Embracing High Speed, Low Power, Complex Security Analytics at the Heart of the Cloud

Sakir Sezer
Chief Technical Officer

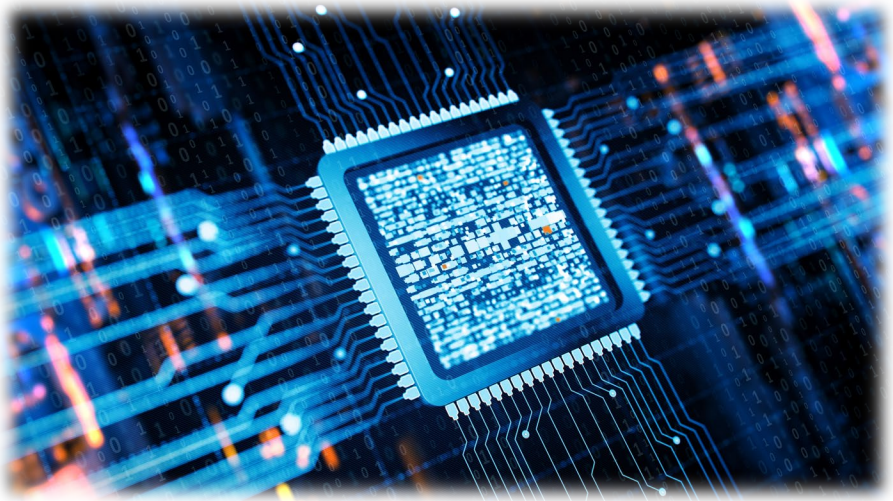

IP-SoC 2019 Conference
December 3rd, 2019

Overview

Titan IC

- Evolution of Datacenter Server Architecture
- SmartNIC - Definition and Architecture
- Challenges of enabling defensive cybersecurity within the cloud
- Enabling Smarter Security on SmartNICs with RXP
- SmartNIC security use cases
- Challenges and opportunities

