Solution Brief

Distributed Services for Cloud Providers

PENSANDO

SOFTWARE-DEFINED | EDGE-ACCELERATED | ALWAYS-SECURE & VISIBLE ANY-ENVIRONMENT

CHALLENGES FOR CLOUD PROVIDERS

Cloud providers have entered a new realm when it comes to customer requirements, accountability and availability. The cloud provider's infrastructure now serves as the datacenter backbone for countless businesses, government agencies, and global enterprises. The challenges of building a cloud infrastructure go beyond providing IaaS with compute, networking and storage:

- **Scale:** Support millions of routes, security policies, network access lists, tunnel endpoints, ECMP paths, virtual block storage, etc. Often these requirements are not mutually exclusive making scale to be the top problem in a hyper-scale datacenter.
- **Performance:** Support for higher speeds of Network I/O up to 100G,, with the lowest possible latency and jitter.
- **Mixed Workloads:** Providing a uniform layer of policy enablement for a variety of workloads including bare-metal, virtualized and containerized.
- **Efficiency:** Optimal resource utilization (cpu cores, memory, storage, and network) is a fundamental competitive requirement. Power envelopes are tight, resource carving must be dynamic, all features should be available simultaneously, and deliver new features optimally.
- **Flexibility:** Business functions in a cloud infrastructure change at a very rapid pace. The usability should be simple and the infrastructure must be flexible to program.

While these challenges are applicable for datapath processing, it is equally important to have software that serves the configuration at a very high rate and supports a nondisruptive upgrade for high availability. The scale also pose challenges to collect voluminous data and export it accurately to provide insights to users.

Another challenge is to secure the infrastructure. While at times, encrypting data in transit within the cloud was considered exotic, some cloud providers started to offer this as a service. This requires providing a private key storage, data encryption, while providing tenant level isolation guarantees.

PENSANDO NAPLES™ DISTRIBUTED SERVICES CARD

The Pensando Distributed Services Card is architected to address all the above challenges for all types of workloads (containerized, virtualized, and bare-metal).

The Pensando Distributed Services Card is based on a domain-specific ASIC, providing highly optimized programmable hardware for packet processing and offering a broad suite of software-defined network, security, telemetry and storage services. The Naples™ Distributed Services Card (DSC) operates at 100G wire-speed with high-performance, low-latency, low-jitter, and the highest scalability targeted for the largest cloud providers.

A key value of the Pensando DSC is not only the comprehensive number of services offered, but also in the ability to chain the services together

in a programmable sequence, without loss of performance at 100Gbps, with a few microseconds of latency.

HIGHLIGHTS

100Gbps line-rate, Cloud Scale networking with:
Up to 1 Million Connections/sec
35 Million Packets/sec
Up to 4 Million IPv4/IPv6 LPM Routes
Less than 5 μsec Latency
Less than 40 nanosecs Jitter
P4 Programmable Data Pipeline
1K SR-IOV VFs
Multi-Tenancy (datapath isolation)
Granular Telemetry on all services/features
Maximum up to 35W power consumption

FEATURES

Networking
- Switching/Routing, SR
- L3 ECMP, L4 Load Balancing
- Overlay Networking (VXLAN, MPLS/UDP, Custom)
- IP-NAT and Port-NAT
- SPAN and ERSPAN (Bidirectional)
- Metrics collection and export in datapath

Security
- Micro Segmentation, NACLs
- DoS protection
- IPsec termination
- TLS/DTLS termination w/ TCP Proxy

Storage
- NVMe/TCP
- NVMe/RoCEv2
- Compression/decompression
- XTS encryption/decryption
- SHA3 Deduplication
- CRC64/32 checksums

DOMAIN SPECIFIC ARCHITECTURE

As Moore's law reaches its limit, use of general purpose computing will not be sufficient to meet the demands of cloud scale deployments. Even when it is feasible, CPU may be inefficient (power/cost vs processing power) for doing IO processing to simultaneous functions at scale with superior performance (100G).

To deliver this level of value for cloud providers, Pensando's Platform includes an innovative Domain Specific Architecture as an intelligent Edge on the server's PCIe bus, and allowing full control of network forwarding and data path/pipeline programmability.

Instantiating an intelligent network Edge on each server allows for inherent scalability.

With datacenters evolving so rapidly, solving any one problem aspect would only create more problems. Pensando has instead taken a scalable and holistic approach, providing a distributed software-defined platform . Not only are standard network services provided implicitly, but the platform agility allows cloud providers to customize, program and control all aspects of network traffic at the Edge.

PROGRAMMABLE DATAPATH

In addition to advanced software-defined network and security services, Pensando provides dataplane programmability through Naples, via a P4-programmable data pipeline that allows for customization of each layer of the cloud provider's infrastructure stack. Cloud providers can now take full control and assume full

ownership of the network and storage stack and all its elements.

