ADVANTECH
Enabling an Intelligent Planet

intel IoT Solutions Alliance Premier

Product Brief

Advantech Packetarium XLc PAC-6009

Carrier Grade Blade Server for Edge Computing and NFV

Table of Contents

Introduction

Subscriber demand and the addition of new services are driving the need for greater scalability and elasticity at the network edge. Applications such as the Internet-of-Things (IoT), augmented reality, content optimization, as well as video and data analytics need to run on server-class equipment close to subscribers where they not only enhance mobile user experience, but also alleviate the mobile core from additional congestion. However, the equipment hosting these new services at cell sites, in base stations, or telecom room aggregation points is limited by different physical and environmental constraints not suitable for standard white-box servers and hyperscale cloud equipment.

What's needed is a new class of system designed to address the requirements of mobile edge computing and the Centralized—or Cloud— RAN (C-RAN). Such a system needs to be rugged, flexible and highly scalable in order to meet edge performance requirements while complying with the related equipment practice and environment. It needs to combine high-performance server processing, high-throughput switching, and carrier-grade availability with Network Function Virtualization (NFV) support and work seamlessly with the existing infrastructure in aggregation points at the network edge.

Advantech has designed the Packetarium XLc Carrier Grade Blade Server to meet all of these needs by bringing together the IT and networking features of the next generation Intel® Xeon® Processor D-1500 product family with its high availability and NEBS-compliant design principles. The resulting versatile and scalable NFV platform, the Packetarium XLc PAC-6009 makes it easier to deploy the same Virtualized Network Functions (VNFs) that are used in the data center or core, anywhere else in the network. The advantage is the flexibility to bring network functions out to base station sites and to enterprises or centralize the functions for greater efficiency.

New Opportunities at the Mobile Edge

As the C-RAN develops into a more open and flexible architecture, it can be used to further optimize existing functions and facilitate the hosting of new applications at the network edge. Running applications at the edge of the wireless network noticeably improves application response times, and delivers a superior Quality of Experience (QoE) to the subscriber.

At the same time, it reduces the volume and the peak traffic levels on backhaul links, allowing operators to upgrade RF and baseband resources instead of their backhaul links to meet additional mobile data demand. It also allows new services innovation for location-aware applications. This enables operators to implement additional value-add mobile services that generate new revenue streams and deliver a better user experience.

The greatest opportunities to increase revenue through new services are expected to come from a wide range of new edge applications such as Domain Name System (DNS) and content caching, video analytics, collision avoidance for autonomous vehicles, and active device location tracking, to name just a few. By implementing the scale-out architecture of the Packetarium XLc, new services can easily be deployed and expanded with sufficient processing headroom to meet sharp increases in subscriber demand. Packetarium XLc also makes best use of energy-efficient CPUs that balance performance and cost to match the needs of applications at the wireless edge while boosting the return on investment.

A Leading Carrier Grade NFVI platform

With a shallow depth of 400mm, straight front to rear airflow and a power consumption of no more than 400W per rack unit Packetarium XLc can easily be installed and operated in data centres, central offices and telecom rooms at the edge of the network alike. The system has been carefully designed to meet Carrier Grade requirements in these environments including NEBS level 3 compliance and five 9's availability.

While the failure of a single system in a hyperscale data center hosting racks and aisles of racks full of NFV infrastructure may not be deemed critical, the situation changes as equipment gets deployed closer to the edge of the network. Each edge site may only host a few racks or not even a fully populated rack of NFV infrastructure (NFVI) in small deployment scenarios. Failure of a system may have a big impact on service availability and user experience in these cases. As many sites at the edge of the network are unmanned, remote locations yielding longer mean time to repair and higher service cost compared to central offices and telecom data centres, availability of each NFVI at the edge of the network becomes even more important.

Packetarium XLc supports high availability via 2+2 redundant power supplies, the ability to withstand a single fan failure, redundant system management, system fabrics switches and control modules. Hot swap support for field replaceable units guarantees maximum uptime.

One fundamental principle of telecom networks is the ability to control network elements and the network itself independently of the traffic load that users impose. This concept is usually referred to as the separation of control and data planes. Congestion of the data plane due to excess user traffic may not impact the ability to steer the network via the control plane. Network Function Virtualization allows the construction of flexible network infrastructure that can be adapted to various traffic scenarios easily. In order to reduce OPEX and CAPEX, these networks will not be built out to handle peak capacities of all supported services at the very same time. Varying traffic situations will no longer require traffic to be routed to network elements running on top of dedicated hardware, but will be supported by spinning up and down virtual instances of these network elements on multi-purpose NFV infrastructure. Such adaptive management of the network infrastructure will rely on the availability of a separated control plane more than ever. As Packetarium XLc has been architected for telecom use from ground up, it has built-in support for separate control and data planes via two independent Ethernet fabrics. To meet the availability requirements inherent to carrier grade NFV, both fabrics implement a dual star topology which avoids a single point of failure in this critical system infrastructure.

