

GPUs at Scale

Trials of a GPUaaS Provider



How did we get here?


- **AI/LLM Boom** Suddenly every company needed GPUs
- **Problem:** Hyperscalers and Unicorns secured all the GPUs, but startups still need access.
- **Solution: Specialised GPUaaS platforms**



High Performance Multi Tenancy

- **GPUaaS at scale** requires multi-tenancy, but AI workloads aren't peaky like web-hosting
- **Problem:** Noisy neighbour scenarios are very common
- **Solution: NUMA aware partitioning, 1 to 1 hardware allocation with CPU core pinning**


- **Lesson Learned: GPUaaS is not the same as other IaaS workloads**



CPU & Memory Optimisation

- **GPUs don't exist on an island**, CPUs and Memory act as data loaders to feed the GPU
- **Problem:** Stability and performance issues can arise from underperforming CPUs and Memory
- **Solution: NUMA partitioning, not scrimping on CPU/Memory spec, boosting CPU clock, hugepages in memory**

- **Lessons learned: AI workloads need fast supporting infrastructure**



Inter-GPU Fabrics


- **Multi-GPU workloads** perform better with NVLink
 - **Problem:** Multi-tenancy on NVLink enabled nodes could not be done in a performance optimised way
 - **Solution: Restrict NVLink instances to whole-node flavours**
- **Lessons learned: Prioritise performance and user experience**



Interconnect Performance


- **Scaling beyond a single node** needs high-performance east/west traffic
- **Problem:** Multi-node jobs that aren't embarrassingly parallel run badly
- **Solution: High performance ethernet and SR-IOV VFLog**

• **Lesson Learned: Think at scale**



GPUs Fail


- **When you are running infrastructure** into the 1,000s of GPUs hardware failures become a reality
 - **Problem:** Workloads will get interrupted
 - **Solution: Educate users on proper checkpointing, ensure that checkpoints are run on appropriately performant storage**
- **Lessons learned: You have to plan for failure**



Conclusion

- Infrastructure at scale is a challenge
- Learn to think at scale, and plan for it from the beginning

- **(or let someone else do it for you!)**



Thank You

Mischa van Kesteren
mischa.vankesteren@hyperstack.cloud