packet NETRONCME

Using Network Acceleration for an Optimized Edge Cloud Server Architecture

My Truong Senior Director, Packet Labs

Ron Renwick *VP of Products, Netronome*

About Packet



packet.com / @packethost





Packet empowers developer-driven companies to deploy physical infrastructure at global scale.



A "Go Anywhere" Cloud Model

Packet's automated infrastructure is delivered in our public cloud, as well as in Enterprise deployment models.





20+ Public Cloud Locations

- Amsterdam
- Ashburn
- Atlanta
- Boston
- Chicago
- Dallas
- Dublin
- Frankfurt
- Hong Kong
- Los Angeles

- Marseille
- New York
- Phoenix
- Pittsburgh
- Seattle
- Singapore
- Sydney
- Sunnyvale
- Tokyo
- Toronto



and and



Why Packet?





packet.com / @packethost

The Experience of the Cloud on Bare Metal

Built by developers, for a developer world - the Packet experience is both intuitive and powerful.

- No Multi Tenancy
- Scalable, Hourly Pricing
- Powerful, Intuitive API
- CloudInit & Meta Data Support
- Native API Libraries
- Leading Integrations

CHOOSE PROJECT -	SERVERS IPS & NETWORKS	STORAGE PROJECT SETTINGS					NEED HELP?	6.
	Q Filter by location, OS, IP, hostname	or tags					+ Add New	
	HOSTNAME .	IPV4 ADDRESS	CONFIG	os	LOCATION	TAGS		
	Robert	147.23.123.67	Type 2A	0	AMS1	VIEW PROGRESS		
	Anouki	148.458.742.45	Type 2A	ø	NRT1	НТТР	1	
	Deimos2	147.23.123.66	Type 2A	ø	AMS1	SMTP	I	
	💈 🌑 My Radical Project Producti	on 147.23.123.66	Type 2A	0	AMS1	EMAIL BACKUP +3	Ξ	



Automating at the Hardware Layer

Lock-In Risk	High Proprietary services are the core business model.	None Workloads are portable to any datacenter.		
Hardware Access	No Access Hardware is not exposed. No chance for optimization.	100% Dedicated Users have full access to hardware.		
Automation Layer	Hypervisor Forced virtualization & multi-tenancy.	Hardware No virtualization tax or noisy neighbors.		
Bare Metal Servers	Google Cloud Platform	packet		



Hardware at Software Speed

By separating the hardware from the software, disaggregation is resulting in significant acceleration of innovation. DevOps users now control the pace of innovation.





What changes at the Edge?

Sub-Scalable Architecture





packet.com / @packethost

Solving For Latency And Edge Cloud Constraints



- Traditional cloud data centers are incredibly large and too far away
- Edge Cloud deployments have stringent cost, power and latency requirements
- Solution: A different architecture in a distributed location
 - Enable smaller and more cost-efficient servers: independently available to developers
 - Utilize network accelerators to enhance compute performance



Bare Metal IaaS For Edge Cloud Apps Using Open19 Design



Operator

- Manages bare metal microservers using SmartNIC
 - ACL rules, Firewall, Encryption
 - Control and data planes and Operator SDN controller
- Offer wide range of customer solutions
 - SmartNIC/non-SmartNIC enabled
 microservers
 - Cost and power-efficient servers

Edge Cloud Customer

- Owns one or more microserver (bare metal)
 - O/S or hypervisor of choice
 - Run low-latency 5G/IoT oriented applications at the Edge
 - Develop accelerated networking apps using eBPF programming
 - Option of SmartNIC or basic NIC
- With SmartNIC
 - Offloads eBPF data plane programming to SmartNIC
 - Programmable RSS
 - Security
 - Upstream OVS-TC and RHEL options.