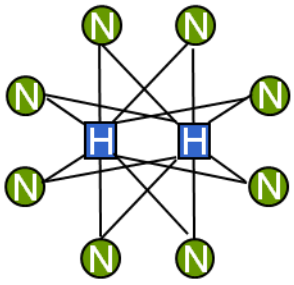

Figure 1 Dual Star Topology with each compute node (N) connecting to both switches (or hubs) (H)

Both the control and data plane fabrics are connected to two redundant switch management blades which also integrate system and chassis management functions for maximum efficiency and density at lowest cost.

Another key aspect of NFV Infrastructure is the ability to scale out rapidly. Building out service capacity quickly encompasses more than optimizing the time to install and provision additional platforms. NFV Infrastructure needs to make best use of available floor space of a facility and optimize the use of the overall supporting infrastructure such as the power distribution and cooling systems. Systems need to be optimized for the optimum density per square foot that can be supported in terms of power delivery and cooling in a given location. At the same time, attention needs to be paid to power efficiency as energy costs are one of the key contributors to the OPEX. With a power rating of 2.4kW per chassis and 16.8kW per 42U rack, Packetarium XLc not only matches the capabilities of legacy and newly built premises but also provides best-in-class compute density.

No matter if compared against legacy telecom platforms such as AdvancedTCA (ATCA), traditional IT approaches such as stacking 1U whitebox servers or recent approaches applied in hyperscale data centres such as Open Compute Project (OCP), Packetarium XLc achieves the highest compute density per kW and per square foot of floor space.

The diagram below shows the related numbers for a full rack of equipment that implements a NFVI with redundant 1GbE control plane and redundant 10GbE data plane and two redundant cloud control nodes. This comparison uses the number of Xeon® Cores multiplied by the CPU's frequency as a density metric as it, on one hand, removes dependencies on different processor SKUs supported on different platforms and on the other hand is an appropriate metric for highly virtualized workloads like NFV. The power budget is based on the consumption of a full rack and does not include the overhead and losses incurred by facility level power distribution and cooling as these may vary case by case. Similarly, the floor space only includes the rack itself and does not include aisle space and space required for supporting facility infrastructure. Details of this comparison can be shared on request.

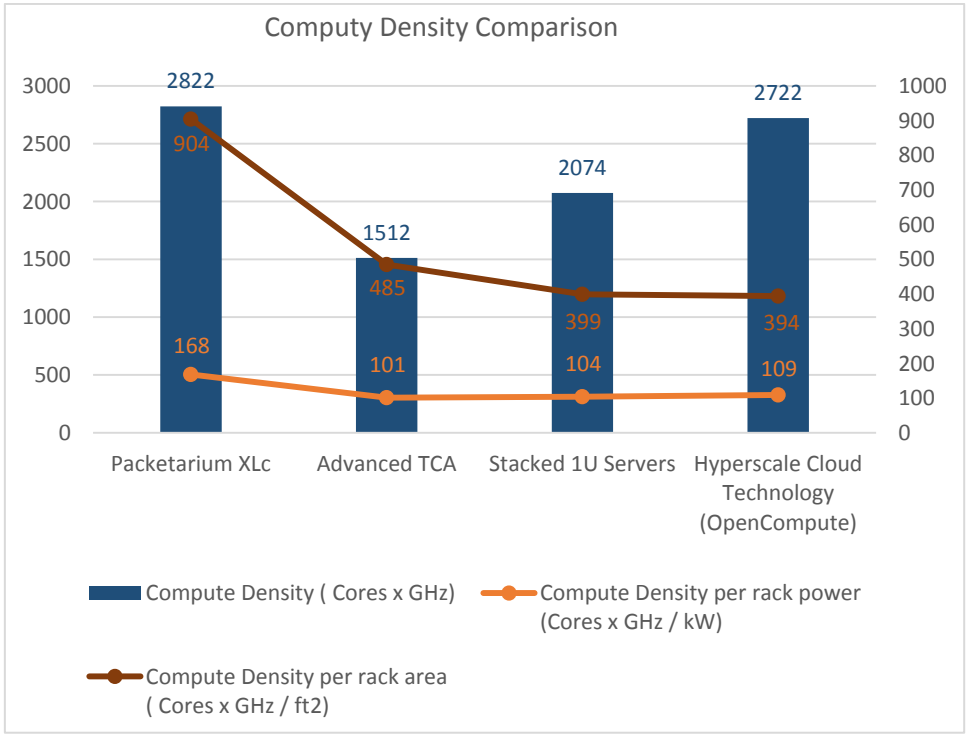

Figure 2 Packetarium XLc Density

Even when looking at a minimum deployment scenario, Packetarium XLc outperforms competing architectures by far. When compared against stacking 1U servers as the “leanest” approach, such deployment would require 2 GbE Top-of-Rack Switches for control plane, two 10GbE Top-of-Rack Switches for data plane networking, two servers as cloud control nodes in addition to the compute nodes that would run the VNFs. Just this basic system infrastructure requires a 6U rack space and an estimated 1.8kW of power and cooling. A single Packetarium XLc system also occupies just 6U of rack space but can host up to 14 Intel® Xeon® processor nodes in addition to hosting the same system infrastructure at an overall power budget of 2.4kW.

