Intel IXP1200 Integrated Router Application

[Application]

This application consists of: a test network; a Cisco* 3600 router; an Intel* IXDP1200 Advanced Development Platform configured as an IPv4 OSPF2 router; a video feed server; and a video feed client. The video feed is routed through both routers. In the presence of a link failure, the video steam is automatically rerouted from the primary to the standby route, demonstrating the dynamic and fault-resilient properties of this routing implementation, and the rapid convergence of the OSPF route selection algorithms. The Intel Pentium* CompactPCI card in the network processor development platform is running IP Infusion's ZebOS Advanced Routing Suite (ARS). Teja NP's IPv4 Foundation Application provides the complete forwarding plane packet processing logic that executes on the Intel IXP1200 network processor.

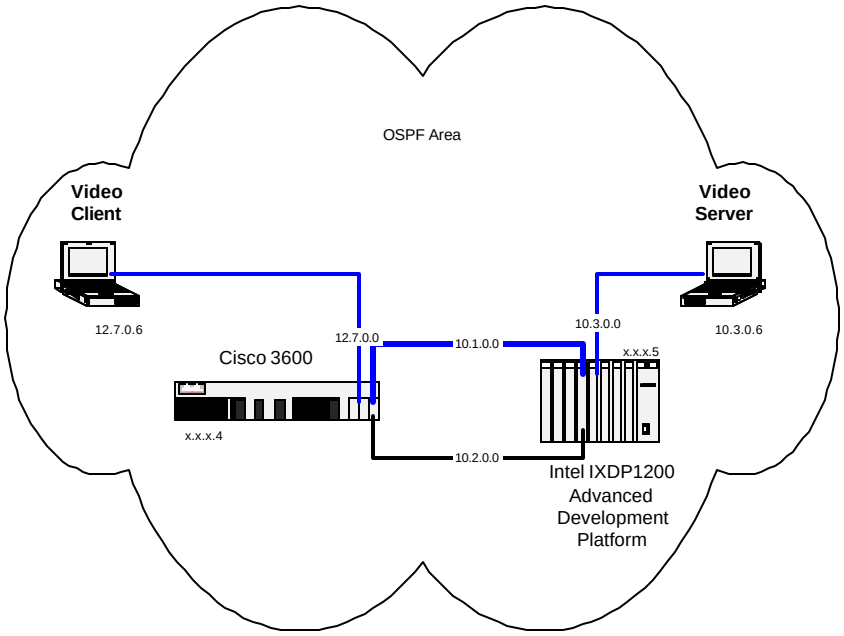

This virtual router architecture defines both a control and a forwarding plane, and in an product implementation would also define a management plane. The forwarding plane is comprised of the six IXP1200 microengines of the IXPM1200 CompactPCI blade that does the 'fast-path' processing and forwarding of packets at line rate (according to RFC 1812). The control plane consists of the complete virtual router application running under Linux on the Pentium III* CPU of the CompactPCI system controller blade, and backplane connected control plane agent and forwarding table manager running under Linux on the IXP1200 StrongARM* core. Routing updates and exception packets are sent from the forwarding plane to the control plane. The control plane computes routes, manages the routing protocol messages, and maintains a route table. The route table is used to create and update the forwarding tables that provide 'next-hop' lookup information for the forwarding plane. For performance and to accommodate Classless Inter-Domain Routing (CIDR), the forwarding decision uses a Longest Prefix Match (LPM) process operating on a unique two-table forwarding table structure. This multiple blade architecture closely parallels actual hardware designs and the NPF and IETF ForCES efforts.

February 2002

Equipment developers want to be able to develop unique, custom products for a range of markets. They need performance and portability, and they need to quickly integrate protocol into their platforms. The challenge for developers is integrating components for their solutions together to create a working, quality system. Teja and IP Infusion have worked together to create a pre-integrated, fully tested solution for IXP1200 network processor software development and routing solutions. This network processor application, Teja NP software platform and integrated IP Infusion Advanced Routing Suite is optimized, fully customizable, portable.

