mode (PIM-SM) and PIM Dense Mode (PIM-DM). PIM-SM initially uses a shared tree that flows from a shared root (multicast router) in the network known as a rendezvous point (RP) and also uses a shortest path tree (SPT) to allow each source to forward traffic to the rendezvous point. The rendezvous point then receives multicast traffic from the source via the SPT and then forwards the multicast traffic down the shared tree to the various receivers. PIM Dense Mode (PIM-DM) uses dense multicast routing. PIM-DM implicitly builds shortest-path trees by flooding multicast traffic domain wide, and then pruning back branches of the tree where no receivers are present. PIM-DM generally has poor scaling properties. PIM sparse-dense mode represents a simple extension to PIM sparse mode operation that enables PIM routers to use PIM dense mode operation for any multicast groups that do not have an RP configured. Sparse-dense mode, as the name implies, allows the interface to operate on a per-group basis in either sparse or dense mode. A group specified as dense is not mapped to a rendezvous point (RP). Instead, data packets destined for that group are forwarded by means of PIM dense-mode rules. A group specified as sparse is mapped to an RP, and data packets are forwarded by means of PIM sparse-mode rules. This is in addition to the sparse or dense mode of operation already supported in the currently shipping versions of ZebOS-XP. **PIM Group-to-Rendezvous-Point Mapping (RFC6226)** Each PIM-SM router in a PIM domain that supports Any Source Multicast (ASM) maintains Group-to-RP mappings that are used to identify a Rendezvous Point (RP) for a specific multicast group. PIM-SM has defined an algorithm to choose a RP from the Group-to-RP mappings learned using multiple mechanisms. This algorithm does not consider the PIM mode and the mechanism through which a Group-to-RP mapping was learned. RFC 6226 defines a standard algorithm to deterministically choose between several Group-to-RP mappings for a specific group. This feature adds support for RFC 6226 in ZebOS-XP.
Unicast Routing Enhancements	**RFC 1793 (OSPF on demand)** This ZebOS-XP feature adds support for RFC 1793 which defines enhancements to the OSPF protocol that allow efficient operation over “demand circuits”. Demand circuits are network segments whose costs vary with usage; charges can be based both on connect time and on bytes/packets transmitted. **OSPF LFA (IP FRR) Loop Free Alternate mechanism RFC 5286** This feature adds support LFA support for OSPFv3 module in ZebOS-XP, as defined in RFC 5286. The OSPFv3 Loop-Free Alternate Fast Reroute feature uses a pre-computed alternate next hop to reduce failure reaction time when the primary next hop fails. It lets a network administrator to configure a per-prefix loop-free alternate (LFA) path that redirects traffic to a next hop other than the primary neighbor. The forwarding decision is made and service is restored without other routers’ knowledge of the failure.

Table 7. Quality Improvements

Feature	Description
Software Quality Improvements	**I. Software Quality Improvements** This release also addresses enhancement requests and defects reported by IP Infusion customer base and QA team in ZebOS-XP 1.1 GA release and releases prior to that. Please refer the ZebOS-XP 1.2 release notes for a detailed description of the defect fixes incorporated in ZebOS XP 1.2.

Please contact us to learn more about our ZebOS® Platform

Phone: +1 877-MYZEBOS
Email: sales@ipinfusion.com
Web: www.ipinfusion.com

U.S. (Santa Clara), +1 408-400-1900
Japan (Tokyo), +81 6661-9231
Korea (Seoul) +82 (2) 3153-5224

India (Bangalore), +91 (80) 6728 7000
China (Shanghai), +86 186 1658-6466
EMEA (Stockholm), +46 8-566 300 42

About IP Infusion

IP Infusion is a leading provider of intelligent network software for enhanced Ethernet and IP services. Tier one and two OEMs rely on IP Infusion's ZebOS software and global professional services to bring products to market faster, and to differentiate them from competitors with less cost. Products built on IP Infusion technology are deployed in networks with five-9s reliability across five continents—as well as a growing number of enterprises—to improve network performance, decrease network infrastructure costs, and grow revenue. IP Infusion is headquartered in Santa Clara, Calif., and is a wholly owned and independently operated subsidiary of ACCESS CO., LTD., of Tokyo, Japan.

IP Infusion
An ACCESS Company
(408) 400-3000
www.ipinfusion.com
3963 Freedom Circle, Suite 200, Santa Clara, CA 95054