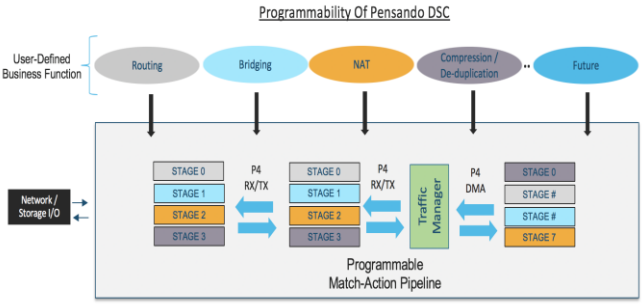

To enable cloud providers to integrate fully with their own cloud control plane, Naples allows cloud providers to own the entire software stack, or use Pensando feature bundles through REST/gRPC APIs.

PERFORMANCE AT SCALE

Cloud deployments require very large scale, and performance while all features enabled simultaneously. Naples DSC has been built with careful consideration to allow for multiple parallel match-action processing engines delivering the cache proximity or locality for IO processing, to achieve the unique combination of scale, performance for all business functions regardless of order.

DISTRIBUTED NETWORK SECURITY

Pensando delivers stateful security groups and firewall with connection tracking for tenant workloads, with the enforcement point completely isolated from those workloads. Naples DSC can be used by cloud providers to offer Security Groups or Network Access Control Lists (NACLs) for bare-metal, virtualized and containerized workloads uniformly.

Pensando's programmable datapath, allows cloud vendors to use a variety of attributes such as VLAN, MAC address, or VF to identify tenant workload for network isolation. Enabling different tenants to securely share a given physical server improving VM density with greater flexibility.

SEAMLESS OPERATIONAL INTEGRATION

Pensando Naples DSC can also be used in bump-in-the wire mode, as shown in the diagram below, to offer various services without requiring installation of any Pensando software on the server/hypervisor.

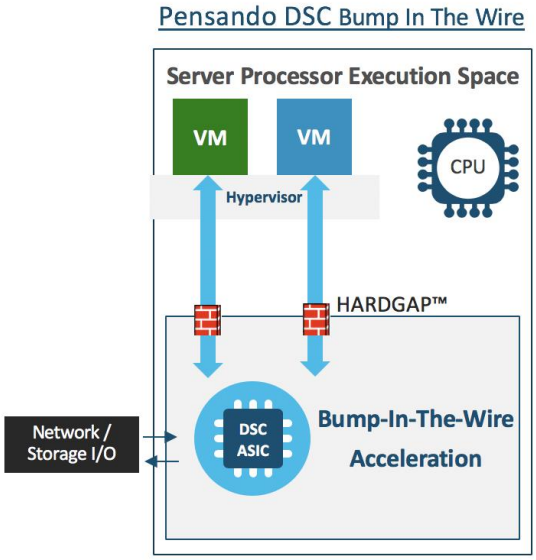

Network policies and firewall rules are deployed to a server with Naples as a remote network line card, rather than a host-based NIC. Policy management and configuration can be integrated with cloud provider's own management/control plane. All servers and associated applications gain all the security and performance benefits without any CPU overhead nor any host OS impact whatsoever. And cloud providers can reap huge benefits through uniform,

OS-agnostic management methods of bare-metal, virtualized servers, and containers.

PENSANDO HARDGAP™ TECHNOLOGY for INFRASTRUCTURE PROTECTION

HardGap™ technology provides PCIe layer secure hardware isolation between any software running on a server from the Naples DSC enforcement engine. This is essential for a cloud provider in order to ensure that tenants can never gain access to the cloud infrastructure, under any circumstances.

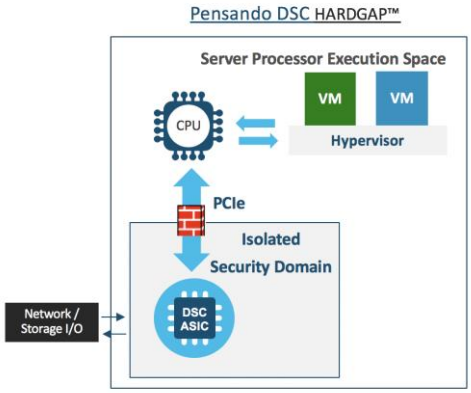

PENSANDO DSC SOFTWARE BUNDLE

Although a cloud provider can develop custom business functions in datapath and software stack on DSC, Pensando offers a fully functional, highly optimized set of various business functions, aimed for the cloud environment. The software stack is accessible using industry standard REST/gRPC interfaces, implementing cloud object model.

SUMMARY

As cloud infrastructure growth continues to accelerate, driven by customers increasing adoption of hybrid cloud architectures, the performance, scalability and functionality demands on cloud providers will not cease.

A next-generation public cloud architecture with Pensando's Distributed Services Card as the foundation, will enable the reliability, flexibility and visibility needed to deliver cloud services. This coupled with orders of magnitude improvement in scalability and performance, will future proof their infrastructure for growth, and position the leading public cloud providers for success.

ABOUT PENSANDO SYSTEMS

Founded in 2017, Pensando Systems is the company pioneering distributed computing designed for the New Edge, powering software-defined cloud, compute, networking, storage and security services to transforms existing architectures into the secure, ultra-fast environments demanded by next generation applications. The Pensando platform, a first of its kind, was developed in collaboration with the world's largest cloud, enterprise, storage, and telecommunications leaders and is supported by partnerships with HPE, NetApp, Oracle, IBM, Equinix, and multiple Fortune 500 customers.

For more information, please visit www.pensando.io