Evolution of Datacenter Architectures

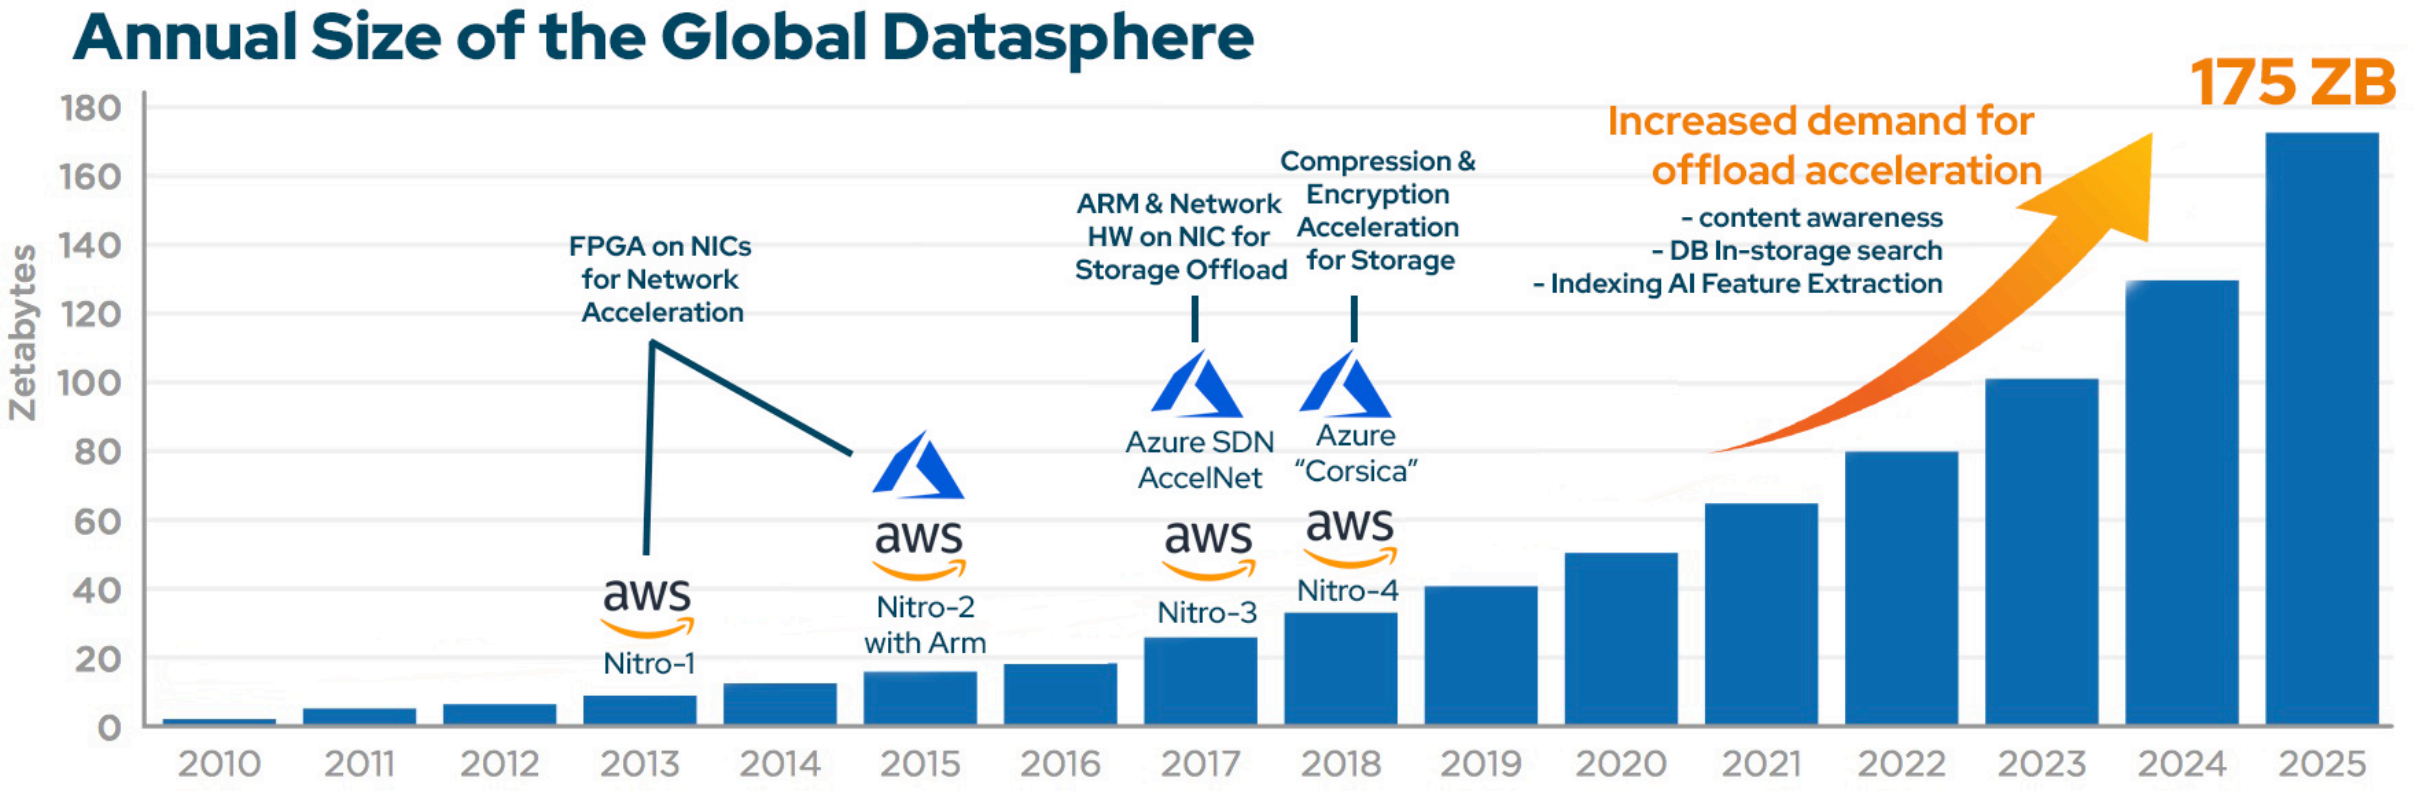

Source: NIC Offload acceleration AWS/Azure – Hotchips 2019

Next Generation Datacenter Architectures

And the evolution of SmartNIC technology

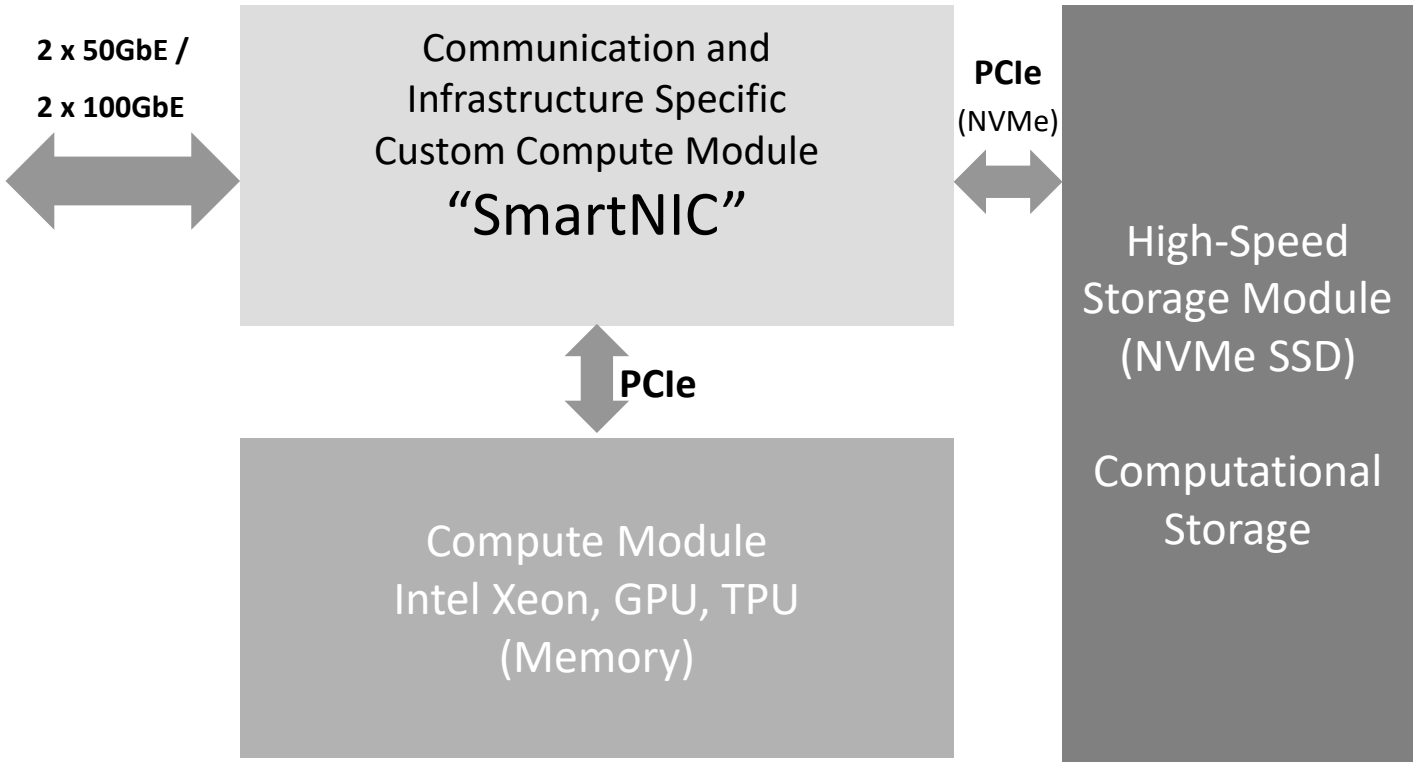

- Driver: Effective utilization of Compute and Storage resources
- AWS Nitro System: deploying custom-purpose accelerators at the NIC
- Azure AccelNet SmartNIC: standard NIC with FPGA for OVS offload & Corsica
- SmartNIC offload trends:
  - Virtual Switching (OVS)
  - Security (VPN)
  - Data Compression / Decompression (ZipLine)
  - Data encryption/decryption
  - Storage Control (file/block/object)
  - Infrastructure control
  - Database Acceleration
  - Computational Storage…

Key Features Defining a SmartNIC

Titan IC

- **Evolution of Datacenter Technology**
  enabling efficient virtualization
- **On-device Processing of Upper-Layer Functions**
  enabling semi-autonomous decision-making
- **On-device Acceleration**
  offloading heavy-duty tasks such as encryption, switching, inspection etc

Common SmartNIC Architectures

Xilinx ALVEO: Fully logic programmable Smart NICs. Enables the customization of all network and application layer functions to achieve the best performance for a given use-case.

Mellanox Innova II: Semi-programmable Smart NICs. Combines highly efficient standard NIC technology with programmable logic for customization of critical network and application layer functions.

