



Metro Ethernet

A New Value Proposition
September 2006



EXECUTIVE SUMMARY

Network infrastructure has begun another cycle of evolution as Service Providers look to Ethernet technology as a viable method of providing both point-to-point and multipoint services over high-bandwidth links at the network edge. These services are necessary for the retention of an existing subscriber base in the face of service commoditization, and the capture of new subscribers from the competition.

The scalability and flexibility to deliver broadband services economically by adopting Ethernet as a wide area transport technology gives Service Providers the ability to offer higher-revenue services such as Video (Video-on-Demand as well as Streaming Video) and Voice over the Internet Protocol [VoIP]. Additionally, the abundance of new fiber deployment to business areas has made Ethernet the logical choice for deployment.

IP Infusion delivers proven technologies in a L2, L3 or L2/L3 implementation that guarantee Return on Investment, speeding a time-to-market and providing a robust feature set that ensures the greatest flexibility in addressing some or all of the new Metropolitan Ethernet requirements.

MARKET DRIVERS

Telecommunication's markets are moving further from the core of the network toward the edge, and out of the Enterprise. With the standardization on Ethernet for the last mile (to the premises), Layer 3 routing protocols are losing ground and importance to Layer 2 transports that provide for low-overhead and prioritized transport of traffic. This is particularly important as network topologies and services converge onto multi-purpose infrastructure delivering all traffic.

Carriers are deploying IP Next Generation Networks to deliver triple play services (i.e.: Voice, Video and Data [VVD]), including video-based services like IPTV. Carrier success in delivering the 'triple-play' of voice, video, and data is not only dependent upon the proper choice of service and content partners but also on the right network infrastructure. This network infrastructure must be capable of evolving as business and consumer needs change, as new services and applications are introduced into the marketplace, and as bandwidth needs grow. Video may play a small role in IP traffic growth immediately, but will figure prominently in the long term.

The initial video applications will be broadcast and on-demand video. Other growing demand applications will be video-based, as well. However the table stakes service is broadcast video and subsequently it has the highest level of expectation from the viewer -- therefore it will be necessary to protect the integrity of this service while seamlessly integrating it with existing voice and data services and introducing additional services.

This rise in broadcast video, video-on-demand and other video-enabled services will place increased demands on provider infrastructures. Traditional wire-line IP architectures are not robust or resilient enough to deliver video and support the viewing experience that is the expectation of the current consumer market. Providers are actively investigating the requirements to deliver video over this all-IP infrastructure. To adequately address the challenges imposed, it will be necessary to provide for high quality services protected from disruption and to do this while scaling to bandwidth in terms of terabits per second.

Only IP Next Generation Networks incorporate all the elements, such as Quality of Service, secure virtualization, massively scalable multicast transport and continuous systems operation which will provide the foundation necessary to offer the high-end, high demand video-based services providers will require going forward. The first area of the network that

will be forced to meet these requirements is the edge, and the requirement will force the use of Ethernet to adequately carry the volume of traffic, at the speeds required, with the types of service guarantees that will provide for the viewing experience expected by the end user. The next area of the network that will be expected to adhere to these increased specifications is the wireless segment responsible for delivering (today) voice, data and audio to the consumer. Already, video is being distributed on a store-and-forward basis to handheld devices capable of full video playback. Proof-of-concept delivery of video streams in real-time is being explored by some telephony providers.

While current wireless infrastructure is sufficient to deliver video, work is underway on more efficient, leaner implementations of transport protocols (e.g.: Digital Living Network Alliance [DLNA]). These protocols will have to be hardened as they develop to provide for the lossless transport that is expected of video streams. In addition, opportunity exists in the securing of such streams to prevent the unauthorized appropriation of user accounts.

Finally, as full-function network services are delivered over the same infrastructure to the consumer, there must be a device at the end that is responsible for the termination of the streams, and parceling out of the individual services to the appropriate device. The requirement for this device will be more than the termination of services, however. It will be expected to be a repository that can send as well as receive, exporting personal multimedia data to the owner regardless of remote location. In that role, it will be required to initiate sessions across the infrastructure, provisioning connections to the receiving endpoint across the global network.

