

# 6WIND Support for Intel® Data Plane Development Kit (DPDK)

# Topics

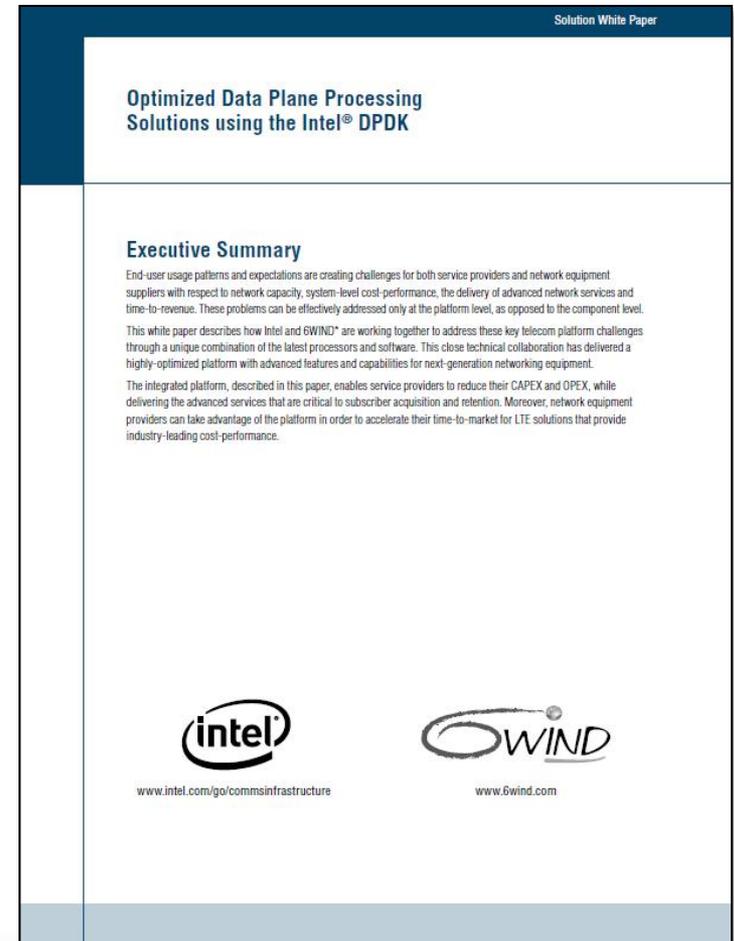
- **Introduction**
- **DPDK.org open-source project**
- **Intel® DPDK library from 6WIND**
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- **Summary**

# Topics

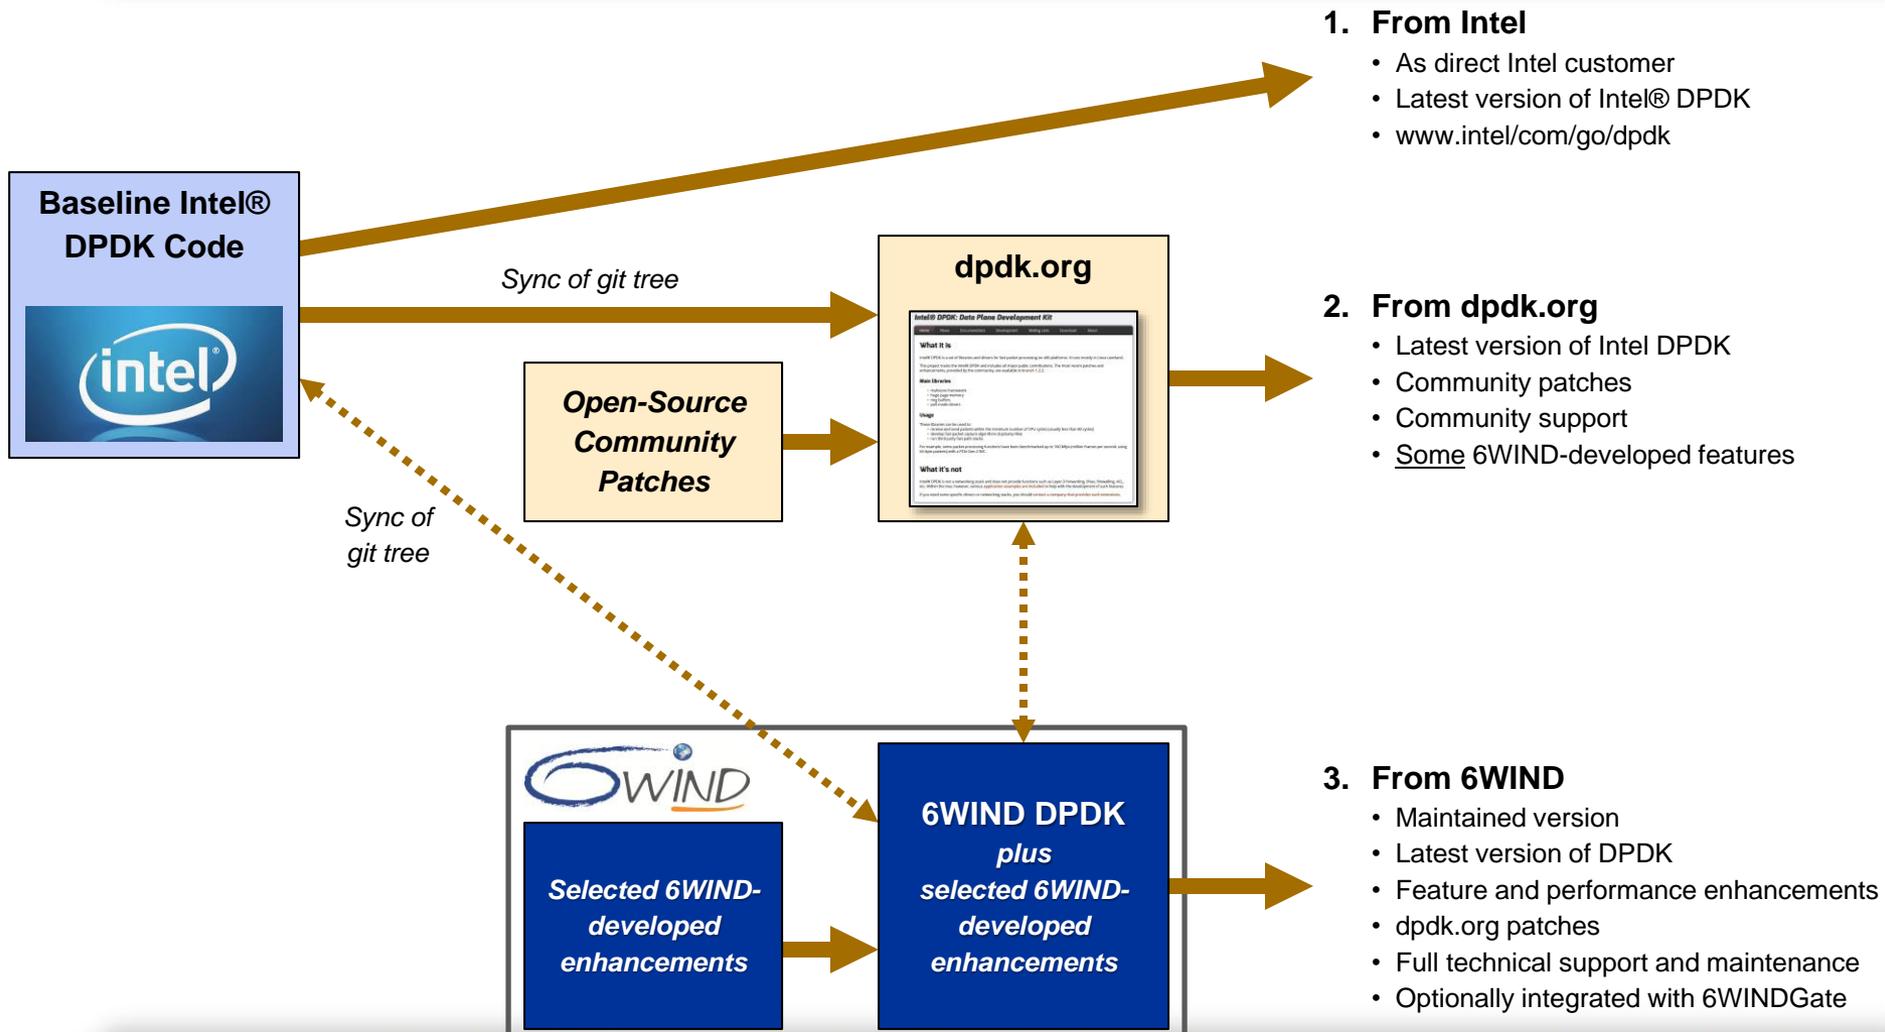
- **Introduction**
- DPDK.org open-source project
- **Intel® DPDK library from 6WIND**
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- **Summary**

# Introduction: 6WIND and Intel® DPDK

- **Long-term strategic partnership**
  - **First** high performance networking software on Intel processors, introduced before Intel® DPDK
  - **First** commercial networking software with Intel® DPDK support, demonstrated at Intel Developer Forum September 2011
  - 6WINDGate deployed on IA platforms by tier-1 networking and telecom equipment OEMs
  - Launched dpdk.org in April 2013



# Your Three Options for Obtaining Intel® DPDK

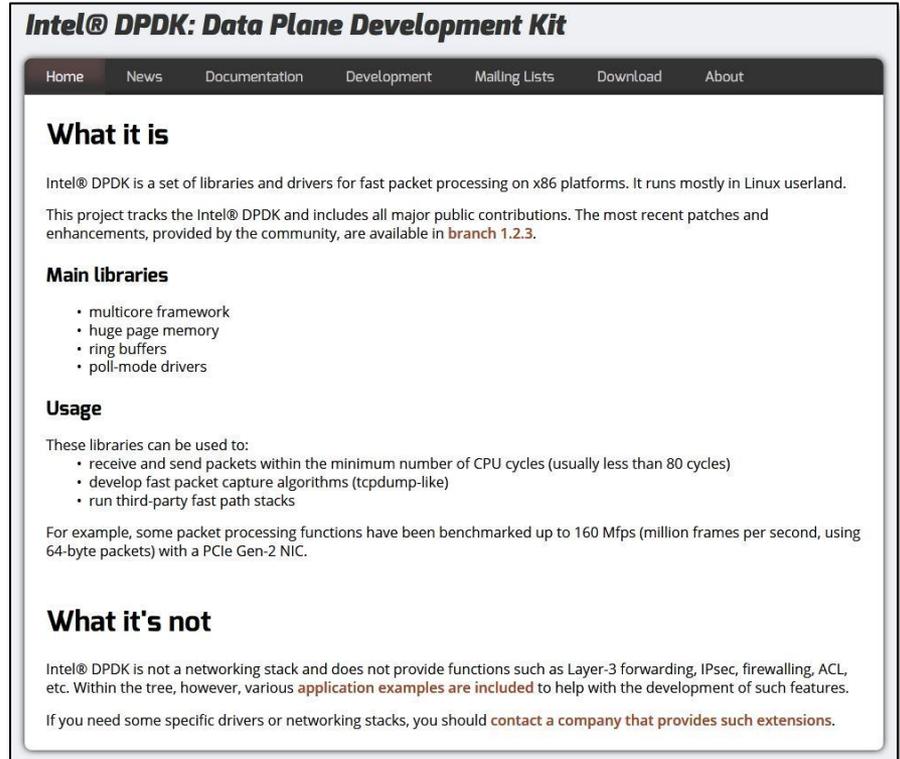


# Topics

- Introduction
- **DPDK.org open-source project**
- **Intel® DPDK library from 6WIND**
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- Summary

# DPDK.org Open-Source Project

- **Open-Source project established by 6WIND**
  - Announced on April 9, 2013.
- **Goal is to accelerate development of DPDK-based networking and telecom products.**
- **Comprehensive set of resources for developers**
  - Libraries
  - Drivers
  - Documentation
  - Development mailing list
  - git repository
  - Technical support.