[Solution Features]

- Teja NP 1.0 for Intel IXP1200 Network Processors
- Teja NP IPv4 Forwarding Foundation Building Blocks
- IP Infusion Advanced Routing Suite
- IP Infusion IPv4 OSPF
- IP Infusion Teja NP Adaptation Layer
- Intel IXA Software Development Kit
- Intel IXDP 1200 Advanced Development Platform
- Intel IXD4501 LAN I/O Card
- Cisco 3600 Router

[Technical Discussion]

Teja NP Layer3 Forwarding Foundation Application is a complete RFC1812-compliant IP Forwarding application using two full-duplex Gigabit Ethernet (GbE) ports, or one GbE and eight 10/100 Mbps Fast Ethernet ports. This foundation application includes forwarding-related control plane elements: a Forwarding Table Manager; remote Route Manager, and a CPDP Agent. The Route Manager services packets to and from the data plane queues and manages route updates from the routing protocols. The Route Manager sends the route updates to the Forwarding Table Manager through Teja's IPv4 Services API. The Forwarding Table Manager uses the Route Manager updates, as well as updates from Teja's ARP tap (which maps IP addresses to MAC addresses) to maintain the Forwarding Table and Forwarding Route Tree entries.

The IP Infusion Advance Routing Suite (ARS) is a complete software implementation of a virtual router that includes the IPv4 OSPF, BGP and RIP protocols. ARS integrates with the control plane via Teja's IPv4 Services API for management of connections, interfaces, route updates, exceptions and statistics. Teja's IPv4 Services API conforms to the definition and capabilities of the Network Processing Forum IPv4 API, as documented by the current work in progress.

Teja IPv4 Services API

Teja offers an API for IPv4 services that is modeled on the emerging Network Processing Forum IPv4 API. The API provides an asynchronous interface for updating routing table information, individual port state, and low-level packet reception and transmission. The API supports multiple line cards connected to a single host.

Modifications to IP Infusion Advanced Routing Suite (ARS)

IP Infusion has created a Teja NP Adaptation Layer that calls the Teja APIs from within the routing engine core to update the data plane's forwarding table. This shim was developed and integrated in a very short time, attesting to the completeness and ease of programming to the Teja IPv4 API.

Packet Flow

Received Packets

A packet will enter the system via one of the ports serviced by the data-plane components. The data-plane component will make a decision on how to dispose of the packet. It may choose to forward it via another port on the

same line card, in which case the host network stack will not be exposed to the packet. The data-plane code may also discover that the packet is a routing update, or control packet destined for the control plane entity. In this case, the packet is forwarded to the control plane as a Teja alert, possibly activating some customer-supplied code.

The control plane will send the packet to the CPDP agent component, which will send the packet (via a Linux socket call) to the CPDP client application. The client application will decode the metadata it receives (primarily received port number) and hand the packet to the stack adapter for injection into the network stack of the control processor. From there, the host stack will process the packet as usual, either by forwarding the packet to another line card or by handling it as a locally bound packet.

If the packet is a routing protocol packet, it will be handed to ARS, which will process it and make decisions about routing policy. The line card control plane is informed of those decisions via the Teja IPv4 services API.