Ease of Deployment

Packetarium XLc can be easily deployed in existing locations as well as new facilities. Supporting standard 19” 600mm rack infrastructure, AC and DC power supply options combined with the bespoke modest power and thermal requirements, facilities do not have be re-modelled or specifically built to accommodate high power gear deployed in non-standard, 1200mm deep racks.

Packetarium XLc systems are lightweight and do not require special procedures for installation. Having shared power infrastructure per chassis and integrated switch fabrics vastly reduces cabling effort which saves cost, installation time and is less prone to errors at higher reliability. Compared to stacked 1U servers, Packetarium XLc saves more than 180 cables per pack and saves 250kg of equipment weight.

Advantech Packetarium XLc PAC-6009 System Overview

Packetarium XLc system is designed to accommodate the highest density of compute available in a 400mm deep 6U carrier-grade chassis. It is a highly scalable platform for deploying Intel® Atom®, Intel® Xeon® E5 series and Intel® Xeon® Processor D -based blades designed for the most demanding NFV workloads. For media processing and content handling needs, it offers a flexible infrastructure for the addition of video acceleration hardware providing unprecedented video encoding capabilities. The system’s 400W per RU power footprint enables deployment in industry standard 19” racks.

The PAC-6009 chassis incorporates a highly versatile and modular design with 9 front slots to host 9 single or dual node Intel® Xeon® processor blades. Generic compute blades run application workloads (VNFs) while dedicated cloud control nodes provide orchestration and virtual infrastructure management functions. Cloud storage nodes for content caching and nodes optimized for media and image processing are in preparation.

Compute blades are available as single or dual node blades:

- **The MIC-8302C node blade** is based on a single Intel® Xeon® Processor E5-2600 v3 with up to 14 cores.
- **The MIC-8304C node blade** is based on a single 8- or 16-core Intel® Xeon® Processor D with additional on-board SSD storage.
- **The MIC-8303C** is a dual node blade equipped with 8- or 16-core Intel® Xeon® Processor D.

The system includes a 9-slot mid-plane connecting two switch / management modules to each of the nodes and the infrastructure components within the chassis.

The PAC-6009 has been designed for high availability and supports hot-swappable FRUs including blades, hard drives, switch / management modules, fans and power supplies.

Figure 3 shows the elements of the PAC-6009 System.

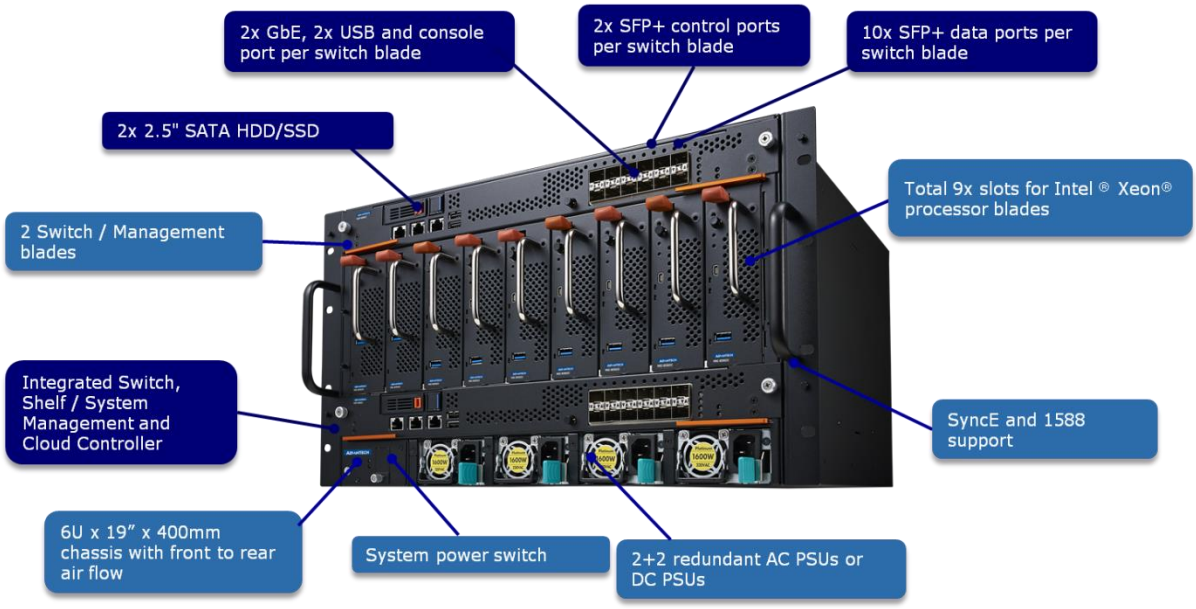

Figure 3 Packetarium XLc PAC-6009 System

The two redundant ESP-9002C switch/management modules handle network communication between node blades on the internal backplane fabric and provide external system connectivity. The switch/management blades also provide system management functionality.