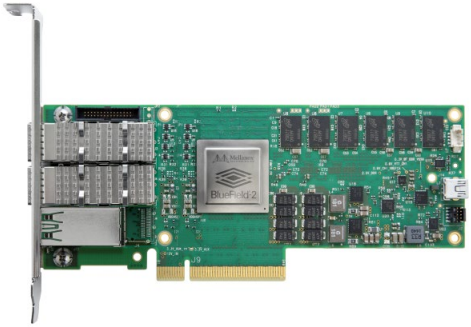

Mellanox BlueField 2: Software programmable Smart NICs. Combines embedded high-performance 64-bit processors (8 to 16 x 64-bit ARM cores) and performance optimized offload accelerators for network and application layer functions.

Titan IC RXP (RegEx offload processor IP) is highly optimized for all three SmartNIC architectures

Traditional Network Security

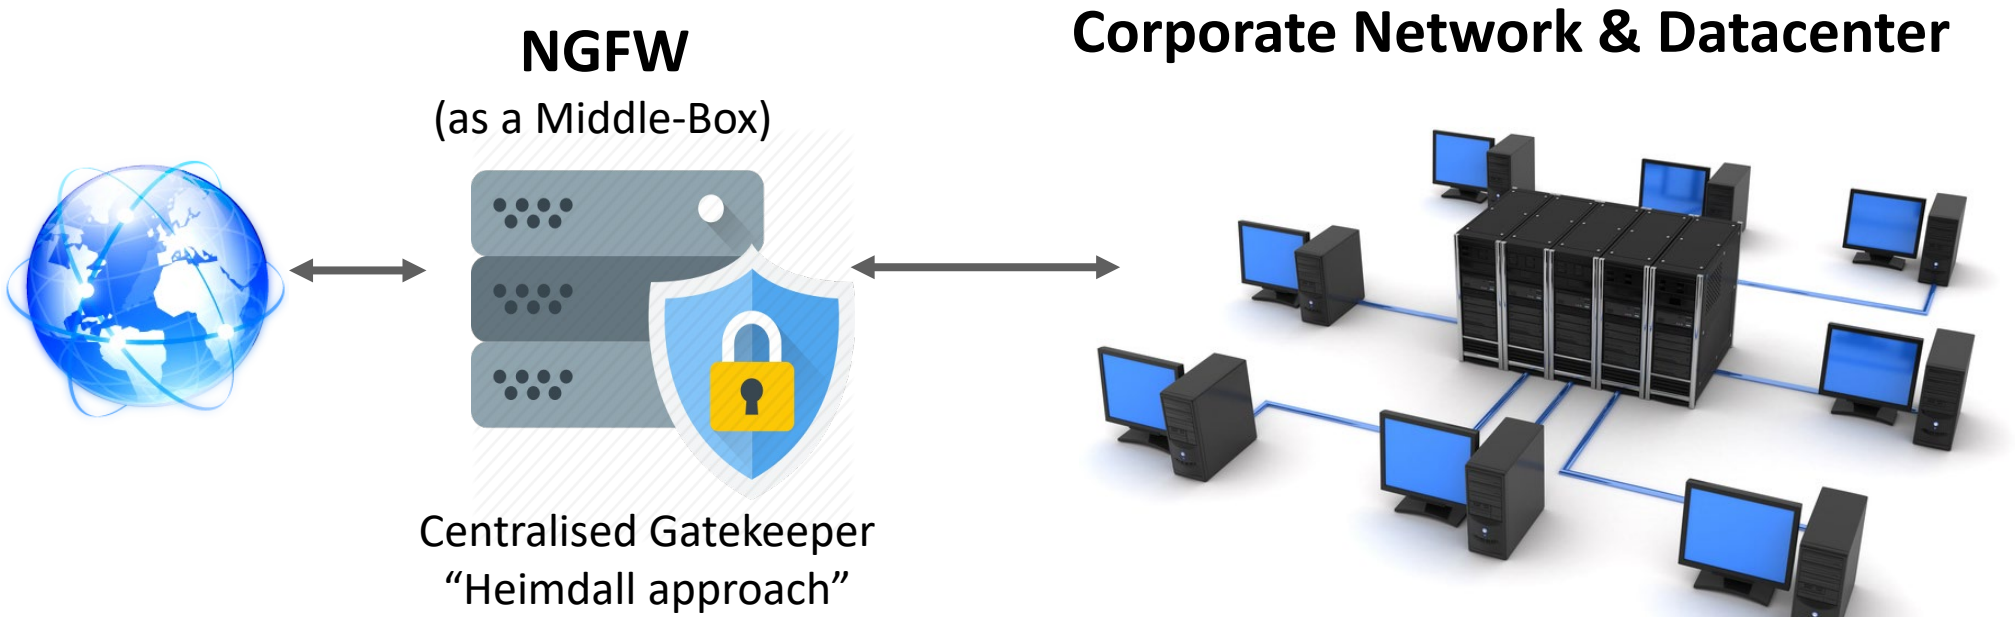

- Centralised, difficult to scale
- Locked to one specific vendor
- Vulnerable to vendor specific DDoS attacks
- Cannot be easily extended into the cloud
- Single point of failure

Cloud Enabling Network and Application Security

Titan IC

Virtual Security Appliance
AWS Instance + AWS Marketplace Apps

Instance + Apps

TREND MICRO

elastic

MySQL

CentOS

SaaS Client
(Enterprise, SME)