CHARACTERISTICS OF METRO-E

The Metro Ethernet market is the network boundary at which the bulk of the implementation for the delivery of NGN services will occur. As noted, traditional wire-line networks do not have the capacity or the transport layers to adequately handle the demands posed by streaming audio or video. Newer wireless networks do not yet have the capacity or the resiliency to deliver streaming audio or video.

Metro Ethernet, based on proven technologies, provides the capacities (predominantly one-hundred or even one-thousand megabit physical layers) and the resiliency (using the established IP transports for carrying traffic) to deliver timing and loss-sensitive data such as streaming audio or video.

Devices participating in the MetroE framework are purpose built (i.e.: a routing/switching platform that has as its primary role subscriber authorization and authentication, application of security and session aggregation prior to their delivery to the core of the network. The platform may participate in the Multi-Protocol Label Switching [MPLS] topology of the network core, but only as a Label Edge Router [LER] for a subset of the global Label Switched Path [LSP] domain.)

These devices carry a software mixture that provides the Quality of Service, secure virtualization, massively scalable multicast transport and continuous systems operations required for NGN implementations:

Quality of Service

Quality of Service [QoS] has traditionally referred to mechanisms intended to maximize the efficiencies with which network resources are allocated. Those sessions deserving greater resources get those resources when needed and lower-class sessions are throttled to prevent resource monopolization.

QoS now assumes a greater role in the network with the advent of streaming services. It is not enough that network resources are allocated according to design (and assumedly need) but also according to type of traffic and on a dynamic basis. NGN data streams require disproportionately large amounts of network resources. It does not benefit the network to reserve those resources when they are not required. Furthermore, resource intensive traffic cannot be allowed to impact other traffic (e.g.: voice) during failover situations.

Therefore, QoS is used to dynamically balance the network, allocating resources as needed, and *only when needed*, while maintaining the traditional class structure as defined in early QoS implementations.

During failover or recovery situations, network resources are guaranteed to critical traffic while alternate paths are chosen for traffic being rerouted. Non-priority ('best-effort') traffic gives up resources to accommodate priority traffic, but existing priority streams continue without interruption, meeting Service Level Agreements [SLAs].

Guaranteed delivery

Given the nature of streaming audio and video traffic, packets comprising the data stream must be delivered, on time and in sequence. Delays in propagation through the network result in jerky or interrupted display at the receiving device, degrading the multimedia experience.

As the transmission is layered upon IP protocols as the transport, if and when there is a missing packet, the IP retry mechanism will be invoked, resulting in [at best] an out-of-sequence packet requiring re-ordering by the receiver, but little or no disruption to the stream. At worst, the resent packet arrives out-of-sequence and after the succeeding packet has been presented, causing disruption.

During fail-over situations, there is a greater risk of packets becoming unsequenced as streams are disrupted, resulting in packet loss as well as introduced latency as streams are rerouted.

Loss-less transport

Streaming traffic is intolerant of packet loss (packet drop). Because IPTV is highly compressed, losing a video packet can result in the loss of valuable encoded information and a visible degradation of video quality, including macroblocking, pixelation, and even loss of a picture frame.

Industry norms have defined the acceptable video quality of experience to be no more than one visible blemish per 2-hour movie. The corresponding network QoS for this must result in an allowed packet loss rate of approximately 1 in 1 million (10⁻⁶). For service providers delivering IPTV, this 10⁻⁶ maximum loss is considered a baseline requirement in the market.

Session Security

With providers depending on the NGN services to deliver new revenues as traditional services become commoditized, the first consideration that will have to be taken into account is security of the network, traffic on the network, and individual sessions across the network.

Security is not limited to preventing unauthorized access to network services or resources.

Security also implies *guaranteeing the integrity* of a stream (or streams) such that in a dynamic network stream (or class) characteristics are not impacted by loss of resources.

Network events cannot be permitted to have negative effects upon existing traffic. Streams directly impacted must be rerouted around failures without loss, and without causing unaffected streams to fall below critical resource levels.

Voice traffic transiting stable paths cannot be allowed to be impacted by video traffic taking an alternate path around a failure and consuming resources on the voice path. The rerouting

of traffic must take into account the available network resources on the alternate path, and not exceed minimum requirements unless there is sufficient excess bandwidth.