The ESP-9002C uses the high performance, low-latency Broadcom StrataXGS® Trident+ BCM56842 providing Layer 2 routing and Quality of Service (QoS) management for the dataplane. Control plane switching is implemented via a Broadcom BCM53346 device.

The ESP-9002C is fully compatible with the Broadcom SDK (Software Development Kit) and offers Layer 2 fast path capabilities. Switch management is handled by an on-board Freescale QorIQ P2040 processor. System and chassis management are based on Advantech's proven and robust SMM-5060 shelf manager. This PICMG3.0/IPMI 2.0 compliant chassis management solution provides fail safe, carrier grade system management including FRU status and health monitoring, hot swap, system thermal control and standards based firmware upgrade management.

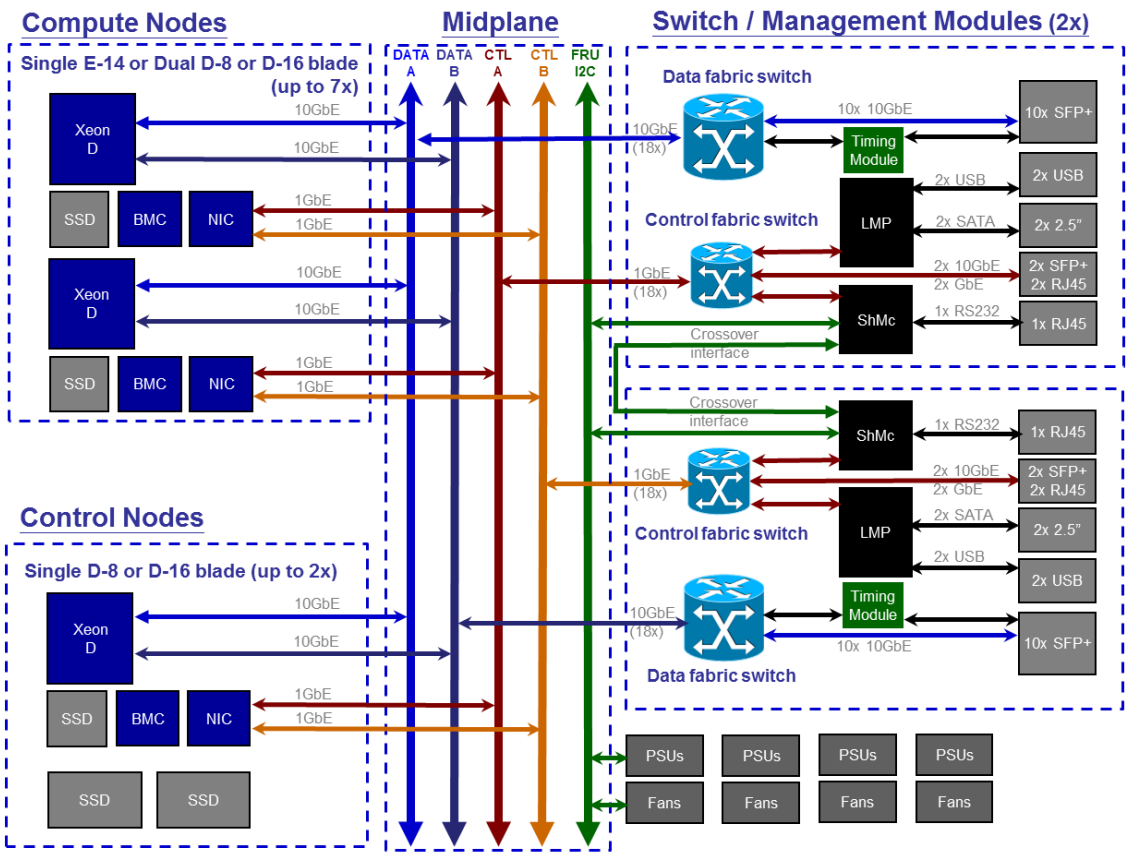

Figure 4 PAC-6009 High Level System Block Diagram

Redundancy schemes and dedicated crossover interfaces allow for failover between switch / management modules and rack level stacking support simplifies overall cluster management.

A total of 10 SFP+ ports are available for external dataplane connectivity and two SFP+ ports for control plane connectivity. 2 GbE copper ports, an RS232 console and USB ports provide optional and direct access to the management subsystem.

The system provides a 2+2 redundant power supply (PSU) set-up supporting redundancy with -48VDC DC feeds or with 220V AC feeds. Support for 2+2 redundant AC power schemes simplifies the deployment in legacy environments that partially only supply two circuits per rack. Moreover, it provides for enhanced system availability compared to N+1 topologies as the interruption of one phase will not compromise the system's power integrity.

Front attached cable trays are available as accessories and simplify rack level cable management.

The main system specifications are summarized in Table 1.

