Solution Brief

Telecom Providers & Medium to Large Vertical Enterprises

Tests Show Scalability of Virtualized StoneWork Multi-Service Router

Tests conducted on the PANTHEON.tech StoneWork Multi-Service Router on Intel® architecture-based servers show scalability across multi-tenant communications service provider use cases

In the rapidly evolving digital landscape, communications service providers (CoSPs) face increasing challenges in meeting customer demands. Traditional dedicated hardware appliance-based customer premises equipment (CPE) does not have the flexibility or scalability to deliver the service agility, scalability, and cost efficiency required to keep up with today's dynamic network requirements.

As a result, virtualized provider edge routers – which are routers that interconnect the networks of two or more CoSPs - have emerged as a transformative approach to address these pressing challenges. By embracing virtualization and adopting software-defined networking (SDN) principles, provider edge solutions offer a more flexible, scalable, and cost-effective alternative to traditional hardware-based customer premises equipment.

This solution brief featuring PANTHEON.tech, an Intel® Network Builders Gold Tier ecosystem member, presents the benefits of utilizing the virtualized StoneWork MSR and validates the performance and scalability of StoneWork MSR control and data plane on the latest generation Intel® architecture processor-based servers, enabling provider edge network functions. The aim of the document is to inspire CoSPs to use this solution to drive innovation in their network deployments.

By transitioning from hardware-centric approaches to software-defined solutions, CoSPs can harness the power of virtualization and cloud technologies to offer advanced services, streamline operations, and significantly enhance customer experiences.

Provider Edge Network

Provider edge (PE) networks are defined as the boundary separating one CoSP's network from that of another CoSP. PE networks play a vital role in bridging the gap between CoSPs and their customers, facilitating a wide array of services, such as internet access, virtual private networks (VPNs), and voice services.

These networks ensure the smooth and secure transfer of data packets across diverse network environments, delivering consistent, reliable, and high-quality services to end-users. In modern networking architectures, PE networks effectively address a range of critical challenges faced by both CoSPs and their customers, including:

1. **Connectivity and Routing**: As the amount of data traffic and other network requirements grow, the task of effectively handling and optimizing how devices and networks are connected, and how data is directed from one point to another, becomes more challenging. PE networks leverage advanced routing protocols and techniques to optimize data flow, reduce latency, and prevent network

congestion, leading to improved overall network performance and seamless data transfer.

2. **Security and Data Protection**: With the ever-increasing frequency of cyber threats, safeguarding data transmissions and protecting sensitive customer information is utmost importance. PE networks implement robust security measures, such as Virtual Private LAN Service (VPLS) and Virtual Routing and Forwarding (VRF) mechanisms, to isolate and safeguard customer data from potential breaches.
3. **Scalability and Flexibility**: As businesses expand and customer demands increase, traditional hardware-based solutions struggle to keep pace with the required scalability.

PE networks offer enhanced flexibility and scalability by embracing virtualization and software-defined networking (SDN) principles. This allows CoSPs to quickly adapt to changing network demands, deploy new services efficiently, and expand their network infrastructure without significant hardware upgrades.

StoneWork MSR

StoneWork MSR is a cutting-edge networking software solution developed by PANTHEON.tech that caters to modern network paradigms while avoiding the pitfalls of monolithic architectures.

Container Orchestration

StoneWork optimizes resource utilization and delivers exceptional network performance by incorporating container orchestration that simplifies the management of containerized applications and automates essential tasks. These tasks include automatic scaling to handle varying workloads, load balancing to optimize resource usage, service discovery for easy communication between microservices, and health monitoring with self-healing capabilities to maintain high availability.

Additionally, container orchestration enables seamless updates and rollbacks through its advanced management capabilities. When a new version of an application needs to be deployed, the orchestration platform follows a controlled and automated process to ensure a smooth transition, minimizing downtime and efficiently managing large-scale microservices-based applications.

Vector Packet Processing (VPP)

Integrating Vector Packet Processing (VPP) into StoneWork MSR further enhances network performance. By utilizing Linux Foundation's widely adopted FD.io VPP capabilities, CoSPs can achieve high-speed packet forwarding and routing, reducing latency and improving the responsiveness of network applications.