Port-based

Port-based security provides for the integrity of the network. Only authorized stations may connect, and when connected have their sessions rigorously set up according to restrictions on where they may go, what services they may access, and with what network resources.

AAA

Authorization, Authentication and Accounting [AAA] overlap physical-layer security (i.e.: only authorized users get access to network resources) but extend further into the delivery mechanisms for control of network access, service delivery, guarantee of services and logging of usage history.

Using AAA providers can guarantee not only the security of the network, but also that users get exactly the level of services for which they contract (minimizing service theft and excess resource consumption) and premium users receive service levels designed to deliver the greatest consumer experience possible.

Virtualization

With physical and logical network topologies converging, there is a growing requirement to provide for 'security' within a platform as well as at the platform UNI.

Virtual Routers are ideal for segregating not only routing domains, but "application domains" as well. Voice, video and data traffic can be segregated, limiting exposure to disruption from failure in external processes. This 'compartmentalization' of services minimizes collateral damage to revenue generating services in failure conditions.

Virtualization by services maximizes the revenue generation ability of the platform by increasing resiliency. Further virtualization by individual multicast and voice domains not only will enhance resiliency, but account security as well.

Scalability

Scalability is becoming a greater issue, as end-users and businesses grow in number and service utilization. Services previously limited to 'hard-wired' end-stations are being extended to stations connected via a number of wireless technologies such as GSM or WiFi/WiMAX. New services are being introduced such as the Simple Messaging Service [SMS] that result in vastly increased amounts of traffic and greater frequencies of transmission between participating end-stations.

Traffic volumes and rates are now multiplying exponentially as traditionally 'dumb' devices become more complex, allowing them to participate in more bandwidth-intensive services such as streaming audio and video.

Ports

As penetration by MetroE increases, it brings Internet and other services (e.g.: VoIP) within reach of larger numbers of subscribers. To handle this growth in subscribers, edge platforms are getting larger and larger, with port numbers reaching or exceeding one-thousand per platform with regularity.

The current market metric for ports is around five thousand on a MetroE platform. With advances in filtering and queuing engines, combined with the translation of wireless sessions into virtual ports, that number can be expected to increase rapidly over the short-term.

Subscribers

As mentioned, subscriber numbers are increasing beyond those envisioned by early platform and network architects. As numbers and complexity of services grows, there results an increase in the amount of processing per subscriber for accounting, administration and security.

On a per-subscriber basis, the MetroE will be delivering access to the network, guaranteed levels of service, high-volume and –bandwidth services like Video on Demand, and session privacy from end-to-end.

Extensibility

The Metropolitan Ethernet networks are in a state of evolution. Originally, they were intended primarily to provide geographically-local access and infrastructure for interconnecting business (e.g.: a Metropolitan Area Network [MAN] carrying Transparent LAN Services [TLS] for urban businesses).

With early implementations based on optical transport, as costs dropped and technology improved Ethernet has become the transport of choice. The reasons are numerous, but perhaps the most important is the virtual extension of all the Ethernet feature sets originally developed for the Enterprise, and the seamlessness with which those services can be carried over into the wide area using Metropolitan Ethernet.

Changes in roles

As such, the MetroE network is changing from access control and local services to a more broadband aggregation and global services distribution role.

While local services were few and limited (e.g.: AAA and simple QoS provisioning) the newer services involve a much more broad portfolio including: greater security, more granular QoS, increased auditing of network usage, global network awareness and distribution of streaming media data streams.

Changes in standards

The delivery of these enhanced services at the local level has resulted in the formation of multiple standards bodies (e.g.: the Metropolitan Ethernet Forum [MEF]), or new working groups in the IEEE and IETF (e.g.: Ethernet in the First [Last] Mile, or Pseudo-Wire Emulation Edge to Edge [PWE3]).

Standards are evolving in areas as diverse as accounting, security, signaling and multicast distribution of streaming data.

Virtual Private LAN Services

At the edge, then, where aggregation combines with delivery of services and accounting for the usage of, and securing services for each subscriber is most important, standards have evolved to extend the environment of the local LAN across wide-area IP connections to create virtual LANs across the public infrastructure.