Figure 5 PAC-6009 System with cable management

Table 1 PAC-6009 System Specification

Slots	9 front slots to host either – MIC-8303C Dual Node Intel® Xeon® Processor D-1500 SoC (D-8 or D-16) compute blades – MIC-8302C Single Node Intel® Xeon® Processor E5-2600 v3 compute blades – MIC-8304C Single Node Intel® Xeon® Processor D-1500 SoC (D-8 or D-16) (Cloud controller node)
Midplane	9 slot mid-plane providing dual star connectivity between Switch boards and blades. Each fabric and star network supports: – Dual GbE control/ management ports per slot – Dual 10GbE data fabric ports per compute slot
Integrated Switching	2 ESP-9002C Switch/Management Modules – Broadcom Trident+ BCM56842 data plane switch – Broadcom BCM53346 GbE control plane / management switch
System IO (per switch blade: total = 2x)	– Data fabric interfaces: 10 x SFP+ ports (10G-SR, -LR, -ER) – Control fabric : 2 SFP+ ports – 2 x RJ45 GbE Management ports, 1 x RJ45 console, 2 x USB
System storage/NAS	Up to 2 x 2.5” hot-swappable HDD trays for SATA disks / SSD per switch blade as integrated system NAS (optional)
Cooling	4 hot swappable fan modules with a total of eight 9238-type high speed fans
Accessibility	
Power	2+2 redundant power supply scheme

	AC SKU: 2+2 3600W CRPS @200VAC~ 230VAC input DC SKU: 2+2 3600 CRPS @ -40VDC~72VDC input
Shelf Management	Based on Advantech IPMI Core G03
Miscellaneous	
Physical Characteristics	Size: 432mm x 6U x 400mm (W x H x D)
Environment	Transportation and storage temperature: -40C ~ 60C, 95%RH Operating temperature: -5C ~ 40C, 95%RH, 55C (short term) Operating altitude: 13000 feet @ 45C max Operating vibration: 0.0005 g^2/Hz, 0.5 Grms, 5-500Hz Operating noise: 78 dBA
Compliance	– FCC/CE/VCCI/C-Tick/KCC/BSMI/CCC/UL/CB – Designed to meet GR63-CORE – RoHS 6/6

Figure 6 shows the front view of a fully loaded PAC-6009 system and images of the node blades and switch / management modules.

Dual switch/management blades

High performance, low-latency switching and management

9 individually hot-swappable server-class node blades

Compute, storage and acceleration

- **10GbE** switching between 10 external ports and 9 internal blades.
- **Broadcom StrataXGS® Trident+ BCM56842 switch** for L2 routing and dataplane QoS management. Broadcom SDK support and L2 fastpath.

- **MIC-8302C** node blade based on a single Intel® Xeon® Processor E5-2600 v3 with up to **14** cores.
- **MIC-8304C** node blade based on a single 8- or 16-core Intel® Xeon® Processor D with onboard SSD storage.

Figure 6 Front view of PAC-6009 System

Figure 7 is a rear view of the Packetarium XLc PAC-6009 System with its 4 rear-serviceable hot-plug fan trays.

Figure 7 Rear View of PAC-6009 System

Packetarium XLc Node Blade – MIC-8302C single Intel® Xeon® Processor E5-2600 v3

Figure 8 MIC-8302C Node Blade

This compute blade is based on a single Intel® Xeon® Processor E5-2600 v3, memory, one m.2 2242 SSD and 1GbE and dual 10GbE internal network connections to each of the switch modules.

Figure 8 shows the MIC-8302C and Table 2 shows its key specifications.

The Intel® Xeon® Processor E5-2600 v3 product family adds 50 percent more cores and cache over the previous generation. It furthermore includes numerous other hardware enhancements to deliver up to 2.2x the performance over the previous generation, significantly boosting output across a broad set of workloads. The Intel® Xeon® Processor E5-2600 v3 product family also delivers an increase in virtualization density of up to 1.6x compared to the previous generation, building on an ever more important capability for high performance server class processing.

Table 2 MIC-8302C Product Specifications

CPU	- Single Intel® Xeon® Processor E5-2600 v3 (max. 14 cores per CPU) - Max. 120W
Memory	- 4 channel memory design - 4 DDR4 ECC RDIMMs @ 2133MHz - Max. 128GB
Midplane I/O	- Up to 4x 10GbE data ports via an Intel® Ethernet Controller XL710 - 2x GbE control / management ports via 2x Intel® Ethernet Controller I210
Management	- Advanced, carrier grade BMC supporting PICMG3.0/IPMI 2.0 (Advantech IP)
Offload	- Intel® DH8950 QuickAssist for encryption acceleration (optional)
Storage	- 1 m.2 2242 SATA-III SSD (max. 256GB)

An upgrade to the next generation Intel® Xeon® Processor E5-2600 product family will be available in 2016. Please contact your Advantech representative for details.

Packetarium XLc Node Blade – MIC-8304C single Intel® Xeon® Processor D-1500

This control blade contains one Intel® Xeon® Processor D, up to 128MB DDR4 DRAM with ECC, one m.2 2242 and two 2.5” SATA SSDs. It features 1GbE and 10GbE internal network connections to each of the switch modules. Figure 9 shows theMIC-8304C and Table 3 shows the key specifications.

