Solution Brief

Distributed Services for Cloud Providers

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 a major problem in hyper-scale data centers
- **Performance:** Support for 100G infrastructure 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
- **Security:** Tenant-level isolation guarantees. Support for encryption of data-in-flight as well as data-at-rest, along with secure key management
- **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 optimally deliver new features
- **Flexibility:** Business functions in a cloud infrastructure change at a very rapid pace. Delivering an agile infrastructure should be simple and flexible with open APIs

Traditional "scale-up" approaches—where networking services are deployed as networking or security appliances, or rely on limited, static functionality in traditional top-of-rack switches—are no longer able to keep up, suffering from either performance or scale limitations as policy tables bloat and the number of active flows reaches into the millions. The limitations and expense of this centralized resource model have led data center architects to limit core network infrastructure functions to simply transporting IP traffic with as little latency and jitter as possible.

Just as compute and storage systems are adopting a "scale out" approach, so too the networking and security elements of the cloud data center must adopt a Scale-out Services Architecture and these functions need to find a new home in this model

The ideal place to deploy these services is the server edge (the border between the server and the network) where services such as overlay/underlay tunneling, security group enforcement and encryption termination can be delivered in a scalable manner. With each server edge node tightly coupled to a single server, it needs to be aware only of the policies related to that server and its users and applications. This approach naturally scales, as more services capabilities come along when new servers are added.

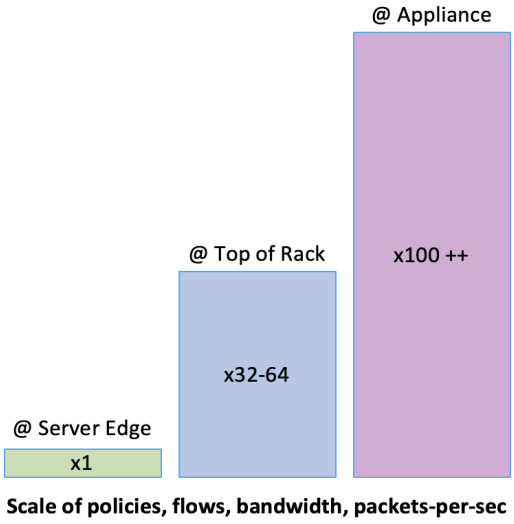

Pensando Distributed Services Platform

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

The Pensando Distributed Services Card (**DSC**) is based on a custom designed domain-specific programmable processor, providing highly optimized hardware for packet processing and offering a broad suite of software-defined networking, security, telemetry and storage services. The Pensando 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.

Performance at Scale

Cloud deployments require very large scale, and performance while all features enabled simultaneously. The Pensando 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 and performance for all business functions regardless of order.

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 limits, the use of general-purpose CPUs to manage infrastructure operations (e.g., networking and virtual firewalls) is not optimal for cloud scale deployments. Even when it is feasible, the CPU-centric approach is generally inefficient (power usage/cost vs processing power) especially given the requirements of 100G infrastructure.

The Pensando Distributed Services Platform includes an optimized domain-specific distributed processing unit (DPU) based on a P4 programmable processor and positioned on the server's PCIe bus, allowing full control of network forwarding and data path/pipeline programmability.

With data centers evolving rapidly, solving any one problem can often create multiple, unrelated issues. Pensando has taken a scalable and holistic approach, providing a distributed software-defined platform. Not only are standard network services provided implicitly, but the platform delivers agility by allowing cloud providers to customize, program and control all aspects of network traffic at the server edge.

Programmable Data Path

In addition to advanced software-defined network and security services, Pensando provides data plane programmability through the DSC, via an open 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.

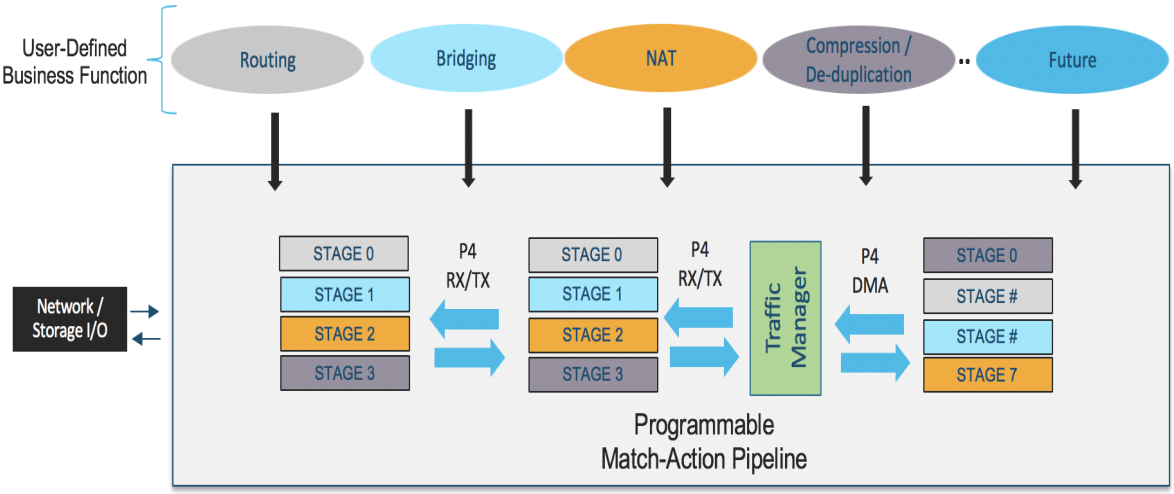

Figure 1: Programmability of the Pensando Distributed Services Card

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

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. The 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 data path 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