Security becomes a heavyweight inefficient software-based virtual appliance

Virtual Security Appliance
Cloud-based Security Middle-box Model

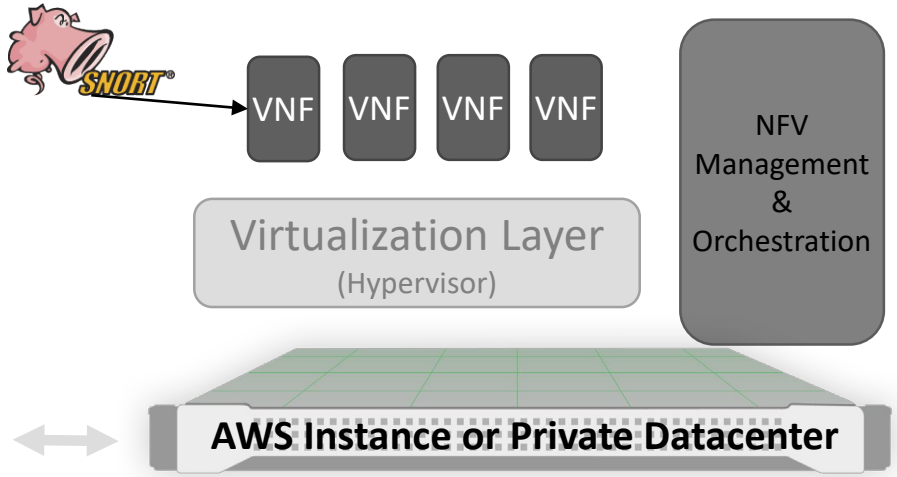

Network Function Virtualization
Virtualized Security Middle-box Model

Centralized vs SmartNIC based Network Security

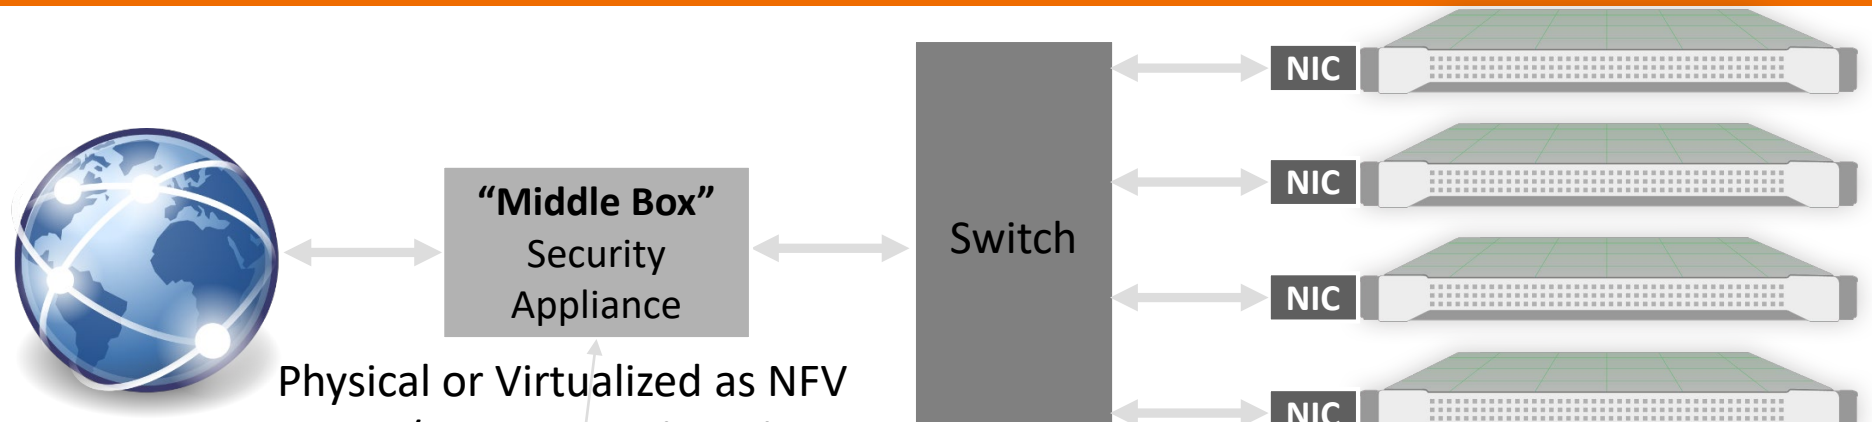

Physical or Virtualized as NFV
or AWS/Azure Virtual Appliance

Security Management
ArcSight splunk> QRadar

Security is an embedded function and integral part of a NIC, customized for the applications on the server

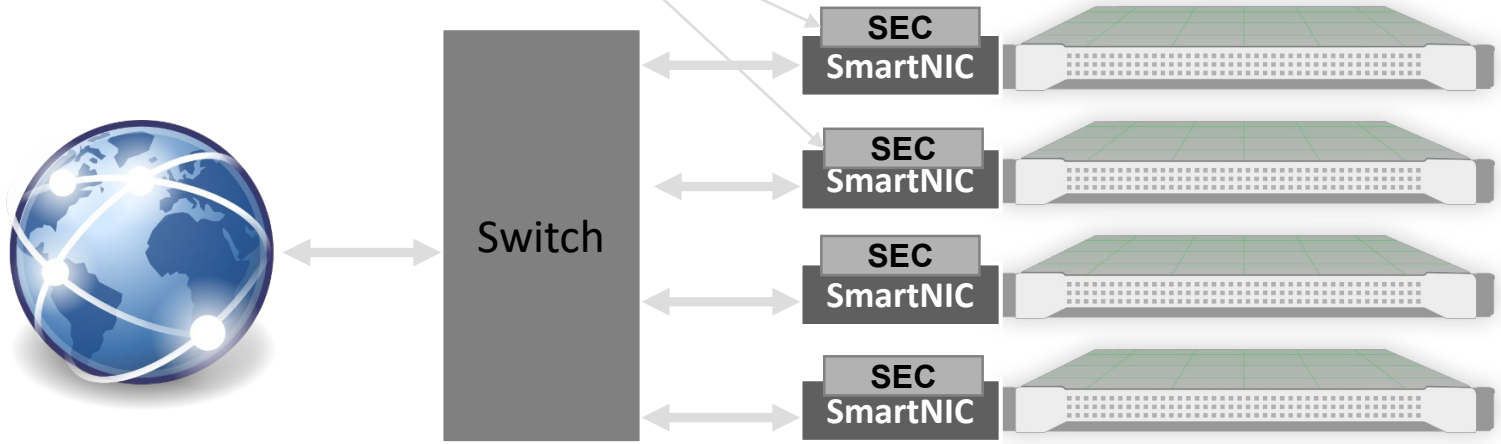

Key Advantages

- Distributed, inherently resilient
- No single point of failure
- Smaller attack surface
- Tailored to the application
- Fully virtualizable without the compute overhead (Advanced NFV)

What is Titan IC RXP?

*RXP, **R**egular e**X**pression **P**rocessor:*
programmable custom-purpose content processor for high-speed pattern matching, supporting PCRE/POSIX regular expressions

- Large number of regex rules in parallel
- Scalable - 100Gb/s +
- Rich set of software support: compiler, API, etc.
- Customizable for target applications, Memory, Performance, Footprint, Power(ASIC)
- Complex RegEx-based pattern matching for:
  - Traditional (ACL) and NextGen Firewall (DPI), Intrusion Detection/Prevention (IDS/IPS), e.g. Snort
  - Application & Protocol Recognition, Application Firewall, detection of SQL injection, Application DoS
  - Database Acceleration (Spark, Elastic Search…), Computational Storage, AI/ML/NLP Preprocessing
  - SDN rule lookup/matching (Multi-Table), …….