Figure 9 MIC-8304C Node Blade

The blade is based on the Intel® Xeon® Processor D family which offers new options for infrastructure optimization by bringing the performance and advanced intelligence of Intel® Xeon® processors into a dense, lower-power system-on-a-chip. The Intel® Xeon® Processor D product family is the 3rd generation 64-bit SoC and the first based on Intel® Xeon® processor technology. Leveraging industry-leading 14 nm silicon technology, this new product family is the first to address a broad range of low-power, high-density infrastructure needs. The system on a chip (SoC) has an integrated platform controller hub (PCH), integrated I/O, two integrated 10 Gigabit Intel® Ethernet ports, and a thermal design point (TDP) of 20 watts to 65 watts. It can run the same instruction set as more robust Intel® Xeon® processors to provide software consistency from the data center to the network edge.

Table 3 MIC-8304C Product Specifications

CPU	– Single Intel® Xeon® Processor D-1500 (max. 16 cores per CPU) – Max. 65W
Memory	– 2 channel memory design – 4 DDR4 ECC RDIMMs @ 2400MHz max. – Max. 128GB
Midplane I/O	– 2x 10GbE data ports via a NIC integrated in the SoC – 2x GbE control / management ports via Intel® Ethernet Controller I350
Management	– Advanced, carrier grade BMC supporting PICMG3.0/IPMI 2.0 (Advantech IP)
Storage	– 1 m.2 2242 SATA SSD (max. 256GB) – Two 2.5" SATA SSDs (max. 2TB)

Packetarium XLc Node Blade – MIC-8303C dual Intel® Xeon® Processor D-1500

This compute blade contains two Intel® Xeon® D processors, up to 128MB DDR4 DRAM with ECC per SoC and one m.2 2242 SATA SSD. It features 1GbE and 10GbE internal network connections to each of the switch modules per processor/node. Figure 11 shows the MIC-8304C and Table 4 shows the key specifications.

The blade is based on the Intel® Xeon® Processor D family which offers new options for infrastructure optimization by bringing the performance and advanced intelligence of Intel® Xeon® processors into a dense, lower-power system-on-a-chip. The Intel® Xeon® Processor D product family is the 3rd generation 64-bit SoC and the first based on Intel® Xeon® processor technology. Leveraging industry-leading 14 nm silicon technology, this new product family is the first to address a broad range of low-power, high-density infrastructure needs. The system on a chip (SoC) has an integrated platform controller hub (PCH), integrated I/O, two integrated 10 Gigabit Intel® Ethernet ports, and a thermal design point (TDP) of 20 watts to 65 watts. It can run the same instruction set as more robust Intel® Xeon® processors to provide software consistency from the data center to the network edge.

Figure 10 MIC-8303C Node Blade

Table 4 MIC-8303C Product Specifications

CPU	– Dual Intel® Xeon® Processor D-1500 (max. 16 cores per CPU) – Max. 65W per processor
Memory (per CPU)	– 2 channel memory design – 4 DDR4 ECC RDIMMs @ 2400MHz max. – Max. 128GB
Midplane I/O (per CPU)	– 2x 10GbE data ports via a NIC integrated in the SoC – 2x GbE control / management ports via 2x Intel® Ethernet Controller I210
Management	– Advanced, carrier grade BMC supporting PICMG3.0/IPMI 2.0 (Advantech IP)
Storage (per CPU)	– One m.2 2242 SATA SSD (max. 256GB)

Packetarium XLc Switch/Management Module – ESP-9002C

Two hot-swappable switch / management modules are used to provide integrated control and data plane switching in an active-active redundancy scheme. Each switch provides up to 280Gbps data plane switching capacity and supports up to 20Gbps bandwidth to each node blade.

Figure 12 ESP-9002C Switch / Management Module

A separate gigabit Ethernet switch provides control plane switching with two GbE interfaces per node.

Both switches are managed by a Freescale QorIQ P2040 local management processor (LMP) with support for Broadcom SDK and Layer 2 FastPath. Support for OFDPA is under preparation.

The following table lists the main switch features:

- **Redundancy**
  - IEEE 802.1D STP, 802.1w RSTP, IEEE 802.1s MSTP
  - Port Trunking with LACP
  - Dual images
- **VLAN**
  - 802.1Q Port-Based, 4K VLAN
  - GVRP
  - Protocol-based/MAC based/IP subnet based VLAN
  - VLAN ID: 1~4094
- **Multicast**
  - IGMP Snooping v1/v2/v3
  - GMRP
  - MLD Snooping/Querier
- **Security Features**
  - 802.1x
  - RADIUS/TACACS+ Authentication
  - SSH, SSL, SFTP
  - ARP inspection, Port Security
  - Access Control List (ACL)
- **Management Features**
  - SNMP (v1,v2c,v3), RMON, SNMP Trap
  - Cisco-like CLI, Telnet, Web-based Management
  - Port Mirroring
  - Multi-function of reset button
  - DHCP option 82
- **QoS**
  - Jumbo frame support up to 12KB
  - SPQ, WRR queue scheduling.
  - 8 Hardware priority queues
  - Broadcast storm control

Support for an optional timing module based on IEEE1588v2 will be available in 2016 to support network synchronization using industry standard mechanisms.