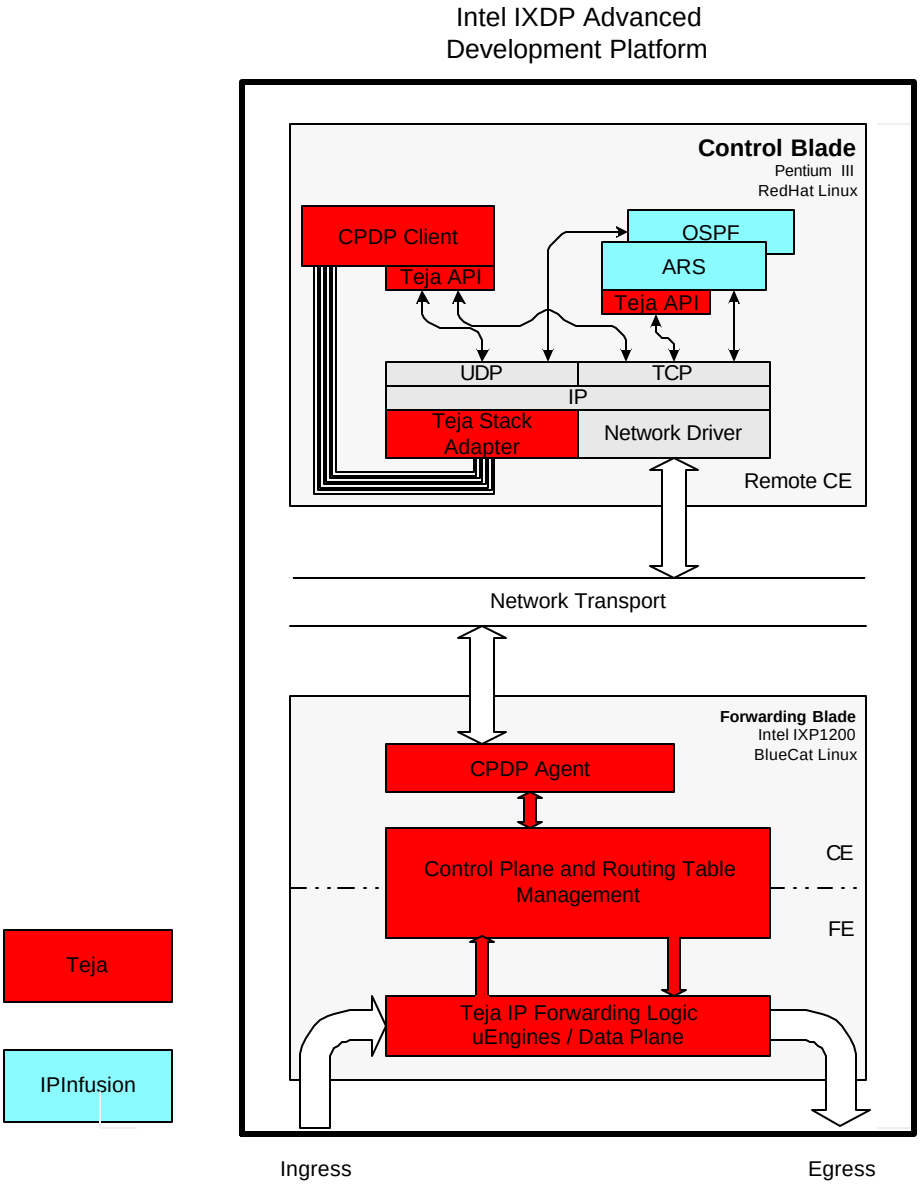

Transmitted Packets

A data packet sent from ARS will be received by the host IP stack, and the routing functions inside the stack will select one of the interfaces presented by the stack adapter as the destination interface. The packet will be passed to the stack adapter, which will send the packet to the CPDP client. The CPDP client will forward the packet to the CPDP agent, which sends the packet as an alert to a component in the Teja NP Layer 3 Forwarding application.

Compatibility

Legacy applications will not need to be aware of the presence of the stack adapter. The adapter will simply appear as another network interface for the control plane operating system, and the operating system's stack will provide ARP, multicast handling, and transmit packet routing to the applications. Applications that need to make changes to the routing base will need to use the Teja IPv4 services API, in addition to the host OS interface.

Routing table updates can be ordered to produce a "edge-to-middle" update simply by applying the host kernel routing table updates in the completion callback from the Teja API. This will minimize the probability of the host routing table and the line card routing table being out of synchronization.

[Results]

The integration of this demonstration application was accomplished in less than one day by members of the technical staff of IP Infusion and Teja Technologies.

[Conclusions]

IP Infusion has worked together with Teja to integrate its IPv4 OSPFv2, BGP-4, and RIPv1/v2 routing protocols into the Teja control plane software. This integration was done in less than a day, using the Teja control to data plane API and was tested using a variety of standard testing tools. This means that developers that are using the Teja Development Environment can quickly, easily, and seamlessly add pre-integrated routing components to their architecture without the cost normally associated with protocol and NP development.

Teja and IP Infusion are working together to integrate and optimize additional protocols into the Teja environment, as well as to promote standards through the Network Processing Forum (NPF).

[Other Resources]

- Teja NP IXP1200 Product Brief
- Teja NP Layer 3 Forwarding Foundation App Product Brief
- ZebOS Advanced Routing Suite Datasheet
- Teja IP Forwarding White Paper
- Teja NP 1.0 IP Forwarding Design Overview documentation
- Teja Technologies Inc, 2 W. Santa Clara St, 6th Floor, San Jose, CA +1 (408) 288-2560
- IP Infusion, 111 W. St. John Street, Suite 910, San Jose, CA +1 (408) 794-1500
- Network Processing Forum
- IETF Forwarding and Control Element Separation (ForCES) Working Group

[7] Teja and Teja NP are trademarks of Teja Technologies, Inc. IP Infusion, ipinfusion and ZebOS are trademarks of IP Infusion Inc. All other brands or product names are trademarks or registered trademarks of their respective holders.