

Alex Saroyan CEO/Co-founder





Talk

Network Architecture for public and GPU-based AI cloud providers

Insights From NVIDIA's Network Automation Partner Friday - 9:30 AM Cloud Monkey Hall

This presentation is about **Networking**

- How to build & operate networks akin to AWS, GCP, Azure.
 On your hardware for your scale
- Specifics of networking for GPU-based AI/ML clouds
- Keep the sales pitch to bare minimum















How to build and operate? Private, Public, or GPU Cloud?

Managed Services for Workloads	VPC	VMs	Storage		K8s	Database	Al Training	Al Inference	
Interface	End-User Interface			C	RM, Billing, R	eporting	Provider Admin Interface		
Automation, Abstraction, & Operations Systems	Compute CloudStack OpenStack Vmware Kubernetes others			Storage CEPH Pure Storage VAST DDN others			Traditional Networking You can't do this with MS Excel and bash scripts		
Hardware	GF	PU & CPU Serve	ers		Storage	9	Netwo	orking	



How to build and operate? Private, Public, or GPU Cloud?

Managed Services for Workloads	VPC	VMs Storag		e K8s		Database		Al Training	Al Inference
Interface	End-User Interface			CRM, Billing, Reporting			Rrovider Admin Interface		
Automation, Abstraction, & Operations Systems	CloudSt OpenSt Vmware Kubern others	Compute CloudStack OpenStack Vmware Kubernetes others		Storage CEPH Pure Storage VAST DDN others			Cloud Networking		
Hardware	GF	GPU & CPU Servers		Storage			Networking		



- Cloud => VPC and underlying Cloud-like Network Automation Decade old now and proven to be very useful technique Even for mid size **Private** & **Public** Cloud Providers
- World-scale web applications => CDNs
- Next networking evolution will be again dictated by a new workload.



H100 Popular GPU





Training Time is what every AI workload cares about AI Training is the most expensive workload by the order of magnitude









Training time = f (#GPUs, Dataset Size, Model Complexity)



- Dataset grows
- Model Complexity grows
- #GPU has to grow

Parallelism works, unless networks gets in the way.

Example:

1024 GPU cluster trains some model in 30 days

If 1 of 1024 GPUs become limited to 25% network bandwidth, then the entire model will take 4 months to train.



Scale Out Networking Multiple GPU Servers





AI Training Parallelism requires 400Gbps GPU-to-GPU bandwidth If one link is slower => entire training slows down to the slowest GPU





400 Gbps, non-blocking, every GPU to every GPU is required at any scale Otherwise – \$\$\$ GPUs won't be utilized at full extent



netris

E-W Network - 400 Gbps GPU-to-GPU

Static Single Tenant GPU Cluster Network One-time complex configuration



- 100-200 Gbps per NIC
- Leaf Spine architecture
- Workload management
- Storage access
- Remote access

- 400 Gbps GPU-to-GPU
- Rail-optimized architecture
- Non-blocking fabric
- Memory expansion: RDMA / RoCE



Dynamic Multi Tenant GPU Cluster Network Real-time complex configuration





Challenge #2

How to dynamically provide connectivity between isolated tenant networks and the Internet? Across overlapping IPs and dynamic and growing number of tenants.

Challenge #1

How to build Automation that dynamically re-configures every switch to move around switch port isolations between various groups as tenants come and go.

Small clusters have 1024 switch ports Large clusters have under 1M switch ports



Dynamic Multi Tenant GPU Cluster Network Real-time complex configuration

Large cloud providers invested years to Build a solution for this problem. VPC. You can build, too. Time & resources permitting.



Internet

Challenge #2

How to dynamically provide connectivity between isolated tenant networks and the Internet? Across overlapping IPs and dynamic and growing number of tenants.

Challenge #1

How to build Automation that dynamically re-configures every switch to move around switch port isolations between various groups as tenants come and go.

Small clusters have 1024 switch ports Large clusters have under 1M switch ports



Dynamic Multi Tenant GPU Cluster Network **Real-time Automated configuration**





Thank You Q&A

- Technical questions
- Business reasons/questions
- What it takes to go into production





Thank You Q&A