The switch/management module also integrates system and chassis management functionality. Low level management is implemented on a dedicated, cost-efficient 32bit ARM processor and some higher level system management functions are implemented on the QorIQ based LMP. More information on management is provided in later sections of this document.

The switch blade provides the following external I/O ports:

- 10 SFP+ data ports
- Two SFP+ ports for control plane
- Two management (RJ45 GbE) ports for system and shelf management
- One system console (RS232) and two USB ports

Packetarium XLc Platform Software Overview

Advantech is integrating leading NFV middleware and applications from a world class ecosystem into the Packetarium XLc platform, turning the PAC-6009 into a fully-integrated, deployable NFVI. Advantech is an active member of the Intel® Network Builder Program and is closely aligning Packetarium XLc with major Intel® platform initiatives such as RackScale, and software frameworks such as Open Networking Platform (ONP Server). The Network Builders Program brings together leading industry members to foster closer collaboration, ensure interoperability and encourage innovation.

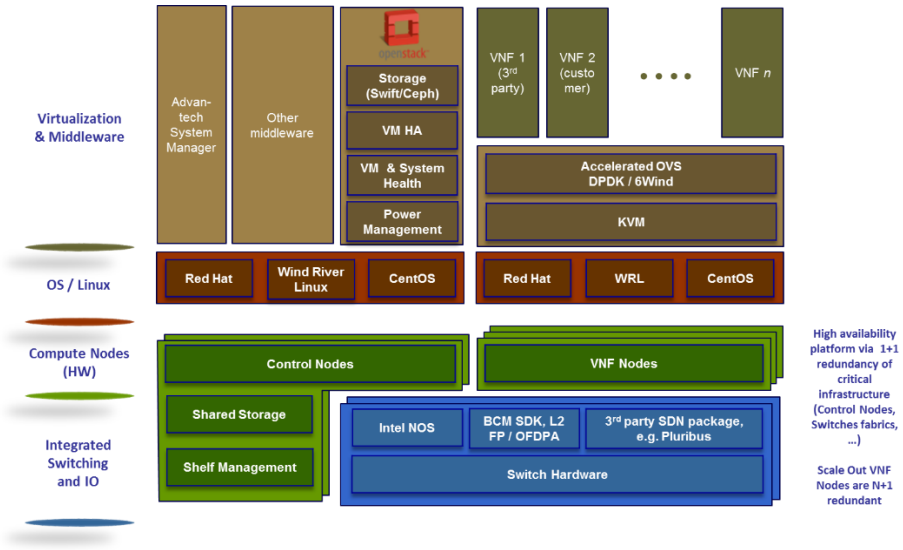

Figure 13 NFV Software Stack

Advantech's Remote Evaluation Services (RES) framework for NFV not only allows us to collaborate with the NFV software ecosystem for integration, interoperability testing and performance optimization but is also a vehicle for efficient life cycle management of the complete NFV platform.

Operating System Support

The general release of the Packetarium XLc PAC-6009 system will include the certification of several Linux operating systems including Red Hat Enterprise Linux and will support the Intel® Open Networking Platform 2.0 (including OPNFV).

Virtualization Support

Packetarium XLc will be integrated with OpenStack including an accelerated OVS based on DPDK and/or 6WIND Virtual Accelerator for higher performance and lower latency communication to, from and between VNFs. The system will integrate Carrier Grade enhancements to OpenStack including those implemented as part of the OPNFV industry initiative.

SDN Support

The integrated switch fabrics will support SDN control via Broadcom's OFDPA package. Enablement of 3rd party network OS's is under discussion.

Packetarium XLc Carrier Grade Platform Management

The system management of PAC-6009 is based on Advantech's widely deployed SMM-5060 shelf & system manager and is integrated on the switch / management modules. Low level shelf management runs on a dedicated ARM processor while a Freescale QorIQ P2040 runs switch management and higher level system management functions.

High availability of Shelf and System management is implemented by running the modules in an active/hot standby scheme. A low latency failover mechanism is provided by a robust, low level failover interface and using crossover Ethernet connections for more extensive state and log synchronization. The management LAN between the System/Shelf Managers and the front panel management I/O ports are independent of the system control plane for enhanced robustness.

Basic System Management Features

Shelf management has been designed as robust and as reliable as AdvancedTCA without resembling unwanted complexities and costs. By making re-use of existing and field proven Advantech IP based on open standards, both for blade management as well as system management, and by removing unnecessary overhead and employing faster interfaces, Packetarium XLc's management offers speed improvements over legacy architectures without compromising reliability. To make system management easier, the internal management infrastructure is kept "under the hood" and the System Manager becomes a central service/management access point.

The manager monitors system thermal management and overall system health, including the blades, power supplies and fans. The management framework provides a full inventory and event log and handles the hot plug of all Field Replaceable Units (FRUs).