VPLS is a Layer 2 Virtual Private Network [VPN]. It provides to customers all the benefits of a Layer 3 VPN with the added speed of a switched network. The MPLS backbone used in the VPLS architecture is expected to behave like an 802.1a/d Learning Bridge. This gives the end users the appearance of being locally connected to all remote users...

VPLS was conceived to provide a service with multi-point connectivity over a packet switched long haul network. Many current systems for long haul networks use circuit oriented systems like TDM, FR or ATM. The advent of packet switched networks, especially those based on MPLS backbones, has shown a marked reduction in cost of maintenance over the traditional

networks. Hence a move to provide local services over the packet switched networks has proven profitable and VPLS and Ethernet access networks are examples of the same. VPLS works using LAN emulation over the MPLS backbone. MPLS LSPs can be setup using any supported protocol (e.g.: the ReSource reSeRvAtion Protocol [RSVP]) while the Labels specific to VPLS are distributed using the Label Distribution Protocol [LDP]. With the implementation of VPLS, providers have the ability to aggregate large numbers of subscribers, as they are no longer tied to physical ports, and implementation scalability is bounded only by available memory and processing power. With VPLS, providers can deliver security to individual subscribers at little incremental cost, partitioning network resources in such a way that flows exist in their own virtual domains, grouped according to requirements, but separate and distinct from other foreign domains. With VPLS, providers can deliver whatever services are required, with the necessary QoS for guaranteed, on-time deliver of data to the user wherever they are, on whatever device is connected, whenever it is requested.

IPI METRO-E

IPI offers a broad portfolio to address the Metro Ethernet's need. MPLS being a key technology in supporting QoS, IPI offered the first commercial MPLS protocol stack back in 2001. Since then we have enriched and enhanced this product line. Today, we support LDP, RSVP-TE, OSPF-TE, VPLS (both Martini and Kompella), and BGP4, which enable applications like Traffic Engineering, VPNs (L2 and L3), Hierarchical QoS and many others. IPI has been offering L3 routing protocols since late 1990s, which include BGP, OSPF, ISIS and RIP just to name a few. As a result, we have a large customer base exercising our software everyday in the field and in their labs. With the demand for Ethernet moving from the edge into the core, we are very well positioned to provide carrier grade L3 software, enabling manufacturers to focus on developing their value-add applications. IPI also supports a full set of L2 bridging protocols, including for example 802.1X, 802.1p/Q, STP, RSTP, MSTP, and LACP. These are essential elements in supporting switches in a Metro Ethernet Network. Our L2 protocols go beyond the standards to enable QoS on different hardware platforms. In support of the requirement for triple play, IPI not only offers protocol level fast recovery, but also a mechanism for state preservation using check-pointing.