The Pensando DSC can also be used in bump-in-the wire mode, as shown in Figure 2, to offer various in-line services without requiring installation of any Pensando software on the server.

Network policies and firewall rules are deployed to a server as a remote network line card, rather than a host-based interface. 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. In addition, cloud providers can reap huge benefits through uniform, OS-agnostic management methods of bare-metal, virtualized servers, and containers.

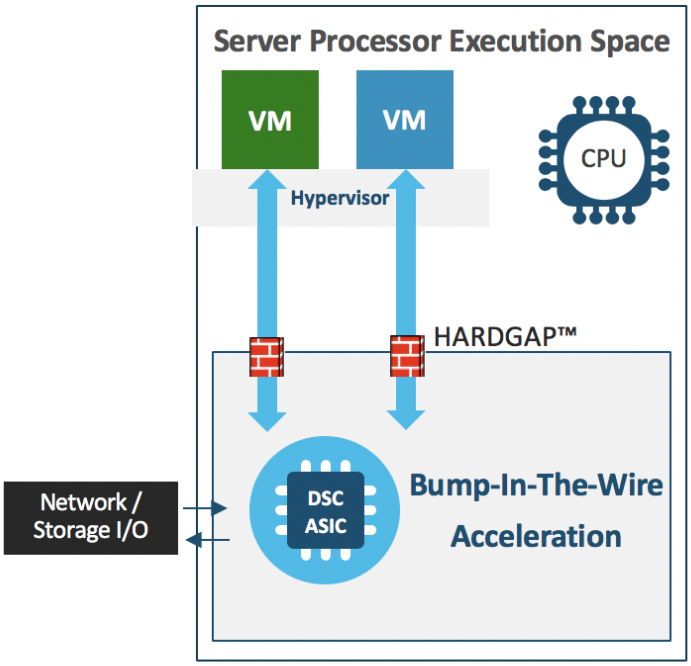

Figure 2: Pensando DSC deployed as a "bump in the wire"

Pensando HardGap™ Technology for Infrastructure Protection

HardGap™ technology provides PCIe layer secure hardware isolation between any software running on a server from the 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.

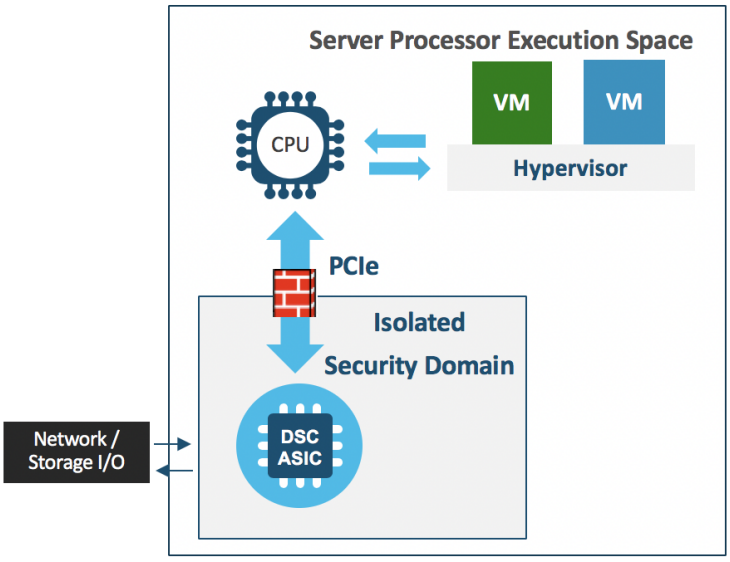

Figure 3: Pensando HardGap technology

Pensando DSC Software Bundle

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

Founded in 2017, Pensando Systems is pioneering distributed computing designed for the New Edge, powering software-defined cloud, compute, networking, storage and security services to transform 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 Hewlett Packard Enterprise, NetApp, Oracle, IBM, Equinix, and multiple Fortune 500 customers. Pensando is led by Silicon Valley's legendary "MPLS" team—Mario Mazzola, Prem Jain, Luca Cafiero, Soni Jiandani and Randy Pond — who have an unmatched track record of disruptive innovation having already built 8 $Bn/Year businesses across storage, switching, routing, wireless, voice/video/data, and software-defined networking. The company is backed by investors that include Lightspeed Venture Partners, Goldman Sachs and JC2 Ventures.
For more information, please visit pensando.io