Advanced Feature Set

The advanced feature set offers additional management capabilities including a **Serial-over-LAN (SoL) Proxy Console Server** which exposes a single SSH shell to the outside world. From within this secure shell, users can access each node's console window for administrative, service and troubleshooting purposes. Console stream buffering and replay for advanced trouble shooting is also supported. This Console Server does not require manual configuration but supports automatic connectivity to nodes in the system.

In addition, a **System NAS Boot server** for all nodes in the system allows compute nodes to boot over the control plane network. It employs a simple table-based configuration denoting which blade in which slot boots which specific image at start-up and enables load optimized BIOS settings for each workload.

The **Upgrade Agent** allows updating the complete system through the system manager instead of updating each FRU piece by piece. This centralized upgrade management helps to shorten maintenance windows and provides resilience such as a system wide rollback capability in case of an upgrade failure of a single component

Management Interfaces

A full suite of management interfaces is provided ranging from a command line interface for debugging purposes, to a Secure Shell (SSH), Simple Network Management Protocol (SNMP) and a Web interface.

System Explorer is a secure Web server with graphical user interface (GUI) that displays status and control information such as views of the system repository, sensor data and system health. It also provides access to system configuration and supports system maintenance tasks such as upgrade management.

Enhanced Platform Design

Advantech's platform solutions are more than just hardware designs as we go much further than an ordinary white box ODM. We ensure that our systems not only have outstanding stability secured by a world class design quality assurance (DQA) process, but are enhanced by building in features which improve availability, serviceability and usability. Features which have a moderate impact on cost as they have been carefully designed by Advantech's in-house engineering teams and are kept consistent across product lines.

Figure 14 Enhanced Platform Features

Advantech provides a comprehensive platform software framework to enhance reliability and improve total cost of ownership (TCO). Features such as redundant BIOS, redundant firmware, remote BIOS upgrades and configuration changes offer failsafe operations for carrier grade needs.

Redundant BIOS is achieved by physical redundant BIOS flashes maintaining one current and one backup copy. A watchdog mechanism detects a failing or corrupted BIOS and allows the node to recover through a rollback mechanism. Advantech provides a dedicated update utility that does not break the redundancy scheme.

Redundant Firmware is maintained by physical redundant baseboard management controller (BMC) firmware flashes containing one current and one backup copy. The system will recover from a failing BMC upgrade through the rollback mechanism.

BIOS upgrade from remote is controlled via the BMC which can upgrade the x86 BIOS. This allows system BIOS to be upgraded as long as standby power is available using the industry standard HPM.1 protocol. The host x86 processor can alternatively perform BIOS upgrades using ABU (Advantech BIOS Utility) for faster updates.

BIOS configuration change from remote allows BIOS settings (including default & current settings) to be changed remotely via the BMC. This offers the capability to match BIOS settings to workload specific requirements for optimum performance

Service Friendly System Design includes front and rear swappable FRUs which can easily be identified. FRU status indicators (LEDs) allow service staff to identify a faulty FRU without having to consult product documentation.

Summary

Packetarium XLc represents a new class of scalable platform that combines high-performance server processing, high-throughput switching and carrier-grade availability. It is a flexible and versatile system design that supports NFV and works with the existing infrastructure in locations such as central offices and telecom rooms, rather than hyper scale data centres.

Packetarium XLc, enabled by a rich middleware partner ecosystem, provides a solid NFV Infrastructure (NFVI) to application developers and content providers enabling them with the cloud-computing capabilities and IT-style service environment they need at the edge of the mobile network so they can drive new revenue streams for operators. In addition, the system addresses real network deployment needs in locations with high mobile subscriber density and where green field installations are not possible.

Packetarium XLc is the first commercial-off-the-shelf (COTS) telco-grade server of its class to extend Network Function Virtualization (NFV) beyond the core network to both edge and access equipment. The system scales compute performance over nine Intel® Xeon® processor blades. Higher processing densities and lower power footprints are achieved through dual Intel® Xeon® processor blades that provide a maximum number of powerful Intel® Xeon® processor cores in a compact 6U platform with a reduced depth of 400 mm. The system routes traffic through two redundant switches that connect to the dual-star backplane creating an internal network with no single point of failure. It integrates Advantech's Advanced Platform Management and is designed to meet demanding industry standards requiring five 9's availability and NEBS Level 3 compliance.

Packetarium XLc is a good example of how members of the Intel® Network Builders ecosystem play a key role in providing operators with the tools they need to deploy a scalable and flexible NFV infrastructure built on Intel® Xeon® processors. The system is a key element in Advantech's NFV Elasticity initiative in providing a range of scalable Intel® Xeon® processor-based platforms that can be deployed anywhere in the network to support a tiered model with different levels of service.

Advantech Contact Information

Hotline Europe: 00-800-248-080 | Hotline USA: 1-800-866-6008
Email: NCG@advantech.com
Regional phone numbers can be found on our website at http://www.advantech.com/contact/
www.advantech.com/nc

Ref: PACKXLC-PB-20151208