**Intel® DPDK: Data Plane Development Kit**

Home News Documentation Development Mailing Lists Download About

### What it is

Intel® DPDK is a set of libraries and drivers for fast packet processing on x86 platforms. It runs mostly in Linux userland.

This project tracks the Intel® DPDK and includes all major public contributions. The most recent patches and enhancements, provided by the community, are available in [branch 1.2.3](#).

### Main libraries

- multicore framework
- huge page memory
- ring buffers
- poll-mode drivers

### Usage

These libraries can be used to:

- receive and send packets within the minimum number of CPU cycles (usually less than 80 cycles)
- develop fast packet capture algorithms (tcpdump-like)
- run third-party fast path stacks

For example, some packet processing functions have been benchmarked up to 160 Mfps (million frames per second, using 64-byte packets) with a PCIe Gen-2 NIC.

### What it's not

Intel® DPDK is not a networking stack and does not provide functions such as Layer-3 forwarding, IPsec, firewalling, ACL, etc. Within the tree, however, various **application examples are included** to help with the development of such features.

If you need some specific drivers or networking stacks, you should **contact a company that provides such extensions**.

**Openly-accessible repository for released versions of DPDK**

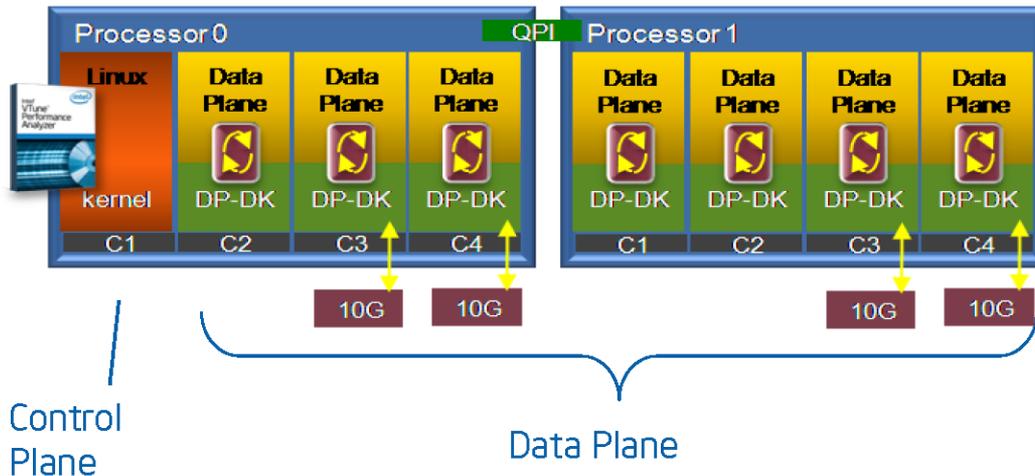
# Topics

- Introduction
- DPDK.org open-source project
- **Intel® DPDK library from 6WIND**
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- Summary

# Topics

- Introduction
- DPDK.org open-source project
- **Intel® DPDK library from 6WIND**
  - Baseline features: [Intel presentation content reused with permission from Intel](#)
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- Summary

# The Intel® DPDK Philosophy



## Intel® DPDK Fundamentals

- Implements a run to completion model or pipeline model
- No scheduler - all devices accessed by polling
- Supports 32-bit and 64-bit with/without NUMA
- Scales from Intel® Atom™ to Intel® Xeon® processors
- Number of Cores and Processors not limited
- Optimal packet allocation across DRAM channels

- **Must run on any IA CPU**

- From Intel® Atom™ processor to the latest Intel® Xeon® processor family
- Essential to the IA value proposition

- **Focus on the fast-path**

- Sending large number of packets to the Linux Kernel /GPOS will bog the system down

## Provide software examples that address common network performance deficits

- Best practices for software architecture
- Tips for data structure design and storage
- Help the compiler generate optimum code
- Address the challenges of achieving 80 Mpps per CPU Socket

# Intel® Data Plane Development Kit (Intel® DPDK)

Intel® DPDK embeds optimizations for the IA platform:

## - Data Plane Libraries and Optimized NIC Drivers in Linux

### User Space

*Queue & Buffer Management, Packet Flow Classification, Poll-Mode NIC Drivers (1/10GbE), and more!*

*Simple API Interface, Uses standard tool chain (gcc/icc, gdb, profiling tools)*

## - Run-time Environment

*Low overhead, run-to-completion model optimized for fastest possible data plane performance*

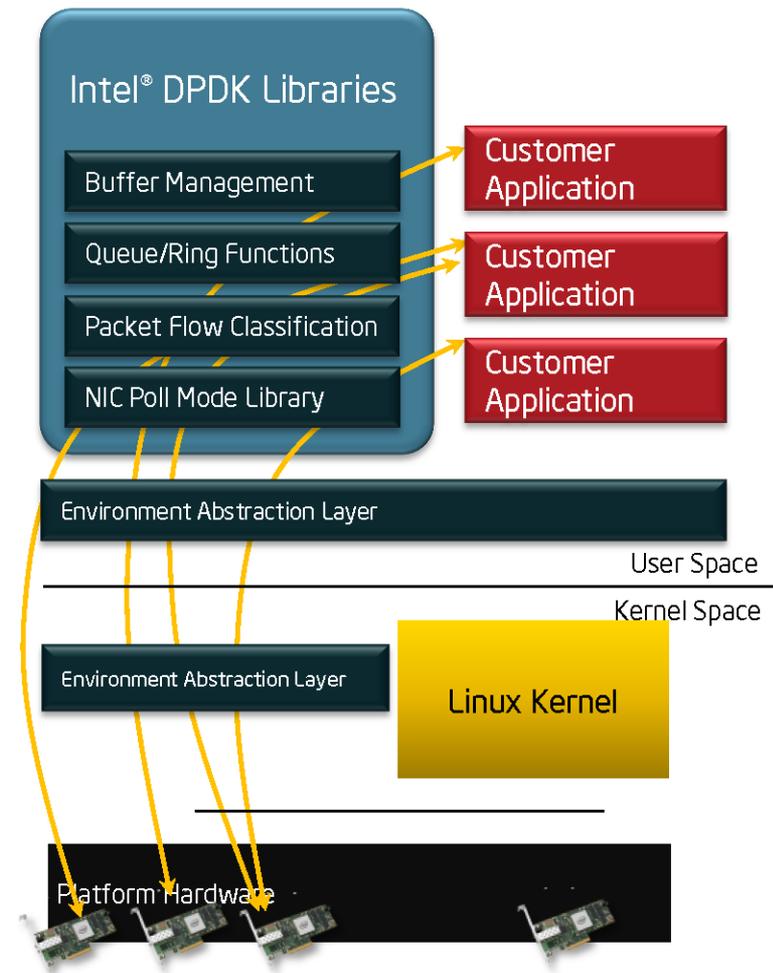
## - Environment Abstraction Layer and Boot Code

*Primarily platform-specific boot guidelines and initialization code, eases application porting effort*

## - BSD-licensed & source downloadable from Intel and leading ecopartners

*Provided under a very flexible BSD licensing model*

*Offered as a free, unsupported standalone solution by Intel or as part of commercial solutions and offerings from leading ecopartners*



The Intel® DPDK is a great starting point for customers and the industry in general - delivering breakthrough packet processing performance.

# Intel® DPDK Libraries and Drivers

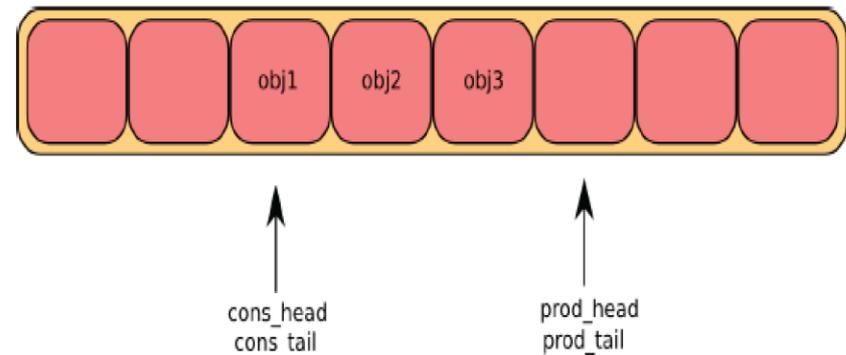
- **Memory Manager:** Responsible for allocating pools of objects in memory. A pool is created in huge page memory space and uses a ring to store free objects. It also provides an alignment helper to ensure that objects are padded to spread them equally on all DRAM channels.
- **Buffer Manager:** Reduces by a significant amount the time the operating system spends allocating and de-allocating buffers. The Intel® DPDK pre-allocates fixed size buffers which are stored in memory pools.
- **Queue Manager::** Implements safe lockless queues, instead of using spinlocks, that allow different software components to process packets, while avoiding unnecessary wait times.
- **Flow Classification:** Provides an efficient mechanism which incorporates Intel® Streaming SIMD Extensions (Intel® SSE) to produce a hash based on tuple information so that packets may be placed into flows quickly for processing, thus greatly improving throughput.
- **Poll Mode Drivers:** The Intel® DPDK includes Poll Mode Drivers for 1 GbE and 10 GbE Ethernet\* controllers which are designed to work without asynchronous, interrupt-based signaling mechanisms, which greatly speeds up the packet pipeline.

