

Xelerated and IP Infusion Reference System

Key Features

- ◆ Xelerated's Xelerator™ SDK 1.2
- ◆ Xelerated's X10 forwarding-plane applications
- ◆ Xelerated's Xelerator™ X10 Reference Board Design Kit
- ◆ Xelerated's X10 Middleware
- ◆ IP Infusion's ZebOS™ Advanced Routing Suite
- ◆ FreeBSD™ 4.4

Architecture Overview

This application brief describes a reference system that integrates Xelerated's X10 family of network processors (NPU) with IP Infusion's ZebOS Advanced Routing Suite (ARS) on a FreeBSD 4.4 platform. The solution enables network equipment OEMs to build full router systems at lowest possible cost by exploiting the already verified interworking between Xelerated's NPUs and IP Infusion's routing software suite.

The reference system includes IP Infusion's IP routing suites for IPv6 and IPv4 with OSPFv2/v3, RIPv1/v2/ng, IS-IS, BGP4+ as well as MPLS with RSVP-TE, CR-LDP and MPLS VPN. The implementation also features Xelerated's tested and verified forwarding-plane applications, which is the code that runs on the X10 NPU.

Xelerated's X10 Middleware is responsible for the control and management of the network processor including management of the Forwarding Information Base (FIB). The ZebOS ARS, which is based on IP Infusion's Network Service Module (NSM), provides management of the Routing Information Base (RIB), route conversion and distribution, industry standard command line interface as well as management of the different Management Information Bases (MIBs). The reference system includes the X10 middleware that is linked with the ZebOS ARS executables. The X10 middleware handles all control traffic between the control and the forwarding plane. The protocol modules OSPFv2, BGP-4 and RIP run on top of the NSM.

The forwarding plane is implemented using Xelerated's X10 reference board as a target. For debugging and demonstration purposes, the X10 simulator could also be used as an alternative.

X10 Forwarding-plane Application

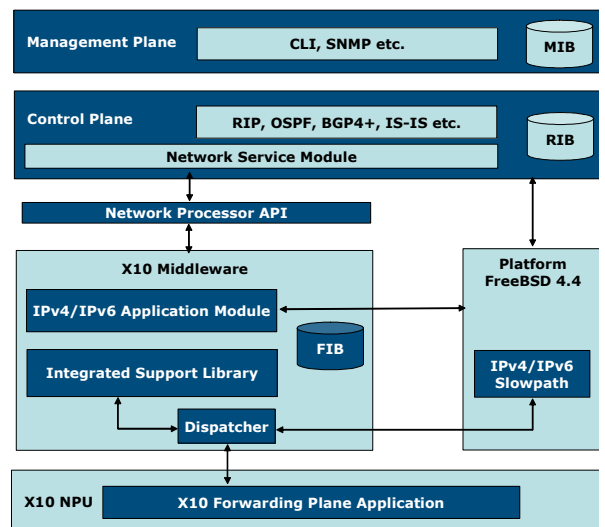
The forwarding-plane application is the code that runs on the network processor. The forwarding-plane code is assembled, linked and debugged using the Xelerator™ Software Development Kit. Today, an extensive set of robust and deployable applications is already available covering IP forwarding (IPv4/v6, ACL, multicast), DiffServ, counting and statistics (RMON, SMON, MIB2), L2 bridg-

ing (802.1Q, 802.1D, 802.3ad), MPLS (LER, LSR), ECMP, VPN (VPLS, Martini), NATP as well as IPv4 to IPv6 transition mechanisms. The forwarding-plane applications are verified and tested by Xelerated.

X10 Middleware

The middleware has the following parts:

- ◆ The *IPv4 and IPv6 application module*. The application module is the control-plane part of the IPv4/IPv6 forwarding-plane application. The forwarding-plane application interacts with all the resources connected to the X10 NPU like TCAM engines, SRAM, co-processors, counters, meters. The application module is responsible for the configuration and management of the X10 by using special headers to carry the information up to the forwarding plane. To make the integration easy, Xelerated provides common data structures that are shared between the control plane and the forwarding plane. The data structures are stored in ordinary header files, which are used by both the forwarding-plane and the control plane.
- ◆ The *Integrated Support Library (ISL)*. ISL contains support functions to access the X10 and its resources connected to it like TCAM, SRAM, CPU and co-processors. The ISL prepares and generates the control messages to configure X10. It keeps a full copy of the different tables, like Forwarding Information Base (FIB) and Access Control Lists (ACLs), to be able to sort the tables according to different policies such as longest-prefix match for the routing table. After the table has been sorted, it is distributed down to the forwarding plane to minimize the traffic between the control plane and the forwarding plane.



Overview of IP Infusion and Xelerated Reference system

- The *adaptation glue* (not shown in the picture). The adaptation glue sets up the dispatcher, opens the routing socket to be able to listen to the slow path, manages the network interface list, the target configuration, and handles all events. The largest part of the code, however, is for parsing and processing messages from the routing socket (slow path), which results in calls to the application module. The dispatcher inspects the packet arriving from the forwarding plane, strips the encapsulation, and either identifies the packet as a control or data packet. Control packets are dispatched to the ISL, while data is dispatched to slow path in FreeBSD.