VPP can be deployed as a container alongside other application containers, providing fast and efficient packet processing while benefiting from the orchestration platform's dynamic scaling and service discovery features. By offloading packet processing tasks and reducing the burden on the CPU, the VPP technology allows the CPU to focus on more complex tasks, resulting in a significant boost to its overall performance.

The control plane management agent, powered by Ligato VPP-Agent, simplifies network function configuration and orchestration, ensuring smooth operation even in complex network environments. VPP is also natively feature rich, with out-of-the box support for L2/L3 forwarding, ACL-based forwarding, NAT, VRFs and bridge domains.

Data plane virtualization offers a significant advantage by allowing CoSPs to boost performance using multi virtual routing and forwarding on a single VPP data path that handles packets efficiently along a processing graph in VPP. Compare this approach to a data path traversing multiple CNFs, which has inherent delays in forwarding packets between individual containers. The data plane VPP functionality is optimized to a single container, providing a single common containerized data plane for all individual and collective control plane functions.

VPP is designed to provide a high-performance, flexible, and programmable networking framework. It has its own data plane implementation, which processes and forwards network packets. However, VPP also supports using Data Plane Development Kit (DPDK). DPDK was created by Intel and is now open source library of software that can accelerate packet processing workloads running on a wide variety of CPU architectures.

This means that instead of relying solely on VPP's native data plane implementation, VPP can utilize DPDK for packet processing.

Integrating DPDK into VPP can provide several benefits, including:

- Improved packet processing performance and throughput.
- Access to DPDK's memory management optimizations.
- Better utilization of hardware acceleration features provided by supported network interface cards (NICs).

Cloud-Native Network Functions (CNFs)

For CoSPs seeking prompt adaptation to customer demands and market changes, StoneWork MSR shows its power through the dynamic integration of cloud-native network functions (CNFs). This adaptability ensures consistent service delivery, even under fluctuating demands, while container orchestration provides a scalable environment for CNFs to operate efficiently. StoneWork's CNF discovery capabilities further streamline integration, reducing manual configurations, and boosting overall network efficiency. By allowing CNFs to share network namespaces (using the same network environment for communication), resource utilization is optimized, and unnecessary processing is reduced, leading to improved packet handling, lower latency, and overall network efficiency.

4th Gen Intel® Xeon® Scalable processor

4th Gen Intel® Xeon® Scalable processor family offer high-throughput and low-latency and are engineered for on-prem deployments. The processor family's architecture combines high-performance processor cores with up to eight built-in accelerators[1] for maximum performance efficiency. Integration of accelerators into the processor redefines CPU architecture and provides a more efficient way to achieve higher performance than relying solely on increasing the CPU core count for workload processing.

The integrated accelerators help reduce a server's power consumption and complexity when compared to a server that uses discrete accelerator cards for additional performance. This power advantage is important to edge servers.

These processors offer up to 80 lanes of PCIe 5.0 connectivity and support Compute Express Link (CXL), a cache-coherent interconnect for processors, memory expansion, and accelerators.

Other seamlessly integrated accelerators include Intel® QuickAssist Technology (Intel® QAT) that speeds up data movement and compression for faster networking, boosts query throughput for more responsive analytics, and offloads scheduling and queue management.

Intel - StoneWork MSR Performance / Scale Test

Virtual routing and forwarding (VRF) is a fundamental technology in modern networking that allows for the segmentation and isolation of traffic within a single physical router or switch. By enabling the creation of multiple customer VRFs (cVRFs), each with its own routing table and forwarding decisions, VRFs offer a powerful solution for multi-tenancy and secure traffic separation in a network environment.

This test explores the capacity and scalability of VRF implementation on a single multi-service router (MSR) instance. Specifically, the tests included how many cVRFs can be configured and effectively carry traffic on the StoneWork MSR, examining the level of workload the MSR can handle and its potential for scaling. The results obtained provide valuable insights for CoSPs seeking to optimize their infrastructure for efficient multi-tenant deployments and accommodate increasing traffic demands.

Edge Multi Service Router Requirements

To ensure secure and efficient communication between customers in a network, certain essential features are needed in the edge MSR. Each customer's virtual local area network (VLAN) requires its own dedicated DNS and virtual DHCP server to allocate private IP addresses within their respective virtual routing and forwarding (VRF) instances.

This segregation of VRFs ensures that customers remain isolated from each other, and each VRF operates within its own separate broadcast domain, enabling the use of individual virtual DHCP servers.