# Component Overviews

- **EAL Memory Management Overview**
- **Queue/Ring Overview**
- **Buffer Management Overview**

# Queue/Ring Management API

- **Effectively a FIFO implementation in software**
  - Lockless implementations for single or multi-producer, single or multi-consumer enqueue/dequeue
  - Supports bulk enqueue/dequeue to support packet-bunching
  - Implements high & low watermark thresholds for back-pressure/flow control
- **Essential to optimizing throughput**
  - Used to decouple stages of a pipeline



# The Buffer Management API (mempool)

- **The Buffer Manager allocates memory from the EAL and creates pools with fixed element sizes.**

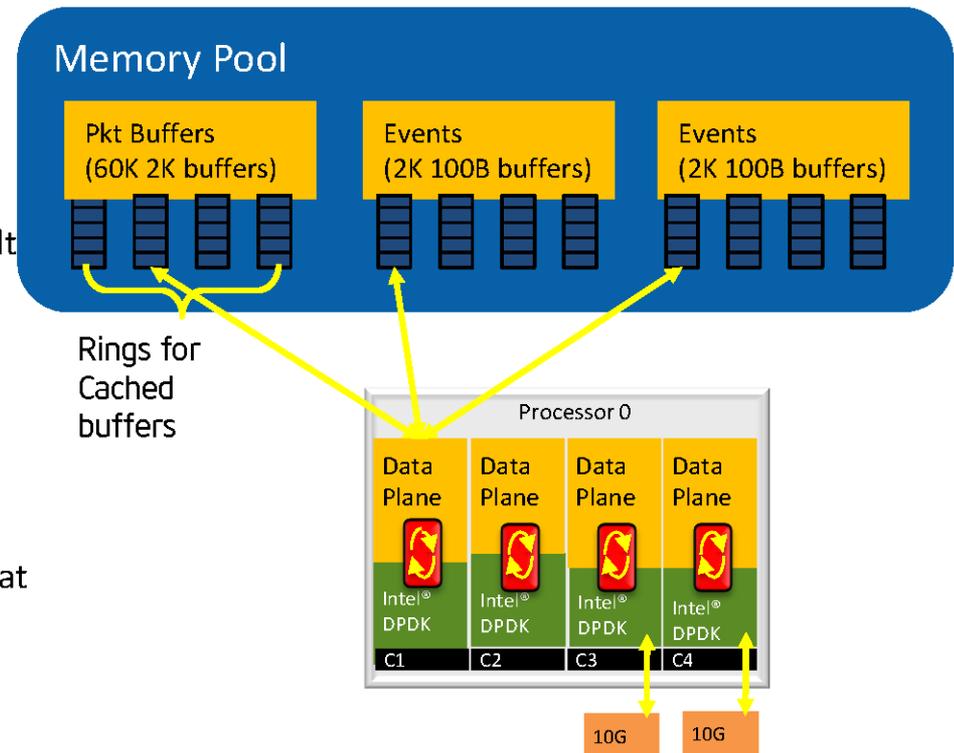
- Typical usage is packet buffers, descriptor ring buffers, etc.
- Intent is to speed up runtime allocation/de-allocation
- Does not support runtime resizing of pools

- **Multi-producer/multi-consumer safe**

- Pools are based on Intel® DPDK rings so are multi-producer and multi-consumer safe
- No locking; use CAS instructions
- Pools can also be used in multi-process environments

- **Optimized for performance**

- Cache alignment
- Per core buffer caches for each buffer pool so that allocation/freeing can be done without using shared variables
- Bulk allocation/freeing support



# Summary

Intel® DPDK enables multi-workload/single architecture potential by making IA extremely competitive for packet processing workloads

Distribution of enabling software under flexible and cost-free licensing model enabling maximum customer usability

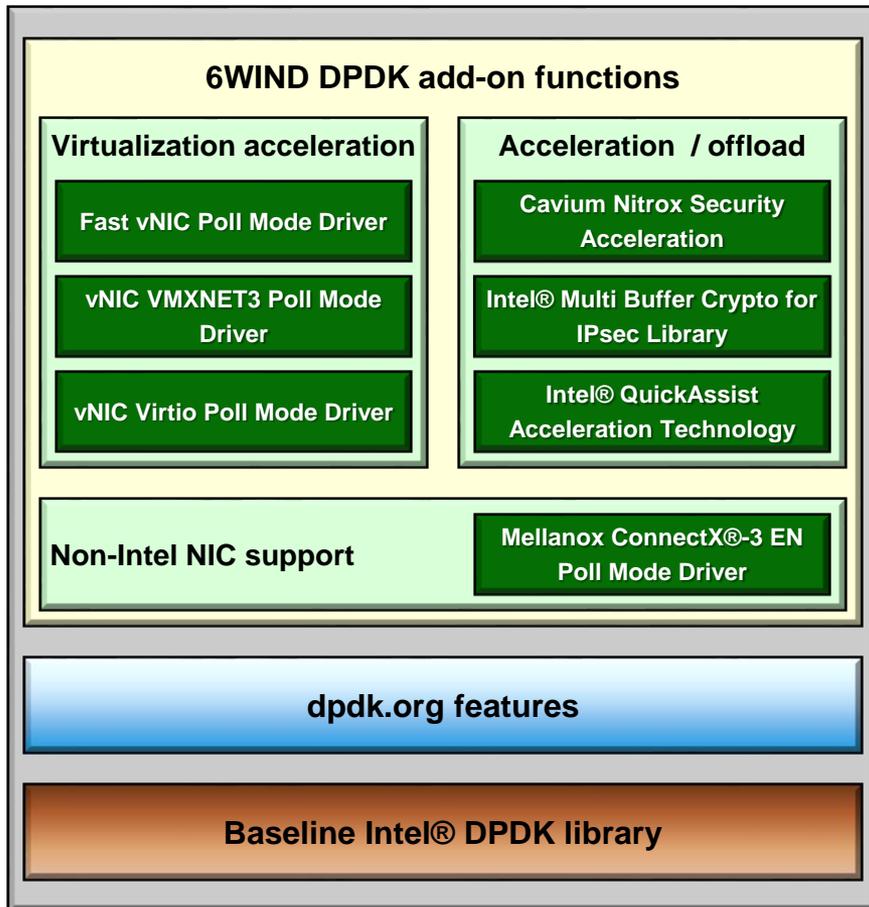
Fully featured and supported IA Data Plane software solutions via Intel's lead Ecosystem partners



# Topics

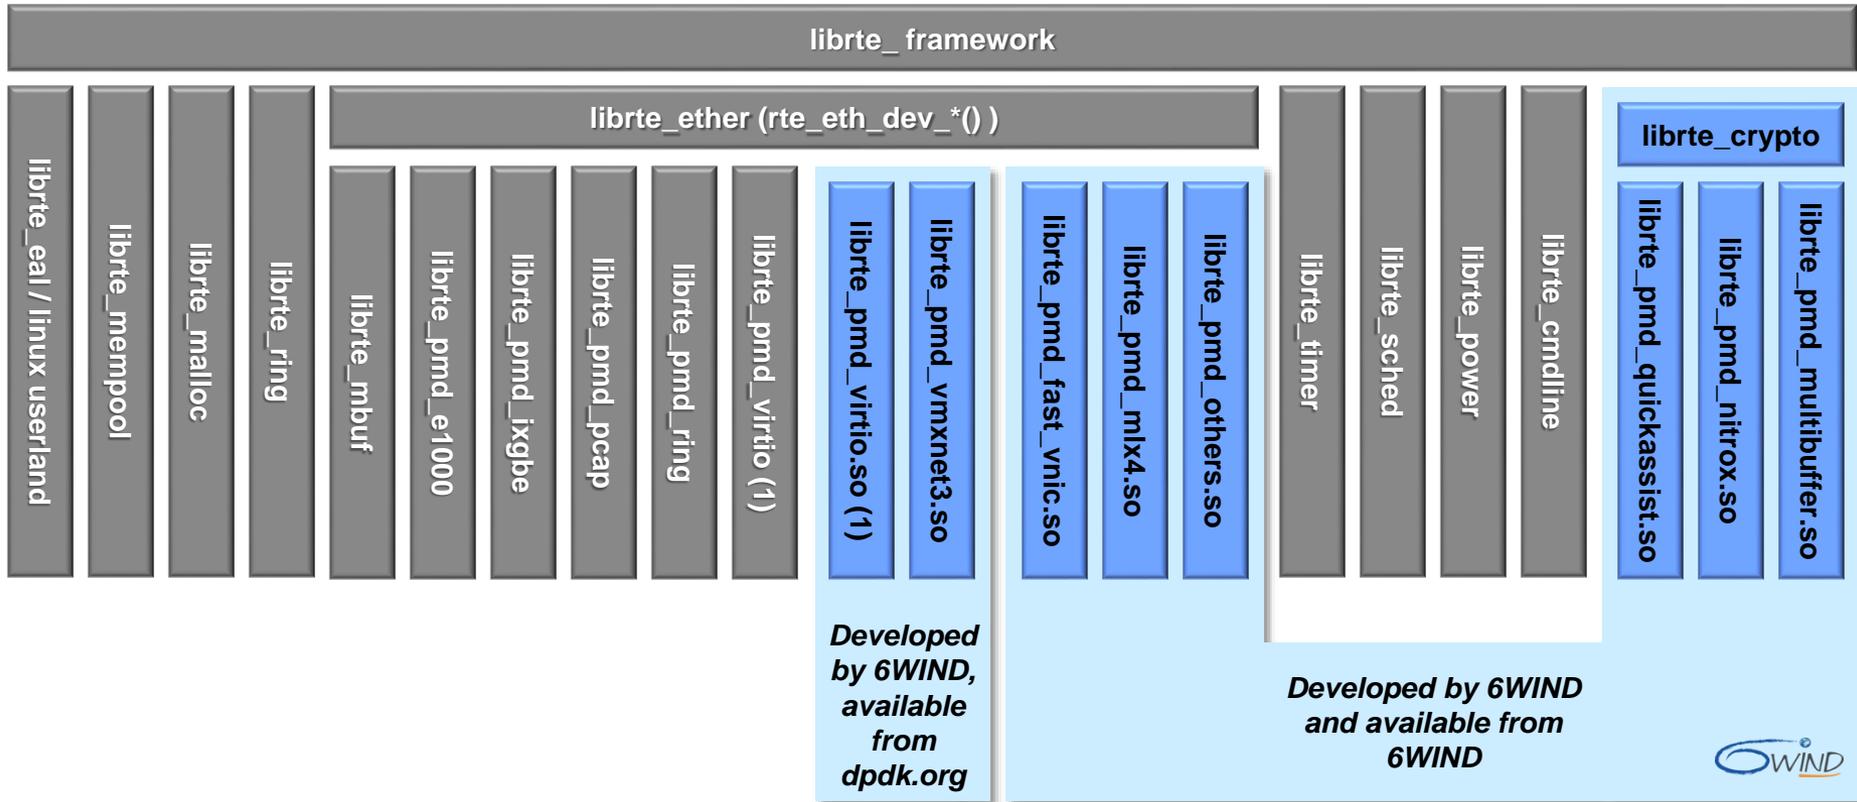
- Introduction
- DPDK.org open-source project
- **Intel® DPDK library from 6WIND**
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- Summary