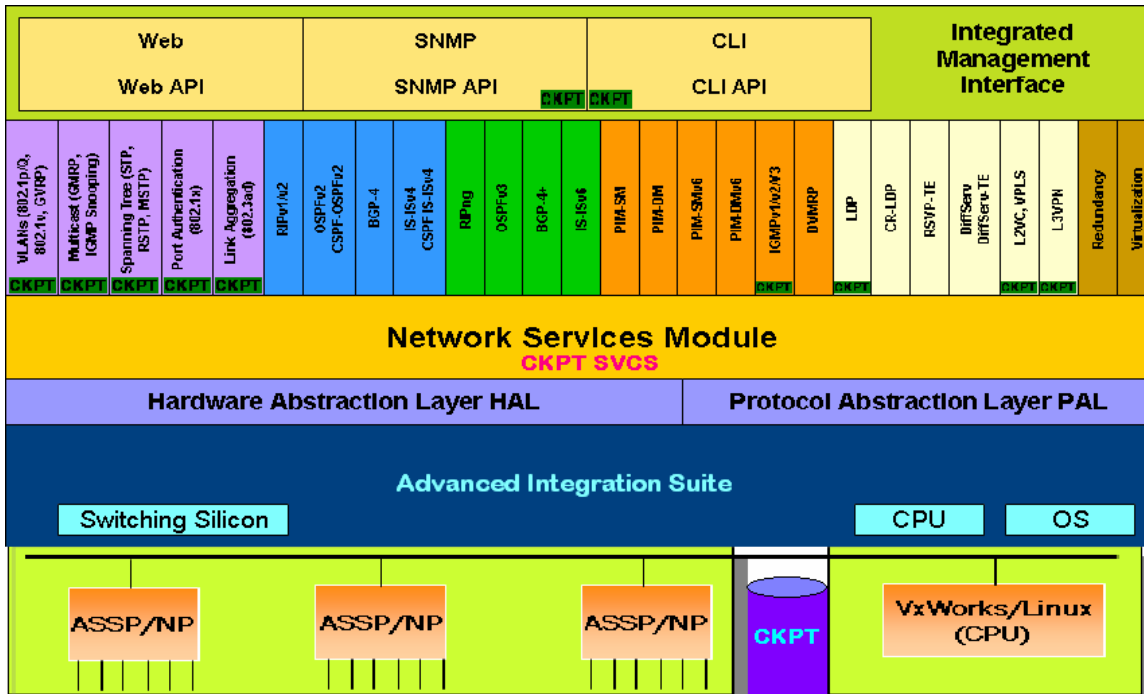


Figure 1 The full L2/L3 implementation.

Quality of Service

Integrated Services (IntServ), Differentiated Services (DiffServ) and Multiple Protocol Label Switching (MPLS) are the key technologies used for supporting L3 QoS in an IP network. IntServ, being used in early days, is implemented using Resource reSerVation Protocol (RSVP) to reserve resources through a network. DiffServ, more accepted and used today, uses the DSCP field in the IP header to mark the packet for the level of service required. MPLS, defined after IntServ and DiffServ, takes advantage of the existing protocols. MPLS enhances RSVP to RSVP-TE for supporting traffic engineering. MPLS uses its Experimental bits to carry the latest 3 of 6 DiffServ Code Point (DSCP) priority bits in the IPv4 header for packet prioritization. MPLS itself also defines the Forwarding Equivalence Class (FEC), which allows packets of the same or similar characteristics to be forwarded in the same manner.

IEEE802.1 p/Q is the standard to support L2 QoS, where the priority bit (p-bit) is used to mark the packet for level of service required.

The mechanism behind all these different QoS technologies is similar and can be summarized in the following:

Classification: based on the SLA set up between the service provider and the user, traffic is classified in different priorities. Traffic can be classified based on physical port, IP port, VLAN, TCP value...

Marking: after classification, packets will be marked accordingly based on the class it belongs to. Marking is necessary for the next routers to forward the packets accordingly.

Queuing: after marking, packets are then put on the priority queue for forwarding.

IPI has a full spectrum of solutions in supporting QoS. At the control plane level, we support protocols such as LDP, RSVP, RSVP-TE, and 802.1p/Q. We support mapping of the priority set up by the different QoS technologies. For example, the p-bit in MAC to DSCP in IP

header to E-bit in MPLS label. This allows the policy defined to be carried over throughout the network. IPI also has a well defined management framework and mechanism, which enable classification, marking and scheduling. It allows user to manage queues independent of the algorithms uses, such as WRR or WRED.

Understanding QoS heavily depends on the support of the underlying hardware, IPI continue to work with key hardware vendors in integrating our QoS solutions to different hardware platforms. We have successfully integrated our solution on the StrataXGSII, -III and Sandburst chipsets from Broadcom, and the DX and EX chipsets from Marvell.

Security

Authentication is one of the key requirements in supporting security. IEEE 802.1X is the most commonly used port-based authentication protocol in wire and wireless LAN. It passes the EAP frames between the RADIUS client and the RADIUS server for authorization. VLAN (IEEE 802.1Q) uses VLAN ID to grant access to the authorized users and deny unauthorized users. In the effort to provide security features in our product line, IPI not only support standards like IEEE 802.1X and 802.1Q, it also takes de facto standard into consideration. IPI has incorporated the Cisco proprietary implementation of root-guard, BPDU-guard and BPDU-filter into our bridging protocol to support hacking avoidance.

IPI allows users to ensure security at multiple levels such as at the physical port level, the virtual port level and protocol level. IPI provides Access Control List (ACLs) for filtering inbound and outbound packets based on different user-configurable criteria, ranging from port- to vlan- to IP address-based criteria. IPI also defines a framework that can be easily integrated with a third party security entity. [We have successfully integrated with IPNet/IPSec/IKE from Interpeak (now WindRiver)].