Titan IC - 100Gb/s RXP Processor

Technology: GlobalFoundries, 28nm HPP

Parameter	Value
Data width	128-bit
Clock frequency	800 MHz
Prefix capacity	16K
Number of clusters	8
TCM:CACHE	2K:2K
Total memory	27,132,864 bits
Memory macro area	14.628 mm^2
Standard cell area	0.935 mm^2
Total post P&R area	19.665 mm^2
Power	4.55 W

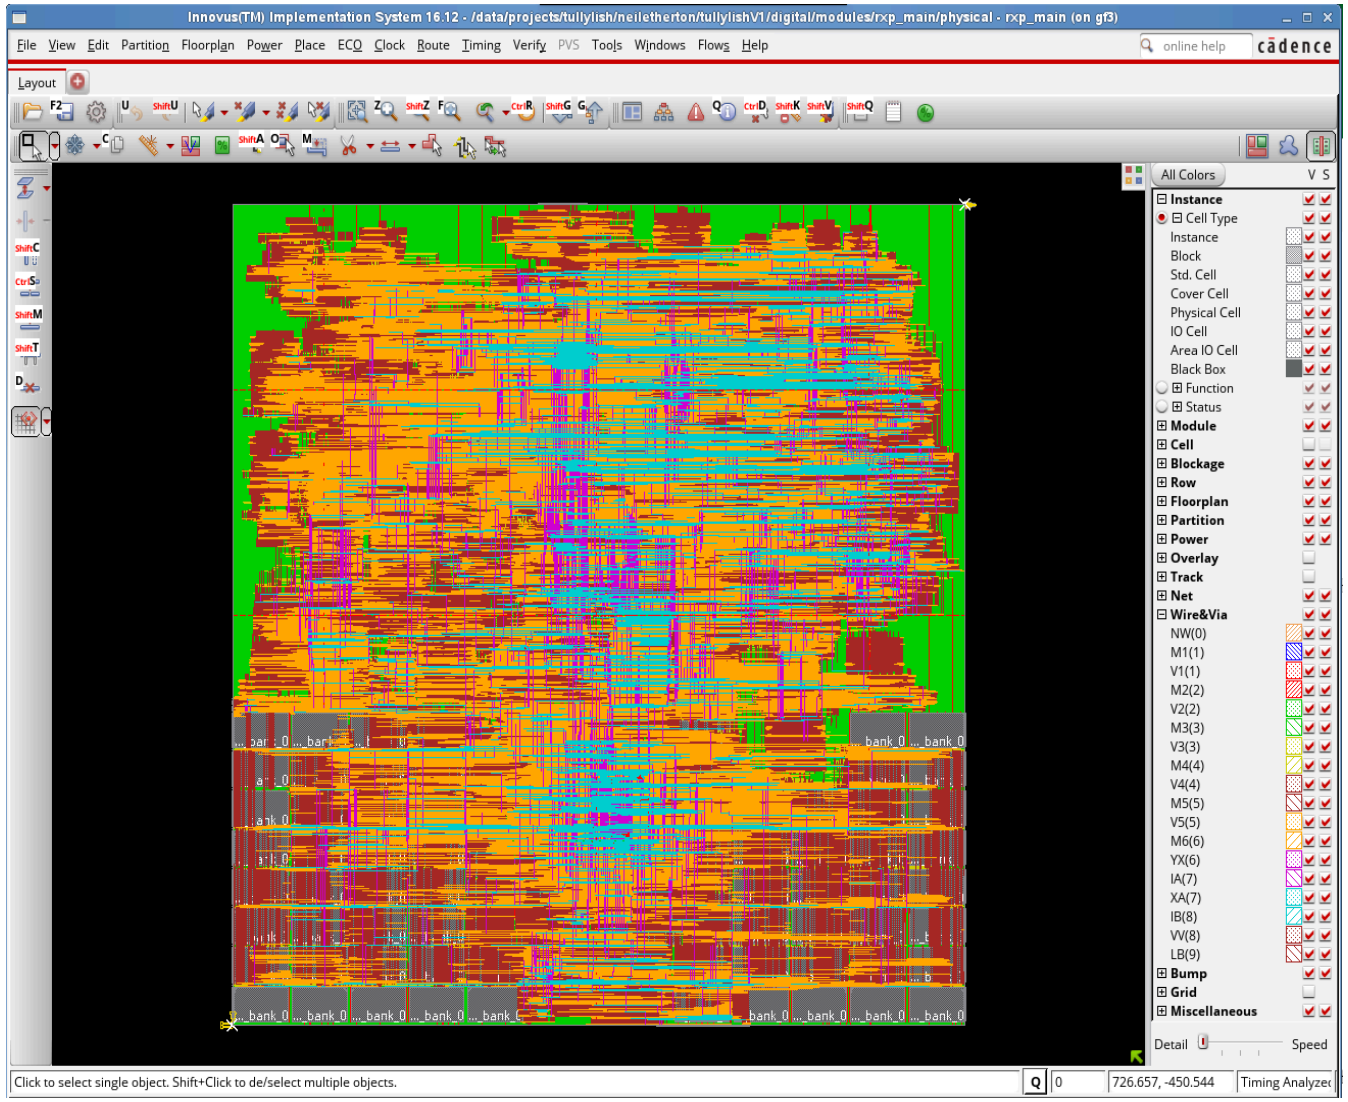

Use-Case: Mellanox BlueField-2

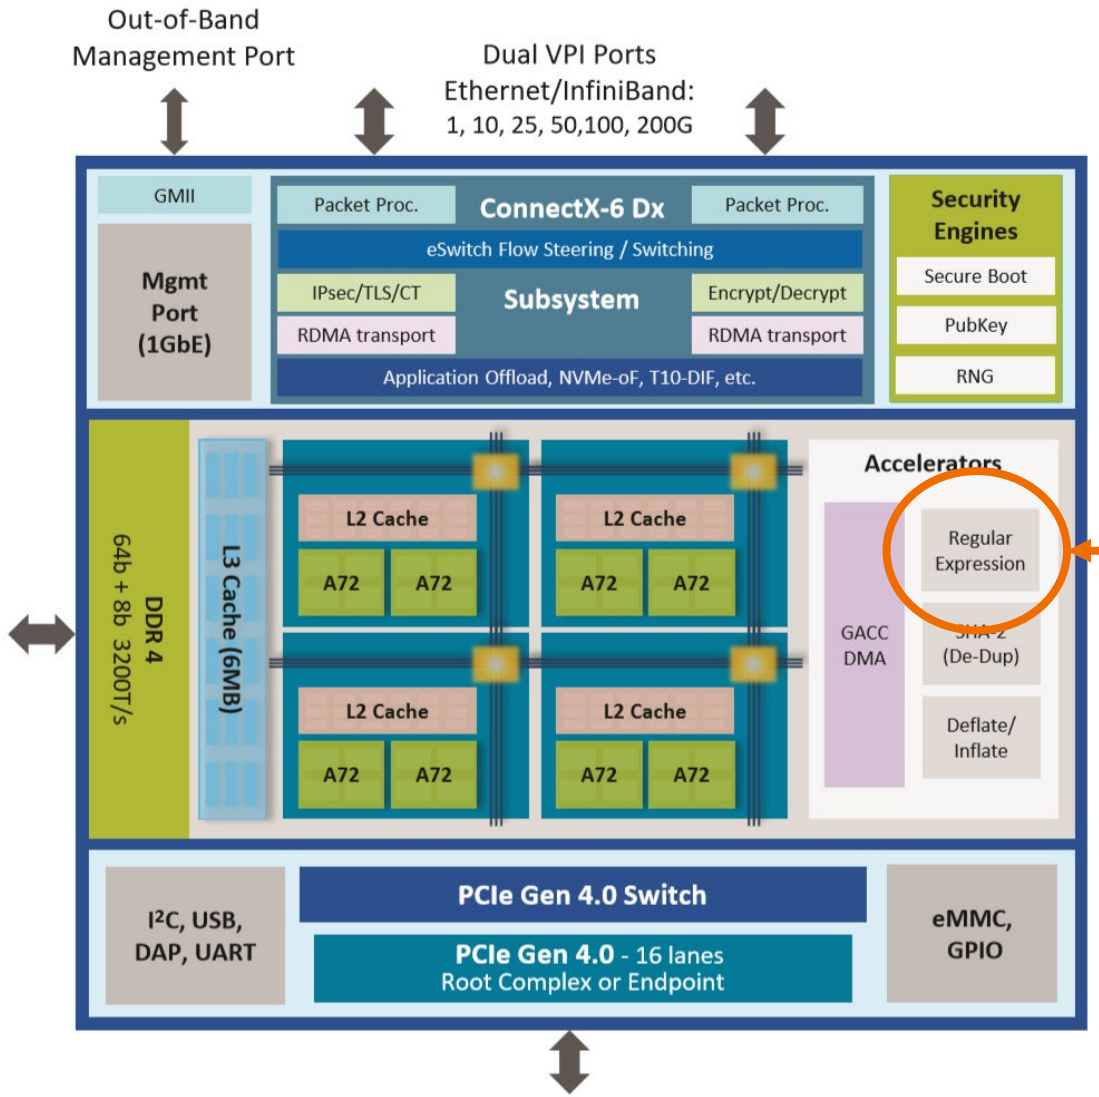

rxp inside

Titan IC RXP

- 50Gb/s RegEx offload
- >1,000,000 rules (External DDR)
- PCRE/POSIX Regular Expression
- Run-time rule update
- Incremental (partial) rule update
- Optimized for Network IPS (Snort)
- NGFW, WAF, SLA policing, etc.

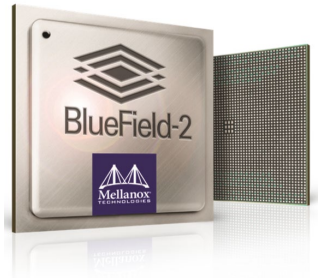

RXP - Soft IP for FPGA Implementation

	RXP Resource Requirements Xilinx KU115, 156 Mhz clock		RXP Resource Requirements Xilinx Vu9P, 200 Mhz clock	
Bandwidth	20 Gb/s	40Gb/s	50 Gb/s	100Gb/s
Rules Capacity (up to)	1 million	1 million	1 million	1 million
# BRAMs	904	1655	586	1172
#URAM	N/A	N/A	297	594
# LUTs	113K	216K	216K	432K
# FFs	130K	241K	255K	510K

- **Key Features**
  - 50G,100G bandwidths
  - Parallel processing of Regex
  - POSIX/PCRE compatible regular expressions
  - Interfaces: AXI, Native, PCIe
  - 100Gb/s uses 2 instances of the 50Gbps 256bit data path IP

Use-case Example 200G Network Traffic Monitoring and Content-Based Selective Traffic Intercepting