# Intel® DPDK Library from 6WIND: Features and Benefits

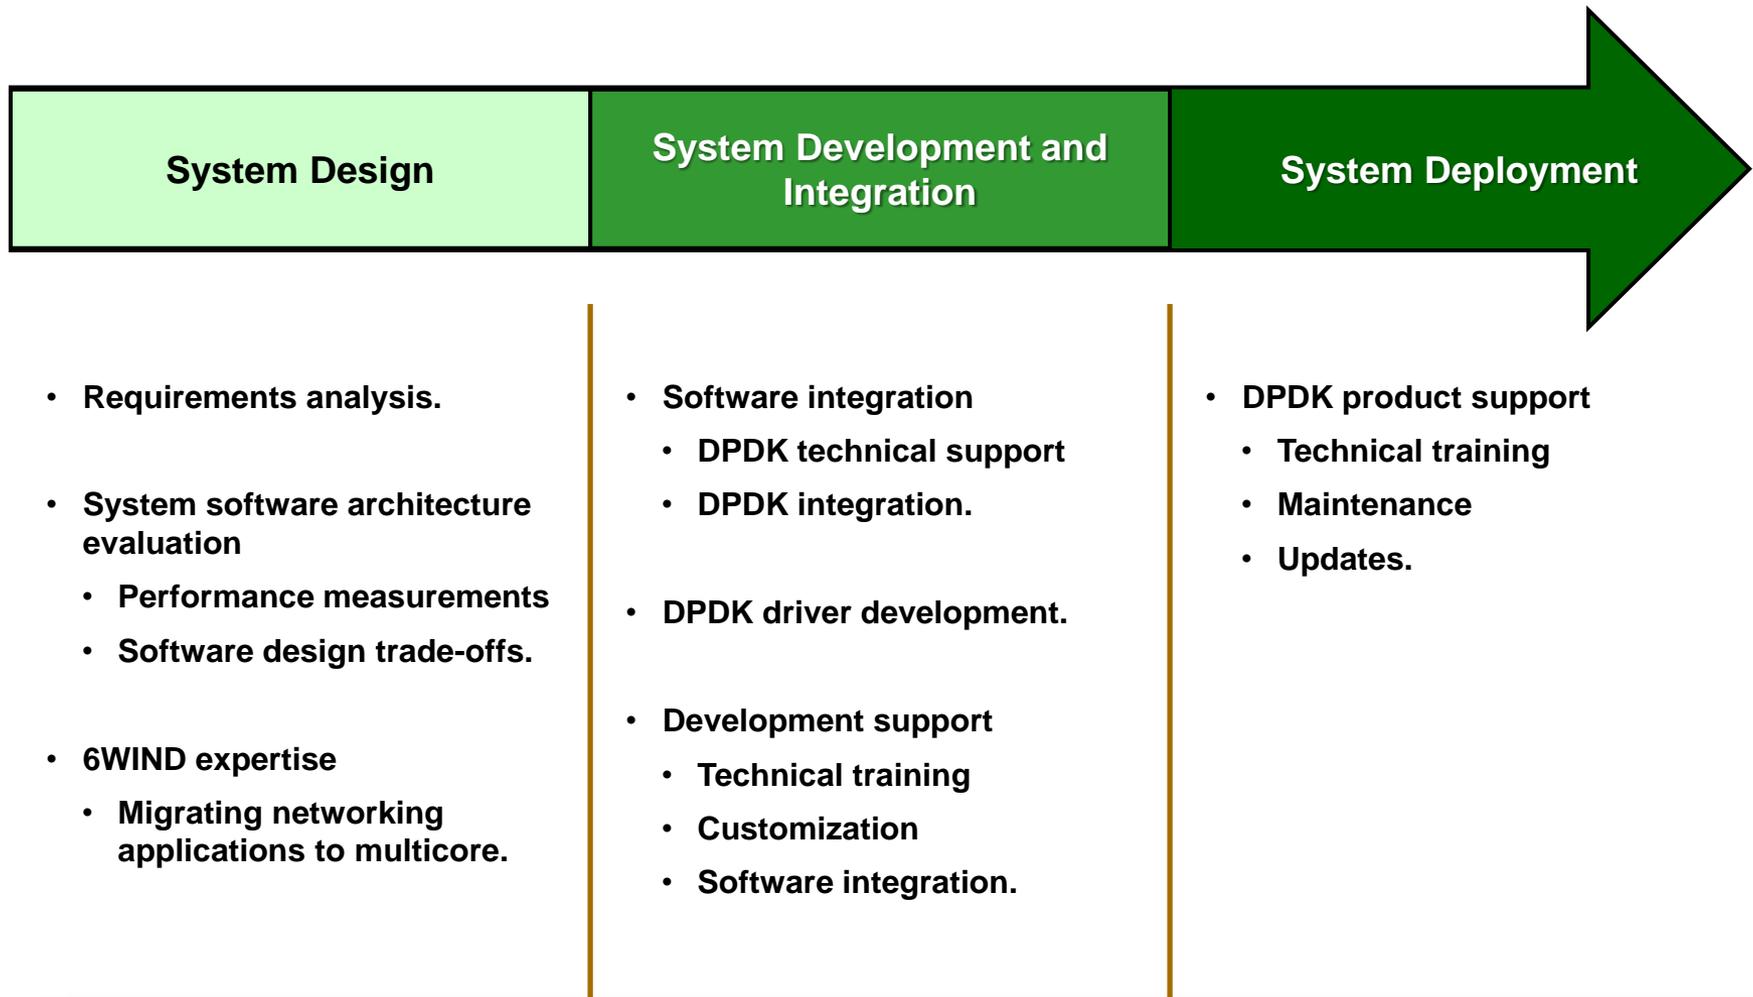


- **Based on baseline Intel® DPDK library**
  - As well as features available via dpdk.org
- **6WIND DPDK add-ons available for increased system functionality, performance and reliability**
  - Poll Mode Drivers (PMDs) for non-Intel NICs
    - e.g. Mellanox ConnectX-3® EN PMD
  - Performance acceleration for virtualized networking
    - Fast vNIC PMD
    - vNIC VMXNET3 PMD
    - vNIC Virtio PMD
  - Crypto acceleration modules that leverage
    - Cavium Nitrox security acceleration
    - Intel® Multi Buffer Crypto for IPsec Library
    - Intel® QuickAssist Acceleration Technology

# 6WIND's Enhancements to Intel® DPDK



# Intel® DPDK Library from 6WIND: Comprehensive Lifecycle Support



# Best-in-Class Technical Support

- **Industry-leading experience from many years of software development for Intel® DPDK platform**
  - Proven with tier-1 networking and telecom equipment companies worldwide.
- **Successful technical support model proven with clients worldwide in all tiers**
  - Experienced with custom support agreements for specific clients when required.
- **Standard support agreement includes**
  - SLA-defined response times
  - Online system for submitting support questions and tracking issues.
- **Dedicated service and support team**
  - Local experts based in Asia, Europe and North America
  - Partnerships with software services companies for value-added services worldwide.



# Subscription-Based Business Model for DPDK

**Note: different business model for 6WINDGate software**

## ■ Developer subscription

- Resources to help develop a product
- Cannot be used for product deployment
- Applies to defined platform configuration(s)

## ■ Quarterly releases of 6WIND DPDK

- Patches and maintenance releases
- Ongoing access to latest version of DPDK

## ■ Standard subscription

- One year duration, web interface for support, single subscriber, single supported platform, limited engineering services, on-line training

## ■ Optional additions

- Up to three years duration, regular support calls, additional subscribers, more platforms, increased engineering services, on-site training

## ■ Production subscription

- Resources to help maintain a shipping platform
- Cannot be used for product development
- Applies to defined DPDK version and platform configuration(s)
- Patches and maintenance releases
- Multiple options for number of deployed units

## ■ Standard subscription

- One year duration, web interface plus 12x5 phone support, quarterly reports, two subscribers, single supported platform, one DPDK version, on-line training

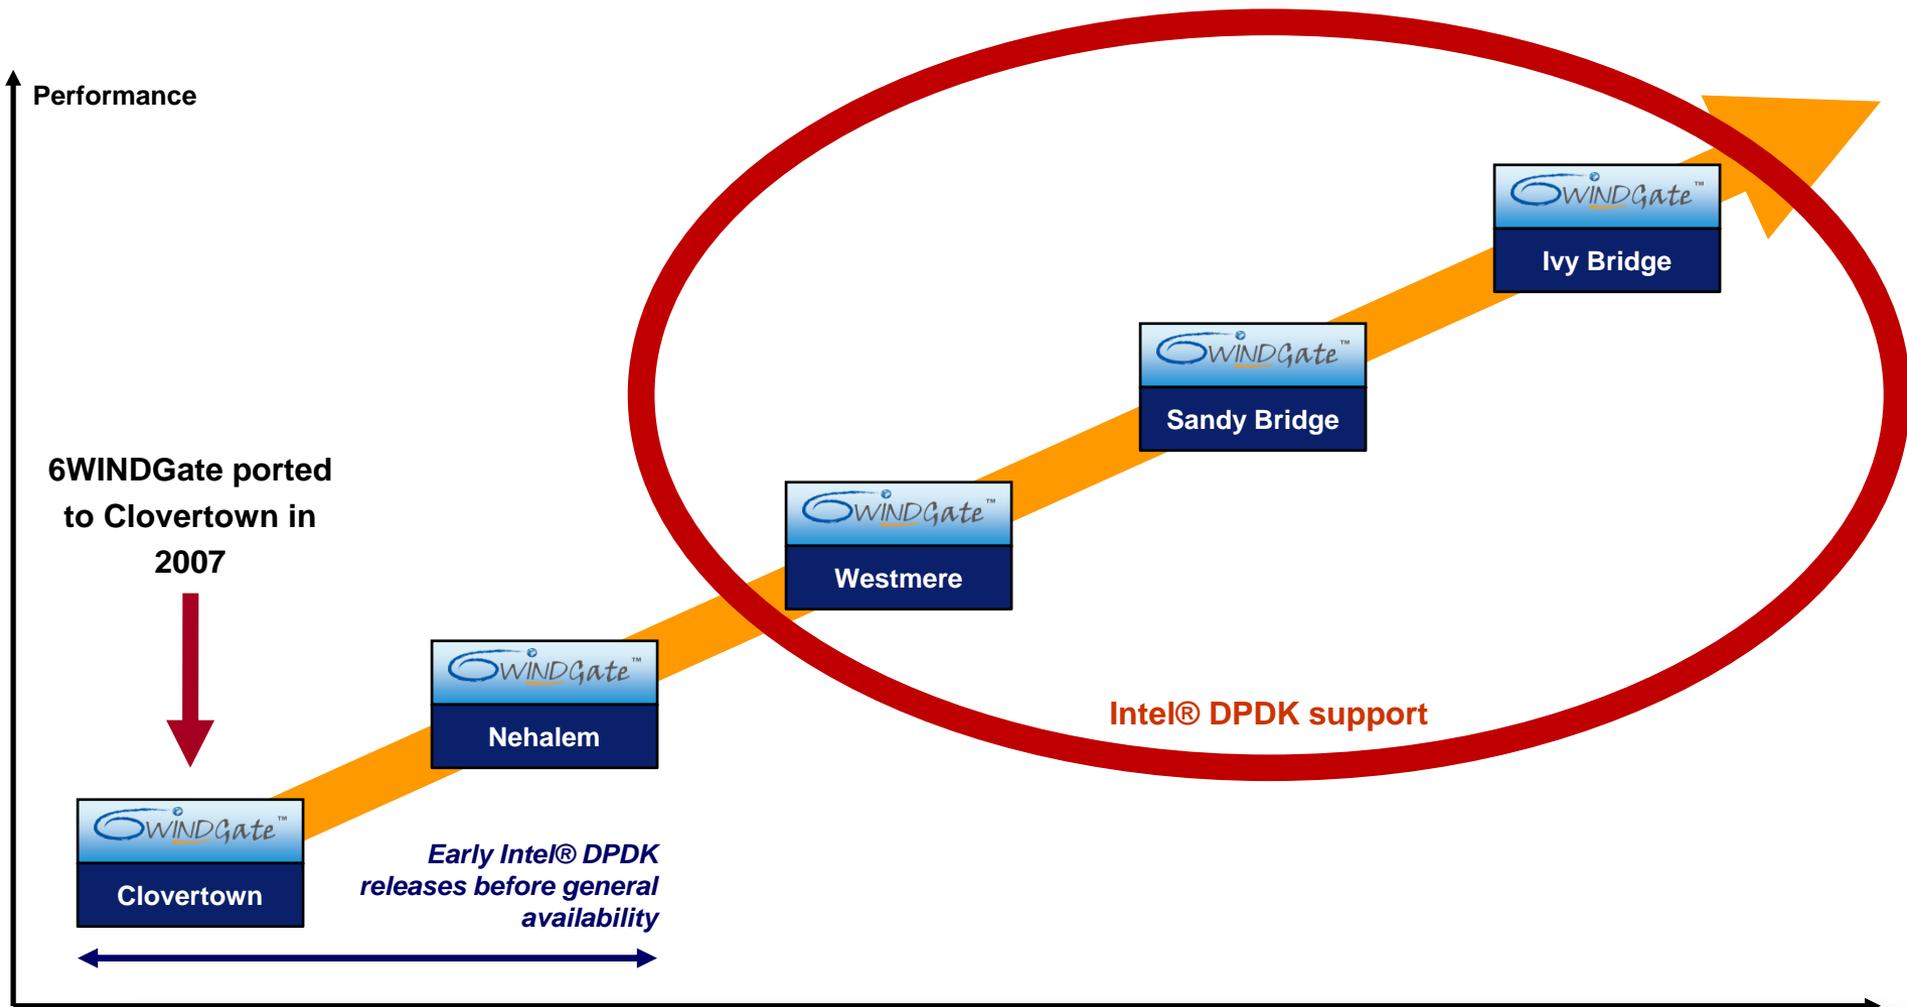
## ■ Optional additions

- Up to three years duration, web interface plus 24x7 phone support, monthly reports, additional subscribers, more platforms, more DPDK versions, engineering services, on-site training, private branch