Scalability

IPI defines a modular architecture, called ZebOS. This architecture allows each protocol to run as a separate process, which enables the partitioning of the software not be limited on one processor. This architecture also allows the control protocol to support multiple data plane, which enables a distributed forwarding architecture. This is crucial in supporting the high number of ports required in rolling out new services like VoIP.

IPI continue to test its software of different network scale. In one of our L2 tests, we obtain results of supporting 5K MAC addresses, 4K VLANs, 48 simultaneous STP ports, 16+ MSTP instances. In one of our L3 tests, we obtain results of supporting 200K RIB routes, 200K FIB routes, 256 VRs. Our software posts almost no limits on the network size.

Extensibility

IPI actively participates in the development of new standards at the IEEE, IETF and MEF, via our founder and CTO, Kunihiro Ishiguro. As a result, we take an active role in upgrading our software according to the latest standards developments.

IPI offers two product families, the ZebOS Advanced Routing Suite (ARS) and the ZebOS Advanced Integrated Suite (AIS). ZebOS ARS is a complete suite of Layer 2 switching, Layer 3 carrier-class routing, multicast and MPLS networking protocol software. The ZebOS ARS supports both IPv4 and IPv6. Each module in ZebOS ARS is scalable, robust, and standards-based, which supports industry standard and best-of-breed operating systems, control, and dataplane processors.

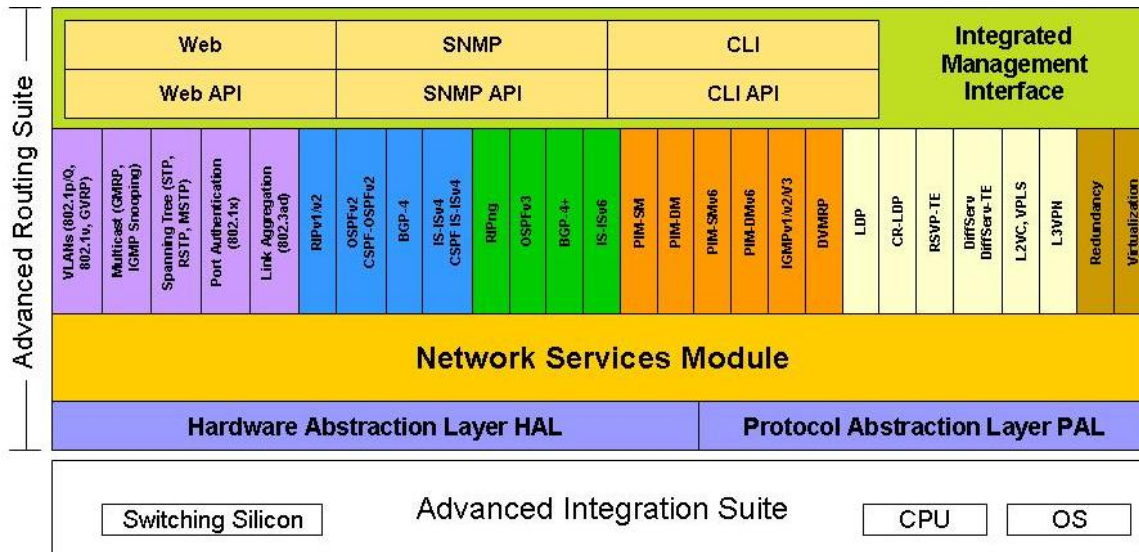


Figure 2 The Advanced Routing Suite

Each protocol module in ZebOS ARS is built on the ZebOS Network Services Module (NSM), which is the base module that simultaneously and independently communicates with every ZebOS ARS routing and switching process. The NSM manages both the route table and each of the enabled protocols; performs route conversion and redistribution; and manages the interface state, routing policies and filtering. The ZebOS Integrated Management Interface (IMI) provides a command line interface and SNMP capability that can be used as is by vendors, or can be integrated into the vendors existing management infrastructure. APIs defined in both NSM and IMI are accessible and extendable.