Titan IC

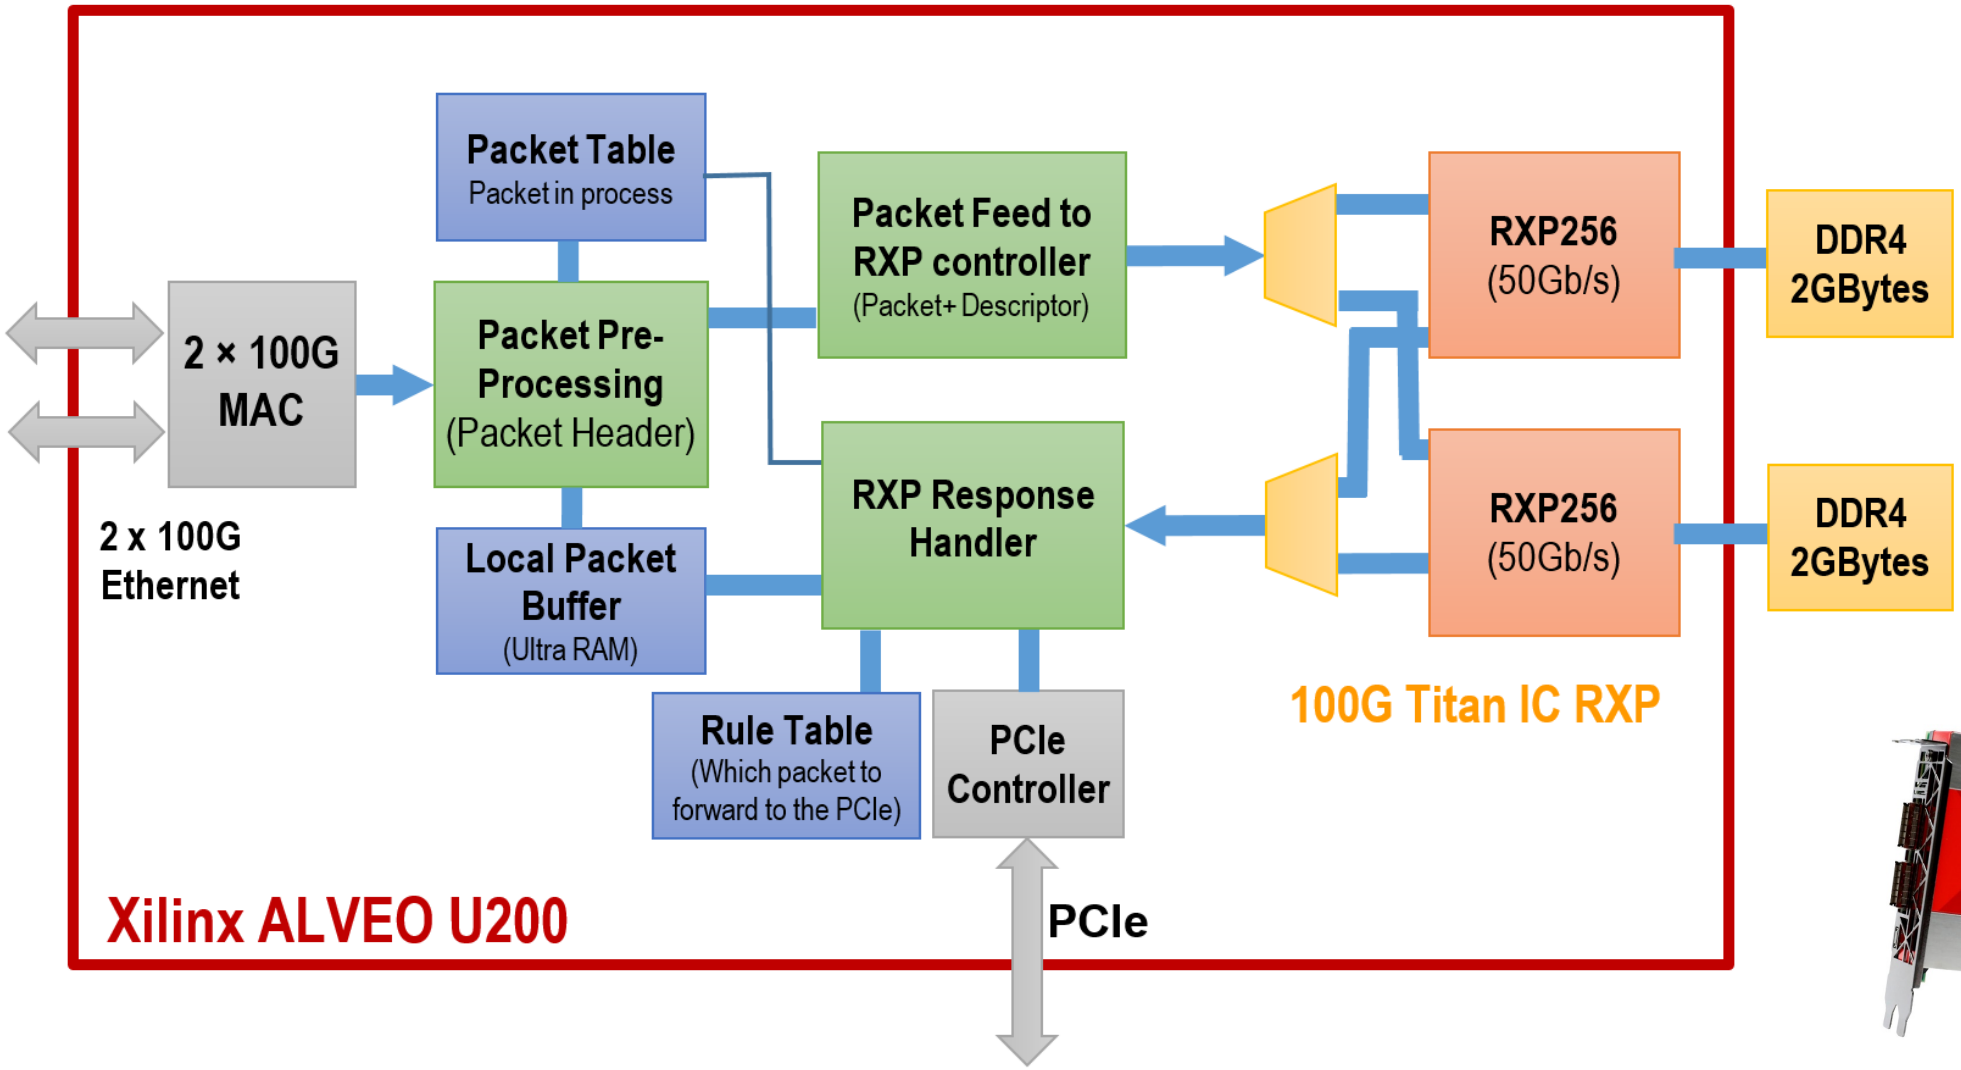

Use Case: SmartNIC Snort 3.0 Network IDS/IPS

- Snort 3.0: - Open source Network Intrusion Detection/Prevention System
- Optimized for real-time detection and prevention of network centric attacks and issues: - buffer overflows, stealth port scans, semantic URL attacks, CGI attacks, etc.
- Snort 3.0 operation can be subdivided into 7 phases
- http processing is stateful and inspection targets reassembled PDUs
- Multiple Snort instances (on multiple cores) can offload many PDUs in-flight