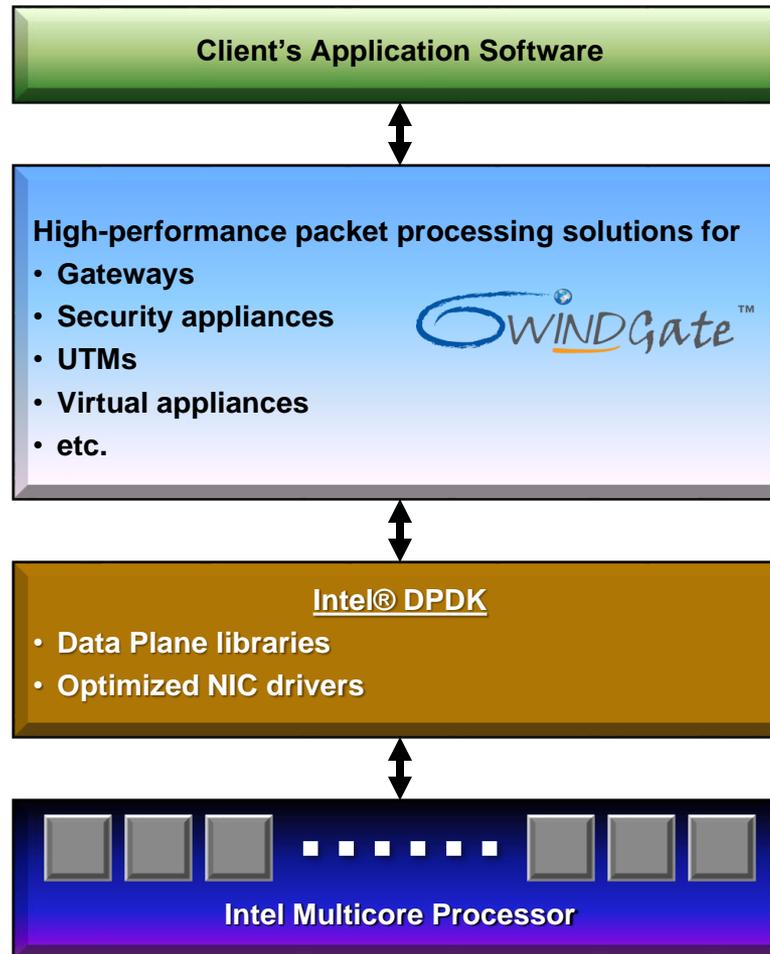
# Topics

- Introduction
- DPDK.org open-source project
- Intel® DPDK library from 6WIND
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- **Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK**
  - Features and benefits
  - Performance
  - Support
  - Use cases
- Summary

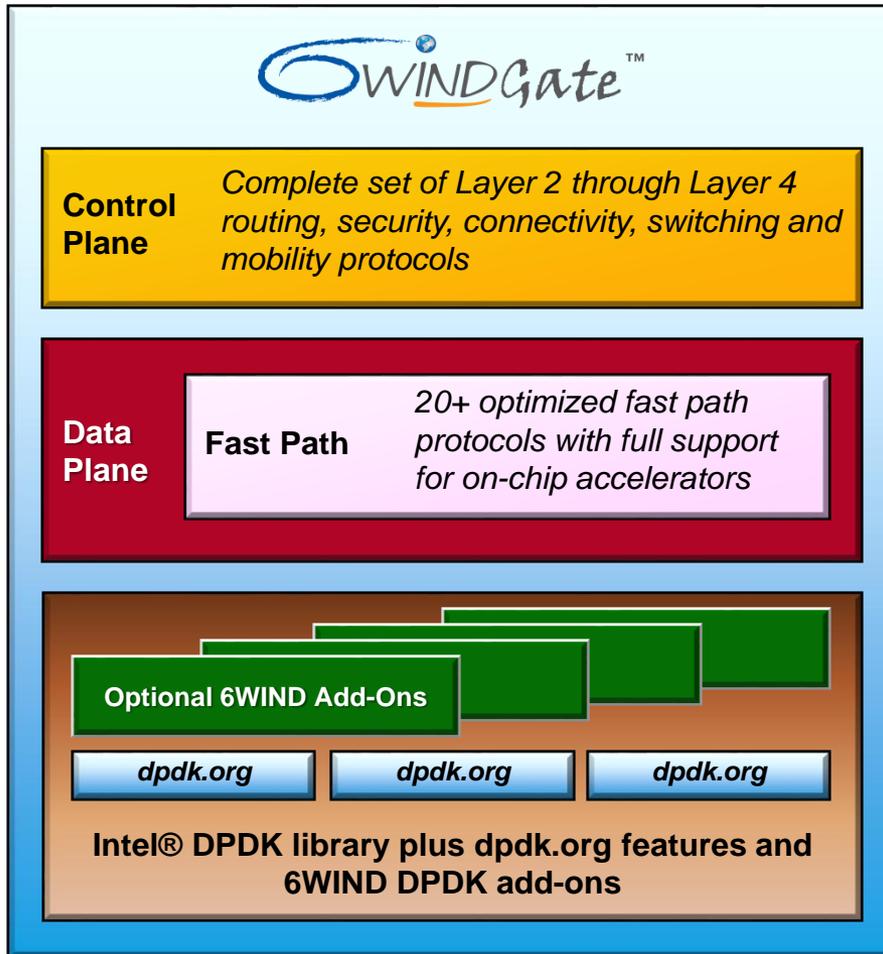
# 6WINDGate Support for Intel® DPDK



# 6WINDGate Provides System Solution Leveraging Intel® DPDK



# Optimized Integration between 6WINDGate and Intel® DPDK



- **6WINDGate provides complete packet processing solution for DPDK platforms**
  - Comprehensive set of optimized networking protocols
  - Full compatibility with application APIs
  - Carrier Grade reliability.
- **6WINDGate fully leverages Intel® DPDK library**
  - Exploits Intel® DPDK data plane libraries for best possible utilization of processor resources
  - Optimized NIC drivers for maximum networking performance.
- **Delivered as integrated solution to accelerate clients' time-to-market.**

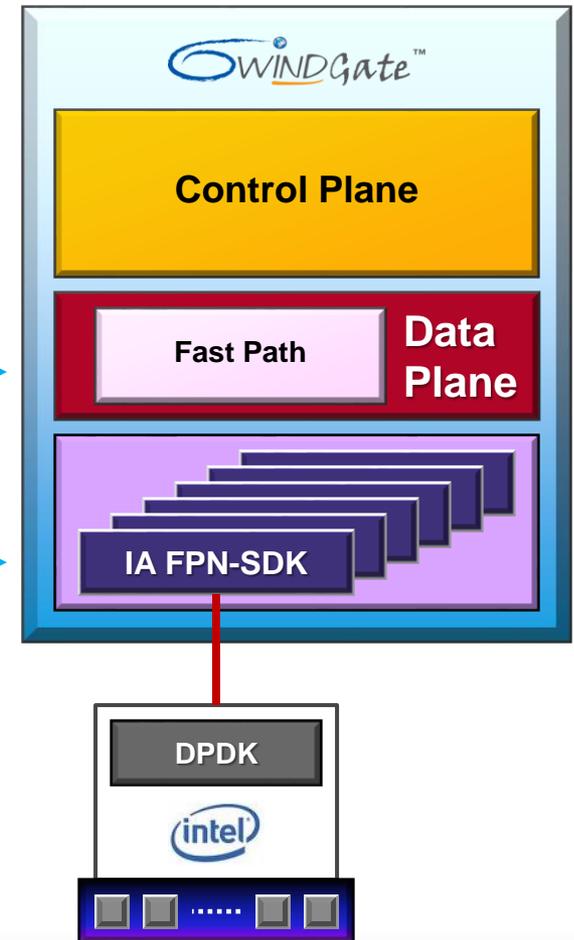
# Intel® DPDK Support Implemented via Optimized "FPN-SDK" Module

## Architecture-independent fast path modules

- Generic, processor-independent source code
- Cycle-level and pipeline-level optimizations.

## Architecture-specific "Fast Path Networking SDK"

- Zero-overhead API for fast path modules
- Support for processor-specific features and resources
- Optimized for Intel® DPDK.



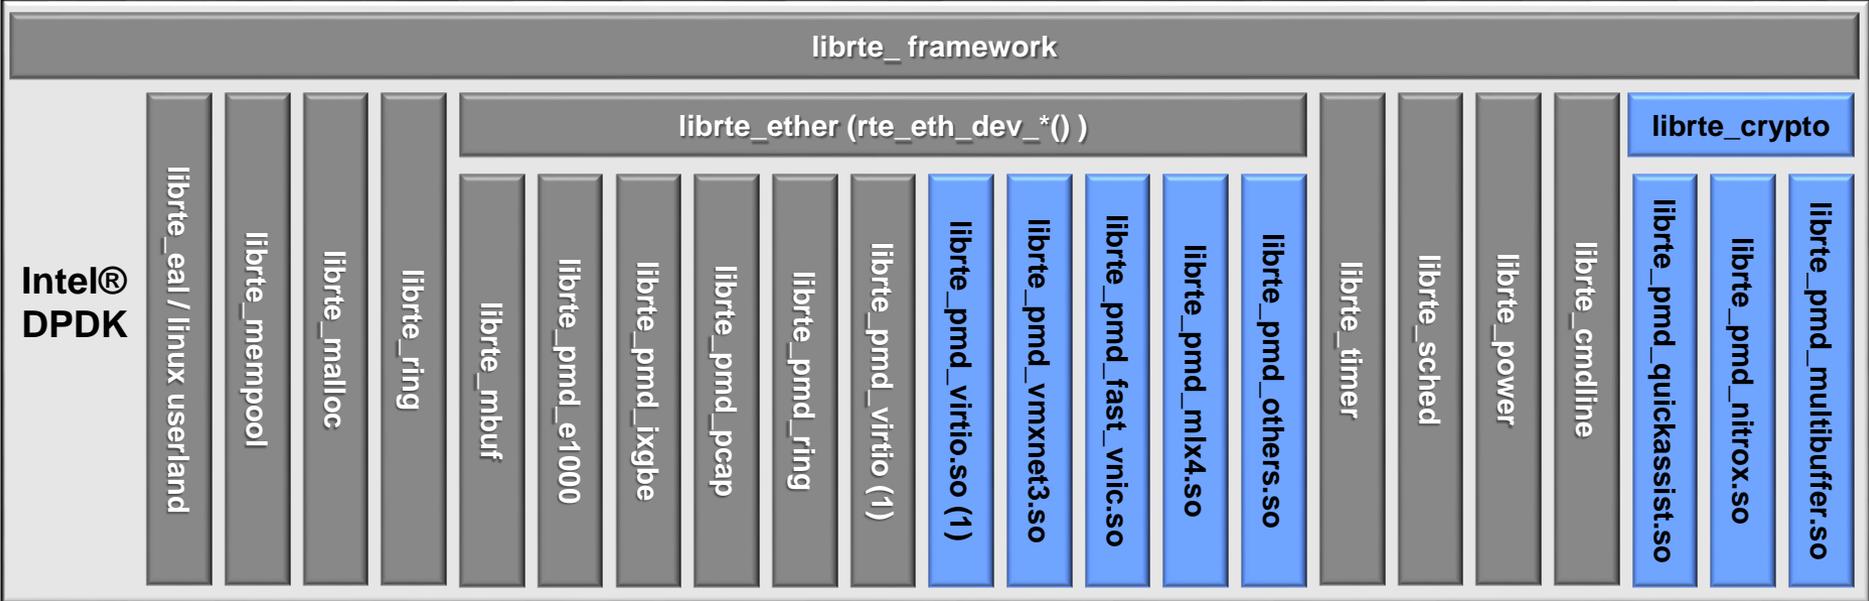
# Detailed Architecture of 6WINDGate on Intel® DPDK