To connect customer VRFs with the public internet and other external networks, network address translation 44 (NAT44) functionality is employed. This mapping mechanism translates private IP addresses from customer VRFs to public IP addresses, facilitating smooth communication between customer networks and the outside world. The translated traffic is then forwarded through the backbone VRF of the underlying network to reach the transit network and the internet.

In addition to these features, access control lists (ACLs) are necessary on the output interfaces of the common backbone VRF. These ACLs serve to restrict access from customer VRFs to destinations within the backbone network, ensuring that customer traffic remains confined to its designated areas and does not interfere with or access resources in the backbone network.

Test Bed Description

The purpose of the test bed is to determine the capacity and scale of the StoneWork MSR's control and data planes. More specifically, how many customer VRFs (cVRFs) can be configured and carry traffic in a single MSR instance. Table 1 shows the configuration of the 4th Gen Intel Xeon Scalable processor-based device under test (DUT) the traffic generator and the configuration of the test bed.

Device Under Test (DUT)	T-Rex (Traffic Generator)	Test Bed Details
CPU: Intel® Xeon® Platinum 8470N	CPU: Intel® Xeon® Gold 6152 CPU @ 2.10GHz	Testing with 10, 100, 300, 1,000 tenants (VRFs)
Memory: 256G	Memory: 384G	Testing booth directions (upload / download)
2 x 100Gbps network controller 0: Ethernet Controller E810-C for QSFP	2 x 100Gbps network controller: Ethernet Controller E810-C for QSFP	Testing from 1 to 10 VPP workers (DUT)
Docker image: ghcr.io/pantheontech/stonework:22.10.1	T-rex version v3.02 in STL mode using 8 CPU cores	Running T-rex console with 8 cores Using namespaces to simulate customers

Table 1. StoneWork MSR test bed description.

		Gbps (L1 overall Throughput)			
Direction	CPU cores (workers)	10 tenants	100 tenants	300 tenants	1000 tenants
Unidirectional (0--->1)	1	27	26	22	17
	2	42	40	39	33
	4	66	66	70	67
	8	89	83	100	98
	10	100	100	100	98
Bidirectional (0--->1, 1--->0)	1	28	26	22	11
	2	38	40	37	28
	4	66	68	68	60
	8	98	95	92	87
	10	98	96	92	85

Disclaimer: The numbers apply to this specific environment and may differ if a different configuration is used.

Table 2. Throughput results using unidirectional and bidirectional data flows.

Test Results

Table 2 shows the results of the tests that were conducted with both unidirectional and bidirectional data flows. Throughput was measured in Gbps at layer 1 and the scalability is shown by the increase in throughput per CPU core in each of the tenant use cases.

Conclusion

The test results verified the performance and scalability of StoneWork MSR control and data plane on a cutting-edge Intel architecture-based server, providing a provider edge router function.

Thanks to the Intel Network Builders testbed environment, PANTHEON.tech was able to confirm the high compliance of its StoneWork MSR solution on Intel hardware, as well as the high scalability and throughput of this solution.

The tests delved into the potential of VRFs and the efficiency in handling traffic on the StoneWork MSR, aiming to understand its workload capacity and potential for scaling.

The test results revealed insights regarding the impact of the number of tenants on CPU performance, highlighting the linear scalability of the architecture. The analysis demonstrated that an increase in CPU cores directly correlates with enhanced throughput, emphasizing the significance of efficient multi-core utilization.

By utilizing a fraction of CPU cores (36%) for the data plane on a 22-core CPU, StoneWork MSR achieves impressive results, indicating the potential for further optimizing performance by distributing the workload between the data and control planes on the remaining cores. With a dual-socket server, there are a total of 44 cores that could be used, but only 8 cores were needed in the worst case to achieve 100Gbps throughput.

These findings provide valuable guidance for network administrators and architects seeking to optimize their infrastructure for multi-tenant deployments and effectively manage surging traffic demands while creating a resilient and high-performing network environment to meet the challenges of the modern digital landscape.

Learn More

PANTHEON.tech

4th Gen Intel Xeon Scalable processor

Data Plane Development Kit

Intel Network Builders Ecosystem

intel.

[1] https://www.intel.com/content/www/us/en/products/details/processors/xeon/scalable.html

0524/TM/DJA/PDF Please Recycle 360210-001US