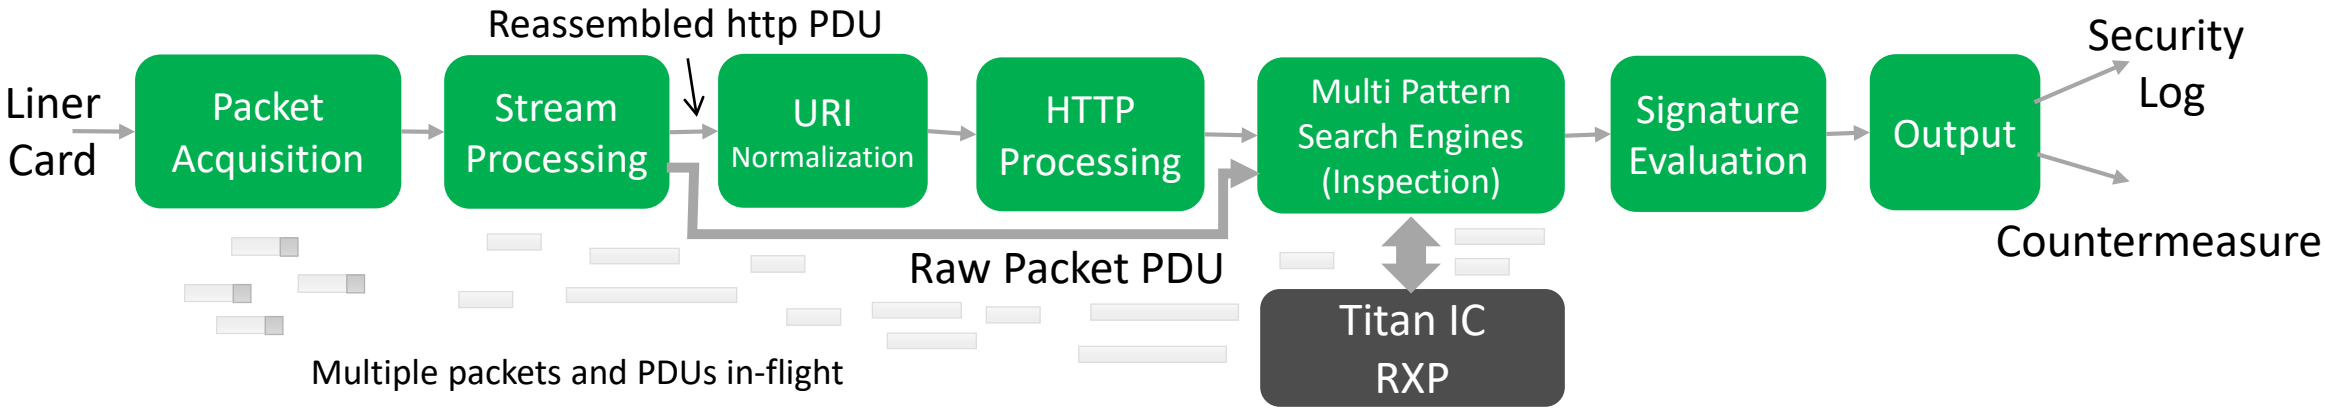

Use Case: SmartNIC Snort Network IDS/IPS

Titan IC

```
alert tcp $EXTERNAL_NET any -> $HOME_NET 1978
(msg:"APP-DETECT Apple OSX Remote Mouse usage";
flow:to_server,established;

content:"mos ";  fast_pattern:only;
pcre:"/mos\s{2}\dm\s\d/";
reference:url,pastebin.com/F81NCiYE;
classtype:policy-violation; sid:20443; rev:2;)
```

Packet Header Rule converted into RegEx

```
/^\xB7\x11\x3A\x5
1\x8F\x75\x45\x53
  .{2}\x07\xBA/
```

Fast Pattern Rules

PCRE Rule

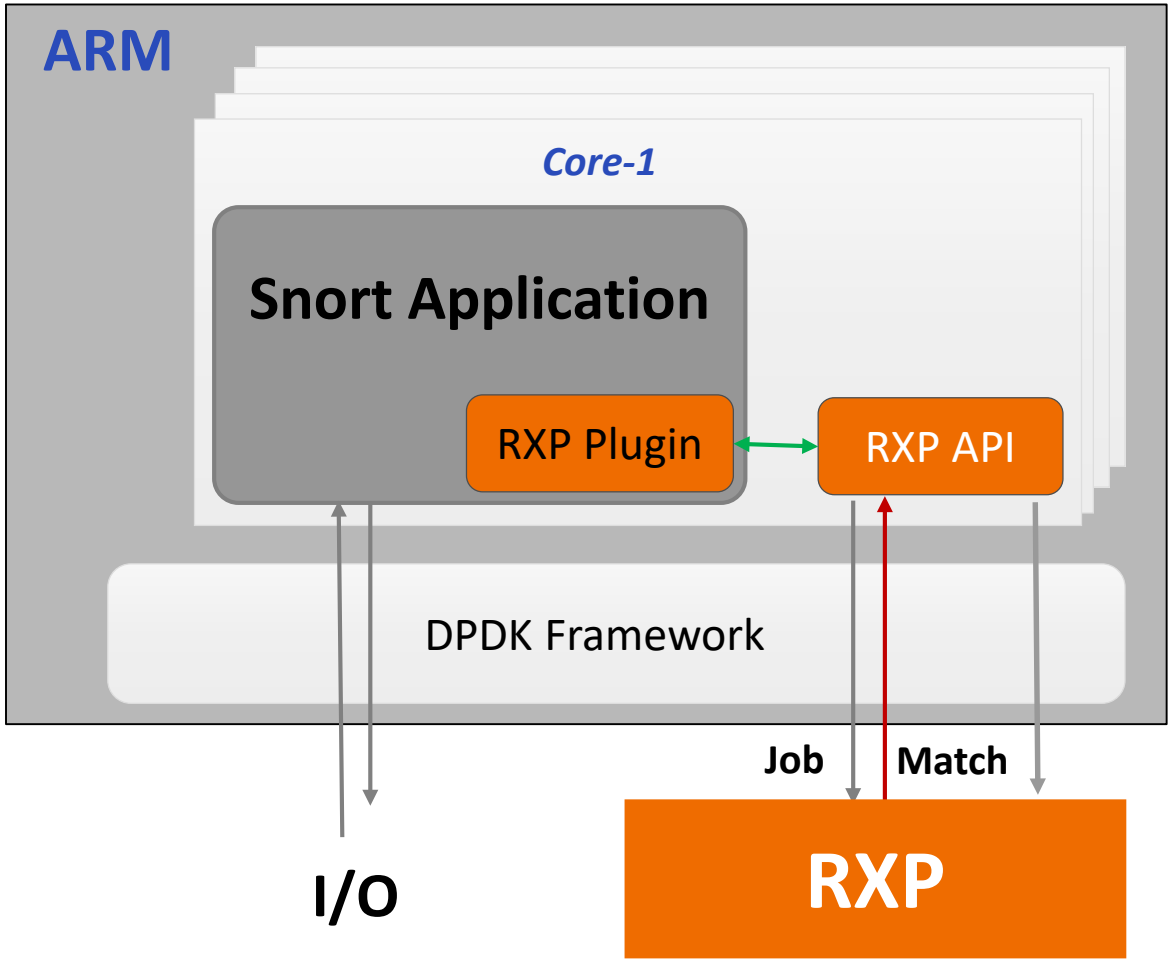

Challenges and Opportunities

Titan IC

Challenges:

- Lack of common approach for User Application Integration
- Lack of SmartNIC User Application Orchestration
- Lack of SmartNIC Native Applications
- Lack of standard interface support for efficient offload acceleration

Opportunities:

- SmartNIC “Open Data-Plane” Framework with offload acceleration
- VNF based SmartNIC open framework for third-party application
  - OVS, NGFW, IPS, WAF, VPN, vRouter, etc.
- Adaptation of established open-source applications and frameworks
  - DPDK Framework
  - OpenStack for NIC orchestration
  - Snort / Suricata (IDS/IPS)
  - ModSecurity (WAF)
  - DB / Spark / ELK Offload (Computational Storage)

In Summary

- Security is an indispensable service underpinning the fabric of Hyperscale Datacenter
- Evolution of Datacenter Server architectures postulates need for scalable security solutions within the server infrastructure
- Exciting new opportunities for enabling critical security functions and new type of unforeseen services on SmartNICs
- Titan IC is providing **underpinning technologies** enabling **critical security solutions** on next generation SmartNICs