Fast Path	IPv4 / IPv6 Forwarding	IPsec IPv4 / IPv6	VLAN	GRE	Ethernet Bridging, LAG	Flow Inspection (BPF)
	QoS	IPv4 / IPv6 Multicast	IP Reassembly	GTP-U	L2TP, PPPoE BRAS	TCP / UDP Termination
	IPv4 / IPv6 Filtering	MPLS / VPLS Encapsulation	NAT	IPsec SVTI	Open vSwitch (OVS) Acceleration	Extended Fast Path (1CP + xFP)
RX/TX plug-ins (load balancers, QoS, custom modules)						



FPN SDK

- rte\_mbuf → fast path buf
- rte\_timers → timers
- DPVI
- fpu RPC



# Comprehensive Protocol Set



## High Availability

HA ARP / NDP	HA Routing
HA Firewall / NAT	HA IPsec
HA Daemon Monitoring System	HA LACP

## Control Plane

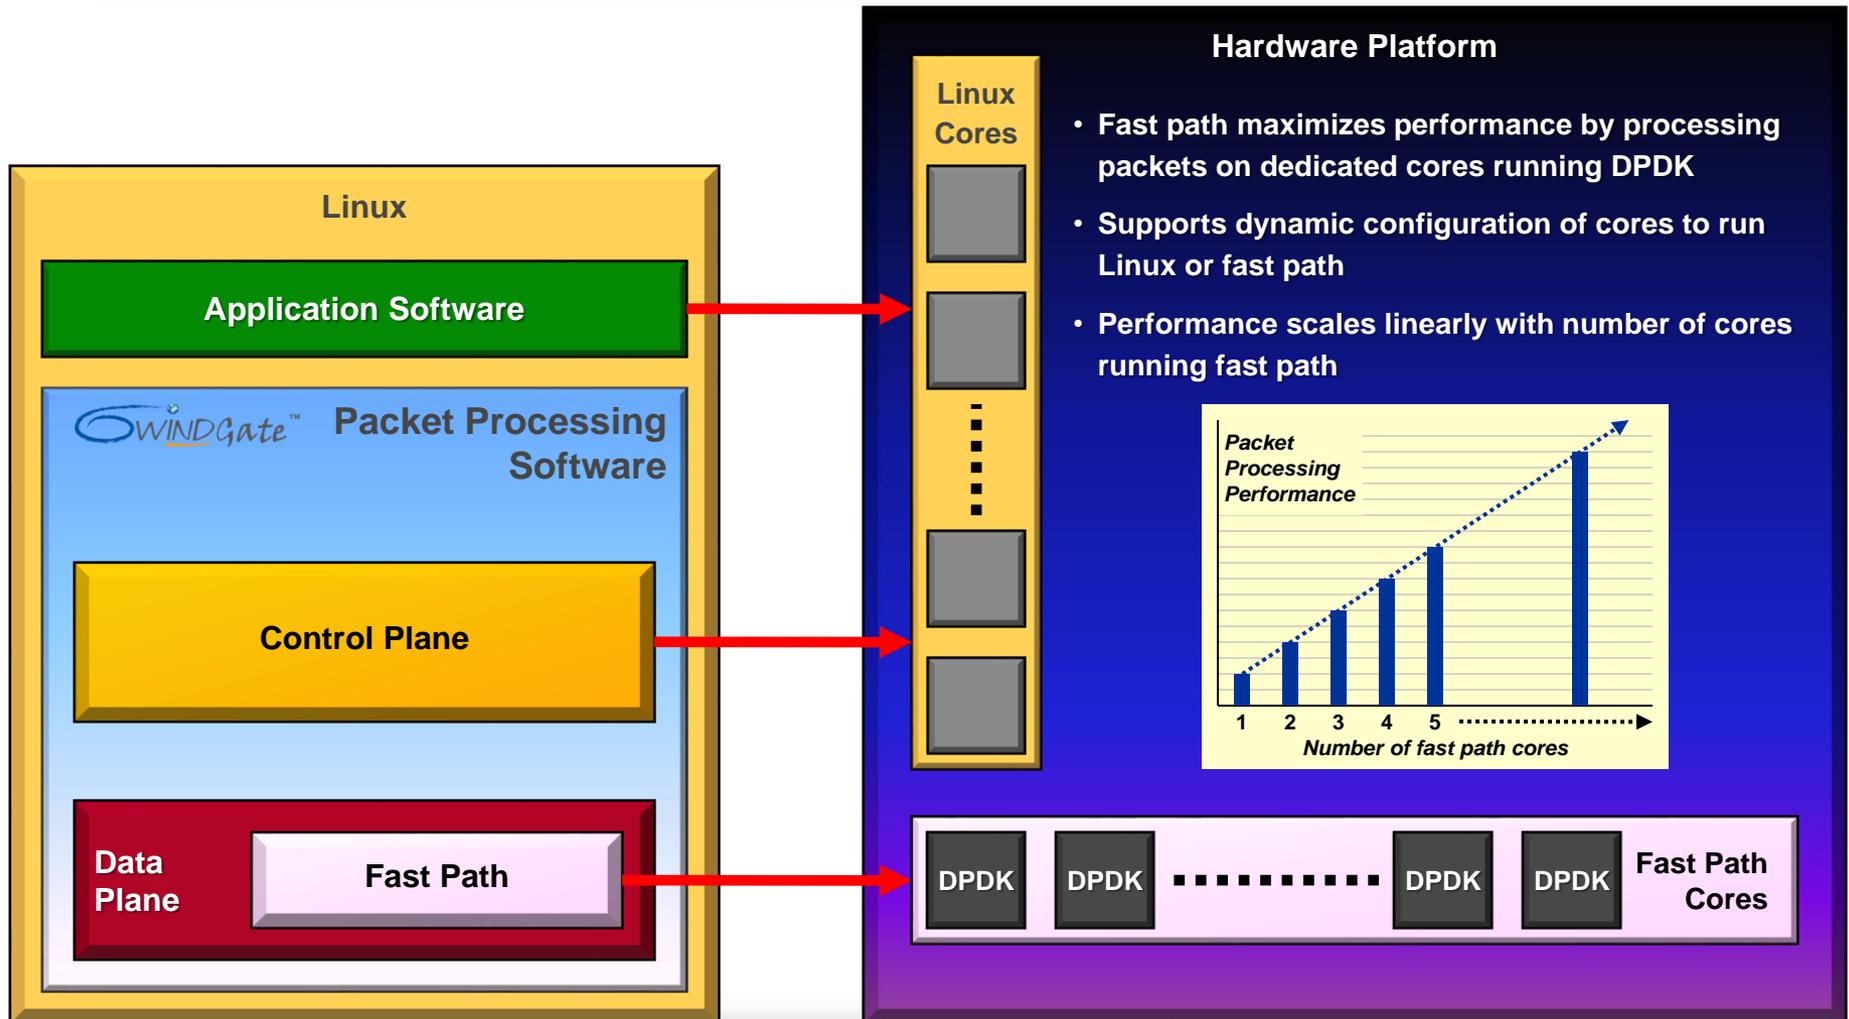
Routing (RIP, OSPF, BGP, Multicast)	Virtual Routing and Forwarding	Security (IKEv1, IKEv2)	VPN Monitoring
BFD	SMR	PPPoE BRAS	L2TP BRAS
VRRP	LACP	Open vSwitch (OVS)	Management

## Data Plane

## Fast Path

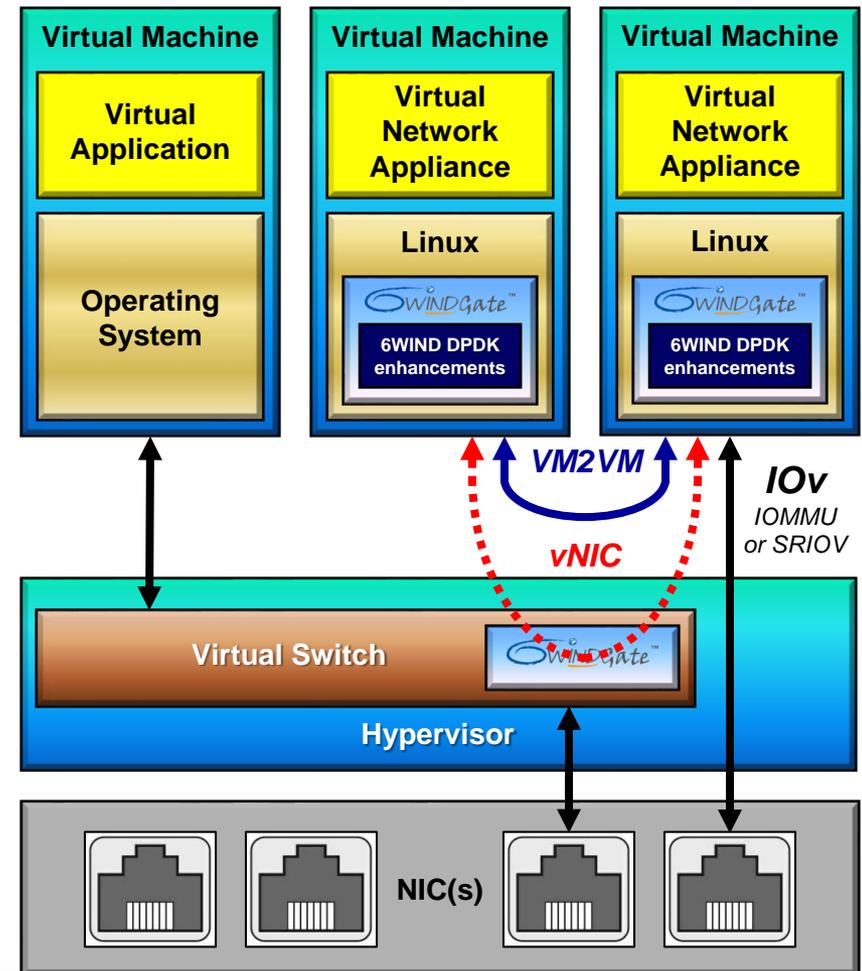
IPv4 / IPv6 Forwarding and Tunneling	IPsec IPv4 / IPv6	VLAN	GRE	Ethernet Bridging, LAG	Flow Inspection
QoS	IPv4 / IPv6 Multicast	IP Reassembly	GTP-U	L2TP, PPPoE BRAS	TCP / UDP Termination (IPv4 / IPv6)
IPv4 / IPv6 Filtering	MPLS / VPLS Encapsulation	NAT	IPsec SVTI	Open vSwitch (OVS) Acceleration	Extended Fast Path (1CP + xFP)