The X10 middleware is implemented as a Xelerator™ module in ZebOS ARS. The code is linked together with the ZebOS ARS executables, and it is also configured when ZebOS ARS starts. For some Xelerated-specific configurations, the ZebOS CLI has been extended.

Network Processor API

The network processor API is NP forum compliant. The NP forum API is based on an asynchronous model where functions return values that are passed back to the caller as asynchronous events. A large part of the code of these functions deals with event handling of the asynchronous API.

FreeBSD 4.4 Platform

The reference system uses a standard Free BSD 4.4 as a platform for slow path. A BSD platform was chosen since the source code is open and freely available, which makes it possible to modify and reconfigure easily, which is an advantage in prototypes. The reference system is portable and could easily be changed to other OSs like VxWorks, OSE-Delta as the customer sees fit.

In BSD, the system is divided into a kernel part and an application/user part. The routing (ZebOS ARS) and management protocols run in user space while the device drivers and IP-stack (slow-path) run in kernel space. The routing tables are maintained by ARS as RIBs by the routing protocols, but the forwarding entries are inserted via a routing socket into the FIB.

Slow-path

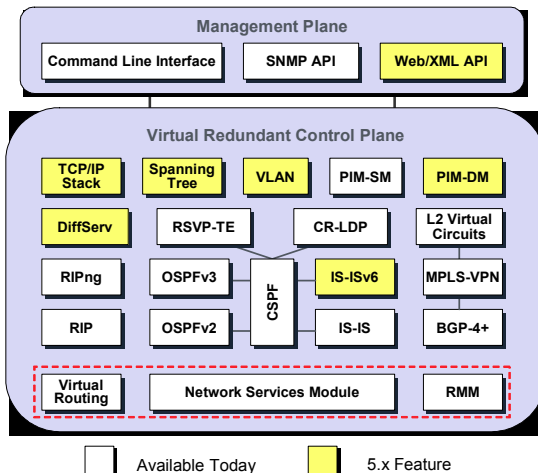
The slow-path uses the IPv4, IPv6 and ARP protocol stacks implemented in the FreeBSD kernel. Packets destined for slow-path processing are sent from the forwarding plane as UDP/IP packets encapsulated within the X10 CPU protocol.

Data packets are sent to the kernel from the dispatcher over a virtual interface. Such a device mimics an ethernet device, but lacks hardware. Instead, a process from user space (in this case the dispatcher) opens a character device to the interface and exchanges data with it, just as the low level part of an ethernet driver would. The dispatcher uses regular open/read/write system calls for this. The operating system, the user space and the IPv4 stack regard the tap interface as any regular ethernet interface.

The IPv4/IPv6/ARP stack processes the slow-path packets. It may forward the packets, process options, or deliver the packets to a user process, such as a ZebOS OSPF daemon.

Control and Management Plane

The control plane hosts the application-level protocols. In this case, IP Infusion's ZebOS Advanced Routing Suite provides the IPv4 and IPv6 dynamic routing protocols, RIP-v1/v2/ng, OSPFv2/v3, IS-IS and BGP-4+. The control plane builds and maintains the Routing Information Base that is later propagated down to the Forwarding information Base in the X10 middleware.



Overview of ZebOS ARS

The management plane consists of CLIs for configuration and maintenance, and an SNMP agent. ZebOS ARS uses several CLIs, one for each routing protocol daemon, and one common ZebOS CLI. With ZebOS CLI, configuration of static routes, interfaces, set debug options can be done. The ZebOS CLI has been extended with some Xelerated-specific commands. These extensions include commands for simulator line-card configuration and debugging.

Supported Targets

The reference system of IP Infusion's ARS and Xelerated's Xelerator™ X10 family of NPUs and associated software runs on several targets, such as Xelerated's clock-cycle accurate simulator that is part of Xelerator™ Software Development Kit. It can also run on the X10 Reference Board that supports 4 x 10 Gigabit Ethernet or 4 x OC 192 Packet over SONET ports. Debugging of software can be done on both targets; control-plane and forwarding-plane integration can also be tested on both targets.



Xelerated's X10 Reference Board

Integrated Development Environment

Xelerator's Integrated Development Environment (IDE) for developing forwarding-plane applications includes: compiling, clock-cycle accurate simulation, and debugging. The IDE is project oriented and covers all the stages of a project: from the early stage of defining the applications to be developed, to the final stage of delivering production-ready applications for your switch or router.

A project starts with an intuitive application wizard guiding you through the initial steps of defining the applications to be implemented, resulting in source files, header files, configuration files and a make environment. Within your project, you can easily handle resources such as: test data packets; control application engine entries and external look-aside engine co-simulation models. In the context of your project, you then implement, debug and simulate your applications until they reach production-ready status.