Accelerating Compute With Disaggregated Architecture

Memory Storage CPU CPU Host Processors Networking Interface (NIC)

Homogeneous Architecture - Historical

- Performance is throttled when CPU manages all traffic
- Scaling performance is breaking the cost/power budget
- CPUs alone cannot keep up with demands of new workload

Disaggregated Architecture



- Disaggregation is key to unlocking performance
- Enable servers to use more cost/power effective CPUs
- Provide lower latency to application and Edge devices

NETRON

Innovative Edge Cloud Microserver Form Factors

Standard 1U Form Factor Solution with 4 Microservers and Integrated Netronome SmartNIC Open19-based Solution with 4 Microservers and Integrated Netronome SmartNIC





Offloaded TLS Enables Superior Application Performance

Queries per Second Latency SW TLS HW TLS SW TLS HW TLS 1,000,000 300 750,000 200 -atency us 500,000 100 250,000 261 204 971.164 812,002

Hardware offloaded TLS provide 20% query throughput improvement

Hardware offloaded TLS provides 22% latency reduction

Benchmarked against Mellanox ConnectX-5, 50GbE

20% improvement on 28-core server = 6 cores released back to application

20% Throughput Improvement

22% Latency Reduction

netronome.com / @netronome

ZPS

eBPF Offload Boosts Performance, Scalability And Efficiency

- eBPF provides transparent acceleration to a wide variety of data center and enterprise edge apps
 - Filtering, Load Balancing, DDoS, Monitoring
 - eBPF-based security for bare metal servers
 - Drivers upstreamed and included in the Linux kernel
- Netronome SmartNICs accelerate and offload eBPF enabling higher performance and efficiency
 - Up to 5X performance for iptables firewalls





>20X eBPF/XDP Load Balancer Offload Performance

^oerformance (Mpps)

- Enables on-the-fly changing of load balancing algorithms
 - direct traffic to available servers based on time of day, geographies served and other security-related criteria)
- Offload of open source eBPF/XDP Load Balancer code to the Agilio SmartNIC yields up to 42Mpps using a single x86 CPU core
- Without offload, 16 x86 CPU cores yield 20Mpps, about **1.5Mpps per CPU core**



eBPF: Agilio vs. x86

XDP: Best-in-Class Latency Reduction

3X Latency Reduction: P(99). For P(99.99) see the next slide



P(99) Latency

netronome.com / @netronome

Best-in-Class Latency Reduction



In an OCP data center allows 2 type 1 servers (Twin Lake) to replace a type 6 server (Leopard/Tioga Pass) within a disaggregated flash architecture.

Saves ~100W (33% of total power) per replacement.

Up to 70X Improvement for P(99.99) Latency

netronome.com / @netronome

Agilio SmartNIC in Web 2.0 App And Data Servers

CONGESTION AND TAIL LATENCY REDUCTION FOR DISAGGREGATED STORAGE





Netronome – The SmartNIC Platform Leader

SmartNIC leader delivering efficient, scalable, programmable and highly secure cloud and hyperscale solutions – from the edge to data center servers





INNOVATION: Netronome invented the NFP as a networking app-specific silicon: it has the highest number of cores and processing threads per \$ and Watt

OPEN SOURCE LEADERSHIP: Prolific open source contributor (10+ orgs), supports an open research community of 200+ users and 70+ projects, founders of ODSA & Open-NFP **PRODUCTS RANGE:** The Agilio SmartNIC platform supports the widest range of eBPF, P4 and C-programmable SDN, NFV, security and visibility apps

Putting Words Into Actions

Calling All eBPF Lovers!





packet.com / @packethost

Edge Architectures are Different

HYPERSCALE MODEL

High Availability for Redundancy of Generic Workload

Application deployed across 3 zones with 2N+ power redundancy. Cooling within a very limited range. Deployment size of 15k -> 100k servers. Built by the megawatt!

EDGE MODEL

Distributed for Optimization of High Value / Specific Workloads

Micro datacenters often utilizing grid power, with potential for high temperature variance. Single deployment could be just a few dozen servers. Complex networking around 5G & IoT.







Developers are Ready for eBPF!



Credit: excellent article: "Learn eBPF Tracing: Tutorial and Examples" by Brendan Gregg



An Architecture for Modern Workloads

Containers, containers, containers. Microservers with SmartNICs are the perfect fit for containerized microservices!





Edge Access Program

Our Edge Access Program accelerates open source and commercial use cases by providing access to edge infrastructure, technology partnerships and expertise.

- Get access to unique, use-case driven edge sites in the United States.
- Bare metal with full automation (early SmartNIC options available).
- Wireless connectivity (CBRS) and Packet Connect (to Azure, GCP, etc.)

Contact sales@packet.com to learn more!





Thank You!



My Truong Senior Director, Packet Labs mdt@packet.com

NETRONCME

Ron Renwick VP of Products, Netronome ron.renwick@netronome.com