# Scalable across Processors, Blades and Racks



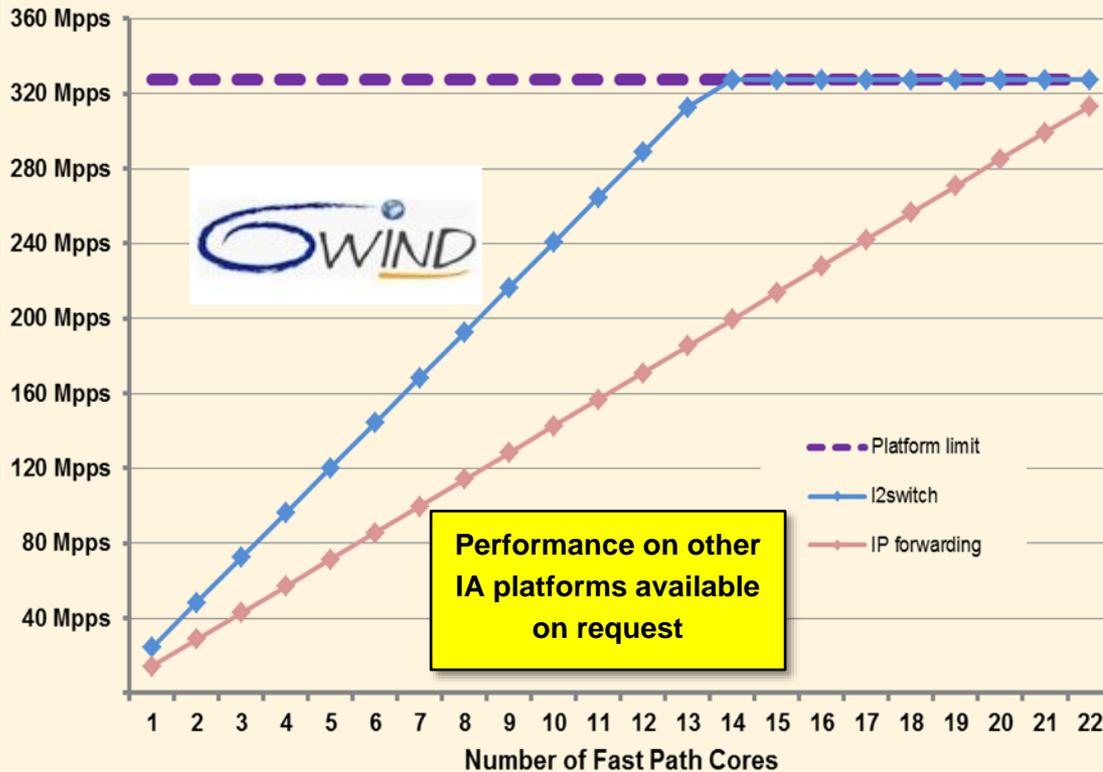
# High-Performance Networking in Virtualized Environments

- **6WIND-enhanced version of Intel® DPDK library removes VM and hypervisor bottlenecks**
  - **I/O Virtualization (IOV)** bypasses vSwitch, removing switching performance constraints
    - Supports both IOMMU and SRIOV
  - **Virtual VIC (vNIC) driver** accelerates VM-to-VM communication via vSwitch
  - **VM2VM driver** enables direct VM-to-VM communication for ultimate performance, bypassing vSwitch
- **6WINDGate accelerates Virtual Network Appliances**
  - Fully-compatible with standard hypervisors
  - Citrix XEN, Red Hat KVM, VMware ESX etc.
- **6WINDGate accelerates Virtual Switch**
  - 10x performance improvement for Open vSwitch, with no changes required to OVS code itself
  - Increases **VM density** and accelerates VM-VM traffic



# Up to 313Mpps on Dual Ivy Bridge Complex (Intel® Xeon® Processor E5-2600v2 Series)

6WINDGate IP Forwarding Performance: Ivy Bridge 3.5GHz



- **6WINDGate IP forwarding performance**
  - 14 Mpps per core
  - 24 Mpps per core with I2switch
  - Up to 313 Mpps with 22 cores (2 cores configured to run Linux)
- **Performance scales linearly with the number of cores configured to run the fast path**
  - Until it reaches the limit of the platform
- **Performance is independent of packet size**

# Comprehensive Lifecycle Support for Integrated Solution



## System Design

- Requirements analysis
- System software architecture evaluation
  - Performance measurements
  - Software design trade-offs
- 6WIND expertise
  - Migrating networking applications to multicore
  - Architecture optimizations
  - Protocols

## System Development and Integration

- Software integration
  - 6WINDGate source code license
  - DPDK technical support
  - DPDK integration
- DPDK driver development.
- Development support
  - Technical training
  - Customization
  - Software integration

## System Deployment

- 6WINDGate product support
  - Technical training
  - Maintenance
  - Updates
  - Roadmap alignment
- DPDK product support
  - Technical training
  - Maintenance
  - Updates

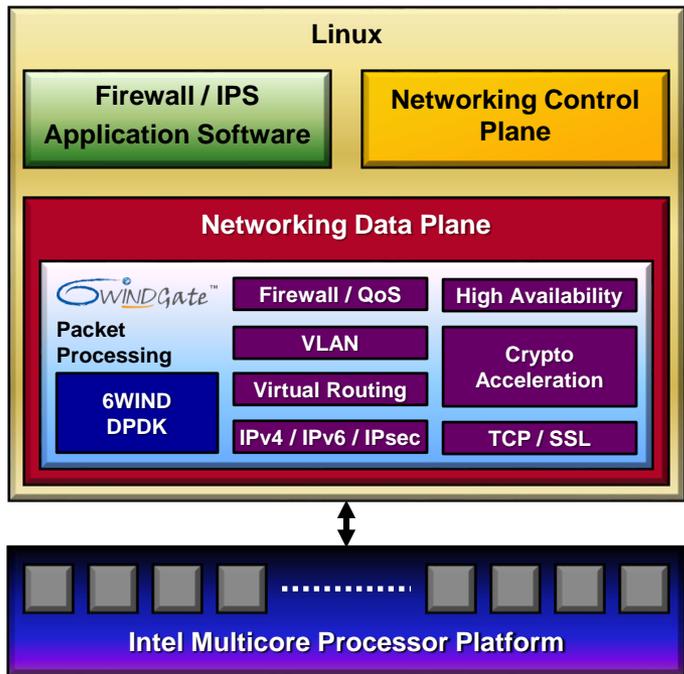
# Best-in-Class Technical Support

- **Industry-leading experience from many years of software development for Intel® DPDK platform**
  - Proven with tier-1 networking and telecom equipment companies worldwide
  - 6WIND provides 6WINDGate pre-integrated with Intel® DPDK to accelerate clients' time-to-market.
- **Successful technical support model proven with clients worldwide in all tiers**
  - Experienced with custom support agreements for specific clients when required.
- **Standard support agreement includes**
  - SLA-defined response times
  - Online system for submitting support questions and tracking issues.
- **Dedicated service and support team**
  - Local experts based in Asia, Europe and North America
  - Partnerships with software services companies for value-added services worldwide.



# Use Case Example: Firewall / IPS

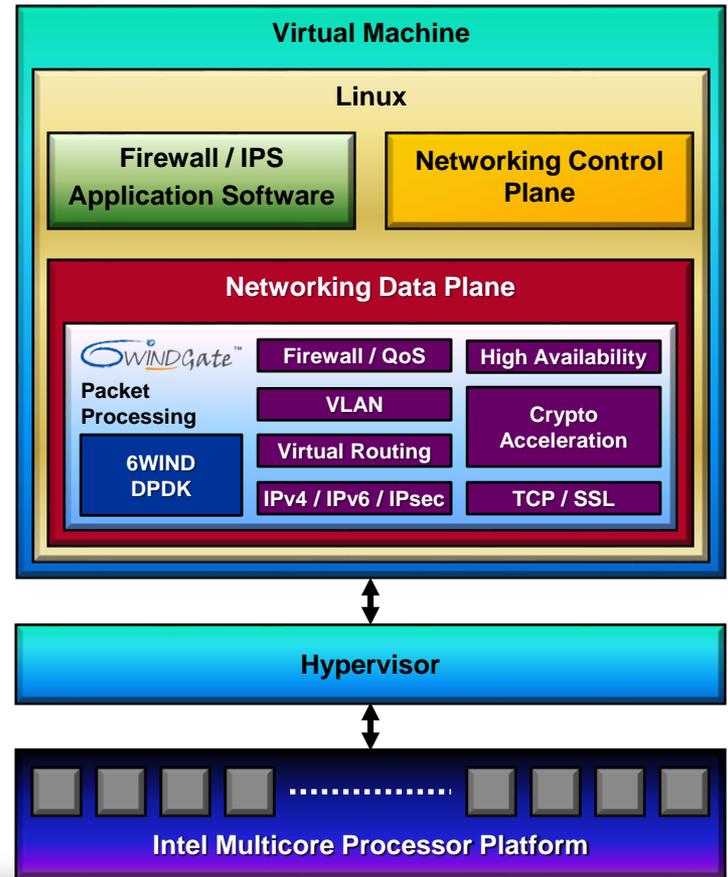
## Physical Appliance



- Leverages high-performance 6WINDgate data plane protocols plus optimized NIC drivers in 6WIND DPDK

- Virtualization acceleration in 6WIND DPDK maximizes performance of virtual appliance