For maximum flexibility and ease of use, the debugger supports both a graphical user interface (GUI) and a command line interface (CLI). To increase the concurrency of hardware and software development, the debugger and simulator support three modes of operation: stand-alone debugging, control-plane co-simulation and hardware debugging.

A unique feature of the Xelerator NPU's data flow architecture is that all compiled code is guaranteed to run at wire speed, so there is no need for a traditional performance analyzer. This removes all temporal dependencies from the code, eliminating the time consuming and complex debugging of asynchronous interaction problems as well as the time consuming process of minimizing load/store stalls, branch delay slots and synchronization stalls. The result is a significant reduction in development cost.

Pre-tested solution

The solution is verified and tested on the reference system. The solution provides an extensive set of robust and deployable applications. Conclusion is that the demonstrated reference system shortens the time-to-market for system vendors.

Product Features Details

Xelerator™ NPUs

- X10s—a single 10 Gbps NPU with 4 SPI-4.2 ports that supports 25 Mpps.
- X10d—a dual 10 Gbps NPU with 4 SPI-4.2 ports that supports 50 Mpps.
- X10q—a quad 10 Gbps NPU with 4 SPI-4.2 ports that supports 100 Mpps.

Xelerator™ Software Development Kit

- Clock-cycle-accurate simulator
- Debugger
- Assembler
- IDE with GUI
- Linker
- Configuration Wizard
- Integration Support Library

Xelerator™ X10 Reference Design Kit

- 2 x 10 Gbps Ethernet / 2 x OC-192 Packet over SONET ports
- 2 ports to a modular adapter board interface
- Board Support Package
- PCB Design Package

Xelerator™ forwarding-plane applications

Layer 2

- IEEE 802.1d/q VLAN tagging and bridging
- IEEE 802.3ad Link Aggregation
- RFC894, A std for transmission of IP datagrams over Ethernet Networks
- RFC1042, A std for transmission of IP datagrams over IEEE 802 Networks
- RFC 2615, PPP over SONET/SDH

IPv4

- RFC1812, Requirements for IP version 4 routers
- RFC2991 Multipath Issues in Unicast and Multicast Next Hop Selection

IPv6

- RFC1981 Path MTU discovery for IPv6
- RFC2373 IPv6 addressing architecture
- RFC2460 IPv6 specification
- RFC2461 Neighbor discovery for IPv6
- RFC2473 Generic packet tunneling in IPv6
- RFC2545 Use of BGP-4 multiprotocol extensions for IPv6
- RFC2893 Transition mechanisms for IPv6 hosts and routers
- RFC3056 Connection of IPv6 domains via IPv4 clouds
- IRFC3306 Unicast-prefix based IPv6 multicast addresses
- draft-ietf-ngtrans-isatap-04.txt Intra-site automatic tunnel addressing protocol

MPLS

- RFC2547bis BGP/MPLS VPNs
- RFC2702 Requirements for traffic engineering over MPLS
- RFC3031 Multiprotocol label switch architecture
- RFC3032 MPLS stack encoding
- RFC3270 MPLS support of differentiated services
- Draft-martini-l2circuit-encap-mpls-04.txt
- Draft-martini-l2circuit-trans-mpls-08.txt

DiffServ

- RFC2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
- RFC2475 An Architecture for Differentiated Services
- RFC3246 An Expedited Forwarding PHB
- RFC2597 Assured Forwarding PHB Group
- RFC2983 Differentiated Services and Tunnels
- RFC3086 Definition of Differentiated Services Per Domain Behaviors and Rules for their Specification

Access Control Lists

- L2/L3/L4 Access Control Lists

ZebOS Advanced Routing Suite

- ZebOS Base Module: Network Services Module

- ZebOS Protocol Modules: BGP4+ Protocol Module, OSPFv2 Protocol Module, RIP Protocol Module, OSPF-CSPF Protocol Module and IS-IS Protocol Module
- IPv6 Routing: BGP4+ Protocol Module, OSPFv3 Protocol Module and RIPng Protocol Module
- ZebOS Multicast Module: PIM-SM Module
- ZebOS High Availability and Virtual Routing Modules: Virtual Routing Module and Routing Redundancy Module
- ZebOS MPLS Switching Modules: RSVP-TE Module, CR-LDP Module, LDP Module, MPLS Layer 2 Virtual Circuit Module and MPLS Forward Module

Address and contact information

IP Infusion Inc.

111 W. St. John Street
San Jose, CA 95113
Phone 408 794.1500
Fax 408 278 0521

Xelerated Inc.

200 Wheeler Road
Burlington, MA 01803
Phone: 781 505 1921
Fax 781 505 1382



sales@xelerated.com

US east coast: (781) 526 7896

US west coast: (408) 371 3851

www.xelerated.com

Xelerated makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising as a result of estoppels or otherwise, to any intellectual property rights is granted by this publication. Xelerated, the Xelerated logo, Xelerator, PISC and combinations thereof are trademarks of Xelerated. IP Infusion, ZebOS, and ipinfusion are trademarks of IP Infusion Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies. Copyright © 2002 Xelerated and IP Infusion Inc. All rights reserved.