AIS is a suite of software platforms called the Hardware Integration Platform (HIP), each of which is created for and pre-integrated with an industry leading merchant silicon, and operating system. These HIPs provide a comprehensive forwarding plane implementation supporting L2, L3 (IPv4 & v6), multicast and MPLS/Traffic Engineering. A ZebOS AIS HIP, when combined with the ZebOS ARS protocol software for the control plane, provides a full system solution for enterprise switching, metro Ethernet, access, edge, mobile wireless and advanced IP services applications.

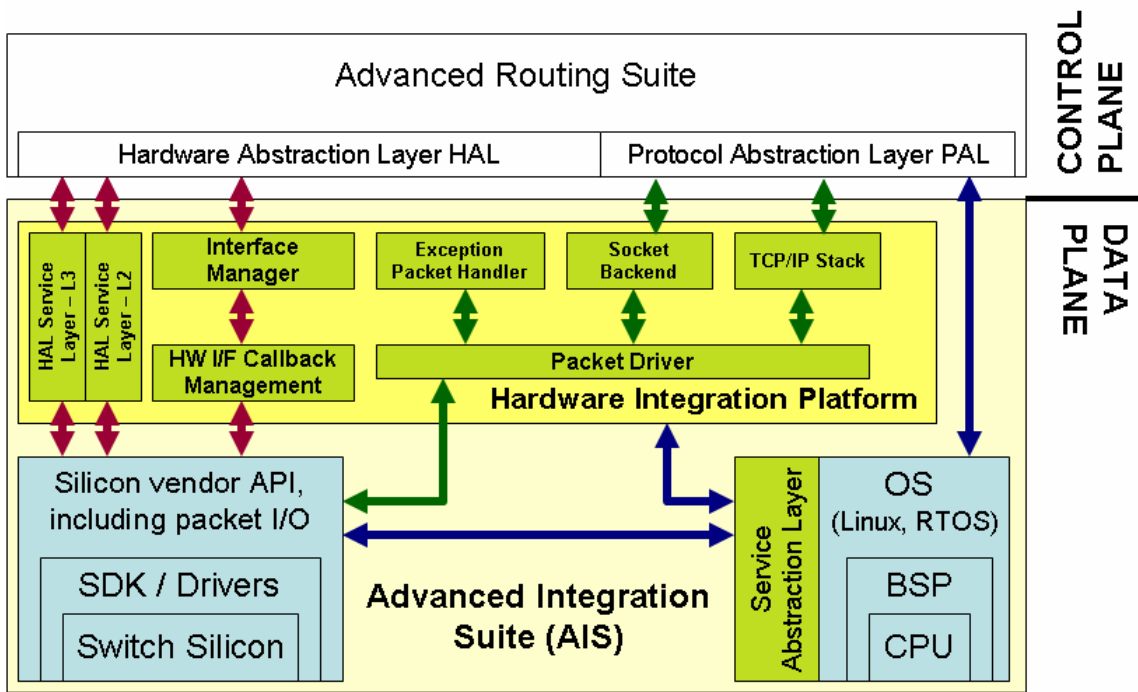


Figure 3 The Advanced Integration Suite

Two key components are defined in the ZebOS architecture to support AIS. Platform Abstraction Layer (PAL) and Hardware Abstraction Layer (HAL). PAL is a set of well defined API, which abstracts the operating systems calls such that ARS can use a unified set of API to access system resources and services (e.g.: memory, timer, ...) on any operating systems. Hardware Abstraction Layer (HAL) is the well defined and extensible API between the ARS (control plane) and the AIS (forwarding plane). HAL isolates all of the hardware platform specific interaction into a small set of well defined function calls for the ARS (control plane). The HAL provides a unified interface for the control plane to interact with the forwarding plane for all L2, L3, multicast and MPLS forwarding needs. The function calls above the HAL runs unmodified for any switching and routing hardware platform. The result is that customers have the full flexibility to select only the required protocol modules in the most cost and code space effective way.

Please visit IPI web site for more information on ZebOS architecture and our product offerings www.ipinfusion.com.



© Copyright 2006 IP Infusion Inc. All Rights Reserved.

ZebOS and IP Infusion are registered trademarks and the ipinfusion logo is a trademark of IP Infusion Inc. All other brands or product names are trademarks or registered trademarks of their respective holders. All specifications within this document are subject to change without notice. Contact Sales for current feature availability.