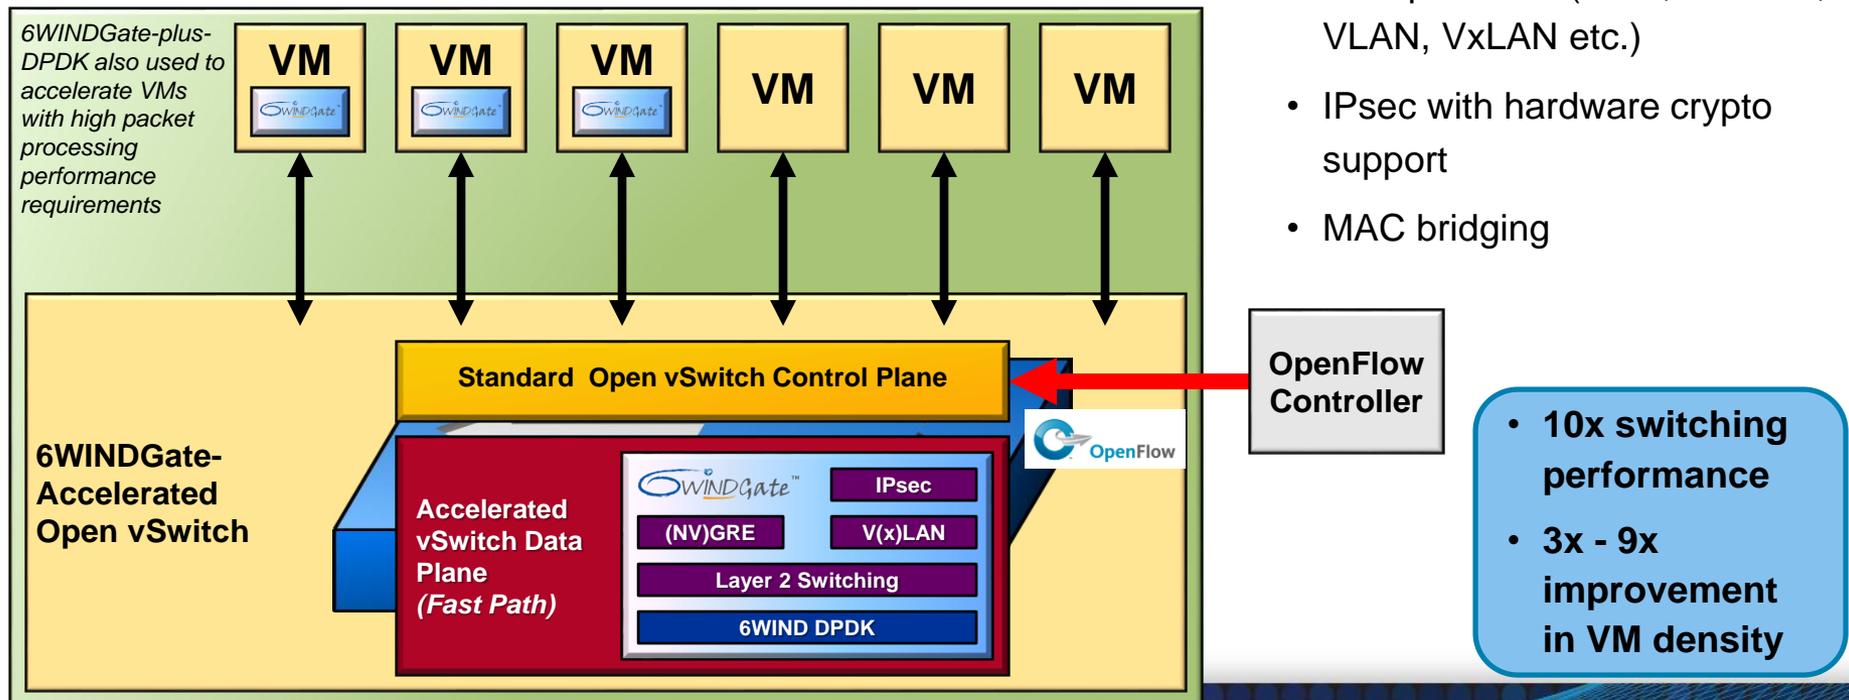
## Virtual Appliance



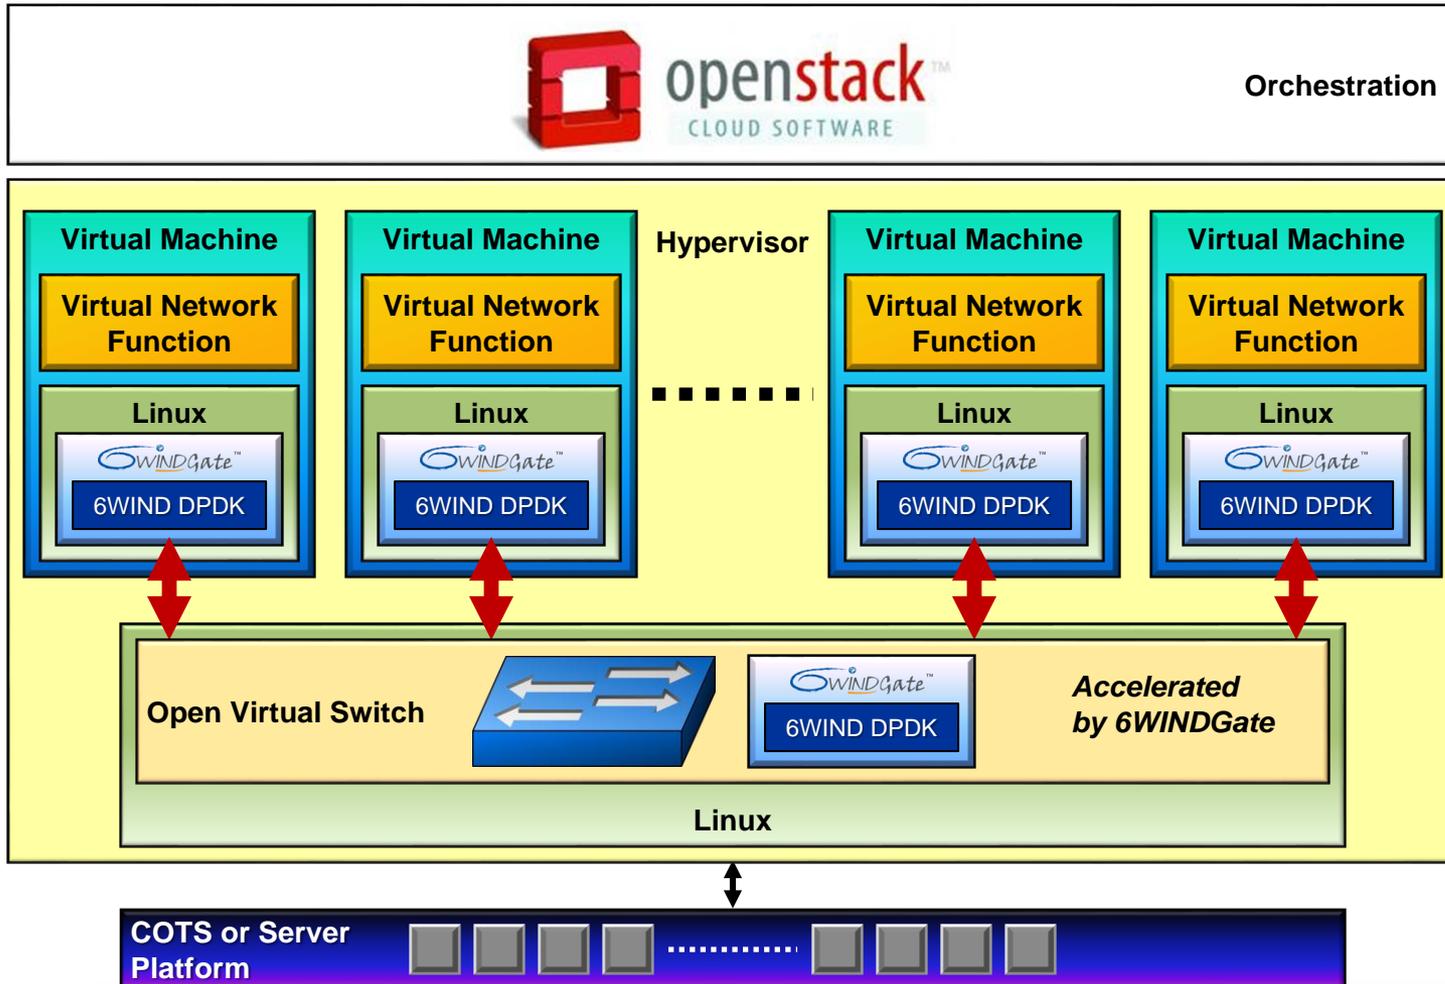
# Use Case Example: Accelerated Virtual Switch for Data Centers

- 6WINDGate-plus-DPDK accelerates Open vSwitch
  - No modification to OVS software
  - Transparent to applications

- Layer 2 switching
  - OpenFlow-compatible
- Additional networking protocols
  - Encapsulation (GRE, NVGRE, VLAN, VxLAN etc.)
  - IPsec with hardware crypto support
  - MAC bridging



# Use Case Example: Network Functions Virtualization (NFV)



- 6WINDGate-plus-DPDK accelerates
- Virtual Networking Functions (VNFs) running in VMs
- **and**
- Virtual Switch responsible for switching traffic to VNFs
- Enables VNFs to deliver comparable performance to physical implementations

# Topics

- Introduction
- DPDK.org open-source project
- Intel® DPDK library from 6WIND
  - Baseline features
  - Optional add-ons
  - Support
  - Business model
- Integrated packet processing solution: 6WINDGate™ plus Intel® DPDK
  - Features and benefits
  - Performance
  - Support
  - Use cases
- Summary

# For More Information

Solution White Paper

## Optimized Data Plane Processing Solutions using the Intel® DPDK

### Executive Summary

End-user usage patterns and expectations are creating challenges for both service providers and network equipment suppliers with respect to network capacity, system-level cost-performance, the delivery of advanced network services and time-to-revenue. These problems can be effectively addressed only at the platform level, as opposed to the component level. This white paper describes how Intel and 6WIND\* are working together to address these key telecom platform challenges through a unique combination of the latest processors and software. This close technical collaboration has delivered a highly-optimized platform with advanced features and capabilities for next-generation networking equipment. The integrated platform, described in this paper, enables service providers to reduce their CAPEX and OPEX, while delivering the advanced services that are critical to subscriber acquisition and retention. Moreover, network equipment providers can take advantage of the platform in order to accelerate their time-to-market for LTE solutions that provide industry-leading cost-performance.

**White Paper**



www.intel.com/go/commsinfrastructure



www.6wind.com



The image shows the 6WIND website homepage. At the top, there is a navigation bar with links for Solutions, Products, Software Defined Networking, Resources, About Us, and Support. The main header features the 6WIND logo and the tagline "Solving the Critical Performance Challenges for Software Defined Networks". Below this, there are several content blocks: "Data Center Networking" with a sub-headline "Accelerated virtual switches for multi-tenant environments" and a yellow "Website" button; "Mobile Infrastructure"; "Network Appliances"; and "Data Center Networking". A central banner highlights "www.6wind.com". Other sections include "Latest From 6WIND", "6WINDGate", "Videos", "Our Customers", and "Hitachi".



The image shows a blog post on the 6WIND website. The page title is "Virtual Switch Acceleration Boosts VM Density in Data Centers". The author is Charlie Ashton, VP of Marketing and Business Development at 6WIND. The post discusses the 6WINDGate™ networking software that delivers a 10x performance improvement for virtual switches. A yellow "Blog" button is overlaid on the text. Below the main article, there are sections for "RECENT POSTS", "RECENT COMMENTS", and "ARCHIVES". The URL "www.6windblog.com" is highlighted in a yellow box.

# Summary: Right Architecture + Right Software = Winning Solution

- **6WIND completes the system solution for IA platforms**
  - Full set of optimized protocols for networking and security products.
- **6WIND delivers the maximum performance for IA platforms**
  - Fully exploits IA processor architecture and Intel® DPDK.
- **Packaged solution accelerates clients' development cycles**
  - 6WINDGate delivered integrated with Intel® DPDK.
- **Stand-alone Intel® DPDK library available for custom development**
  - Value-added feature enhancements and add-ons.
- **Expert technical support for all Intel® DPDK users**
  - From the company with the industry's best experience of developing software for Intel® DPDK platform.

