




# GPUs on OpenStack

Konstantinos Samaras-Tsakiris  
( on behalf of the CERN Cloud team )

## Objective

- Allow spawning OpenStack instances with GPUs
- Evaluate possible technologies and use cases
  - Help from Techlab for resource access and GPU knowledge



## Use Cases


- Collected a few use cases to understand requirements

	GPU	CPU	RAM
CSec	1	8	60
CMS	1	2	20
BioDynaMo	1	2	20

- No current use case for multiple GPUs
- Varying, but mostly low CPU requirements

# Integration 1: PCI Passthrough

- Exclusive access to GPU by single VM
  - Good tenant isolation
- No performance penalty expected
- No control from the host
  - No monitoring?
  - Direct access to the hardware by end users?




## Integration 2: Virtualization

- Share GPU among many VMs
  - Oversubscription of physical resources
- Device visible on host
  - Monitoring possible
- Licensing? (nvidia)

Reference: [https://www.youtube.com/watch?v=Xs0TJU\\_slPc](https://www.youtube.com/watch?v=Xs0TJU_slPc)

# Provisioning

- First node added from the Techlab pool
- Setup fully puppetized
- Flavors express vendor and other details
- 2 CPU/GPU ratios
- CPU matches our “normal” flavors
- Similar solution to AWS
- New generation for each GPU type



	GPU	CPU	RAM
<b>g1.large</b>	1	4	8
<b>g1.xlarge</b>	2	8	15
<b>g1.2xlarge</b>	4	16	30

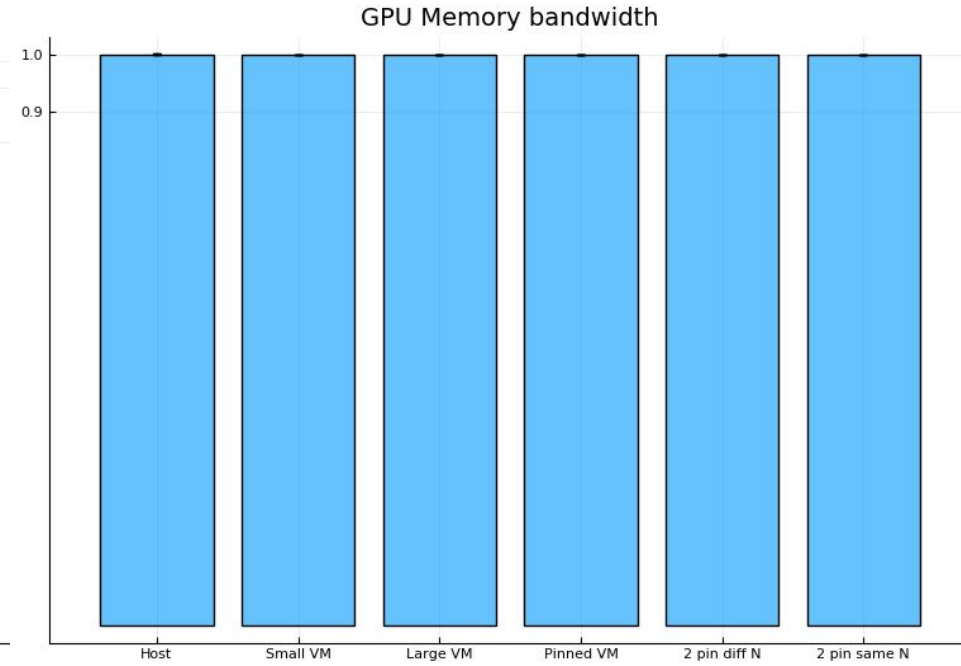
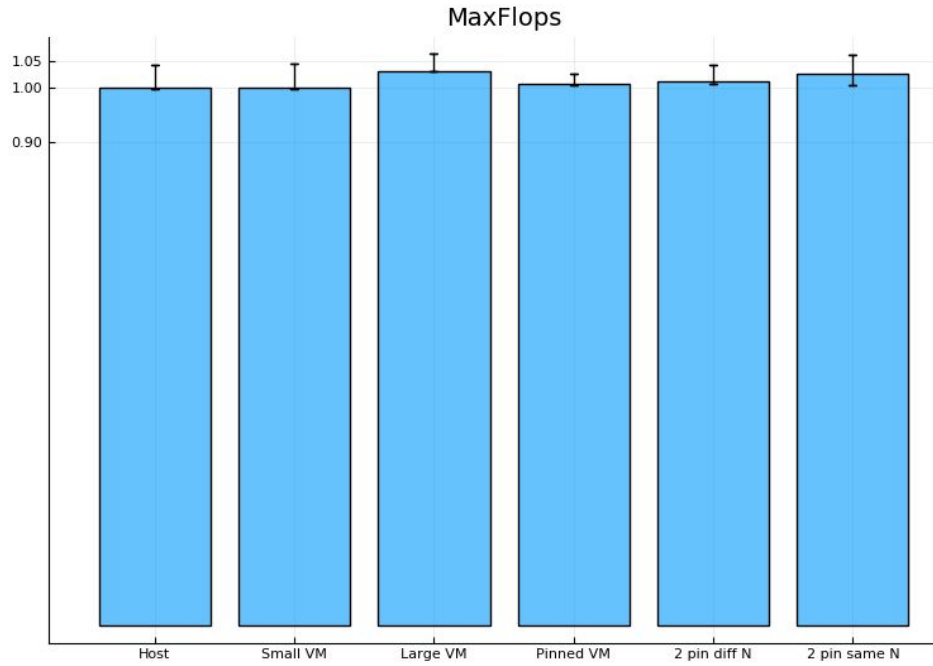
	GPU	CPU	RAM
<b>g2.xlarge</b>	1	8	15
<b>g2.2xlarge</b>	2	16	30
<b>g2.3xlarge</b>	4	32	60

# Benchmark


- SHOC benchmark suite, 3 levels of algorithms
  - Low-level device characteristics, like GPU memory bandwidth
  - Simple calculations, like matrix multiplication
  - More complex programs, like neural network training
    - Not ideal for NN performance measurement!
  - **Small sample size**
- Test cases
  - Host
  - Small VM
  - Large VM
  - Pinned VM
  - 2x pinned
    - Same NUMA node
    - Different NUMA nodes
- Controlled variables
  - Kernel: 3.10.0-693.17.1.el7.x86\_64
  - Nvidia driver: 390.25




# Benchmark Results - Low Level



# Benchmark Results - Simple algorithms



# Benchmark Results - Complex algorithms



## Benchmark Observations

- No penalty on most benchmarks
- Minor improvement in some is probably not significant (small sample size)
  - Uncontrolled variable in host: NUMA node for benchmark
- Non-significant performance boost for larger VM
- Non-significant performance boost for pinning
- Larger variability with 2 VMs on the same NUMA node

## Open Issues & Future Work

- Open Issues
  - Quota Management
  - Instances with multiple GPU
  - Reset GPU state after each use
- Future Work
  - Integrate more Techlab hardware in OpenStack
  - Investigate GPU virtualization options
  - Graphics? Not evaluated yet
- Timelines
  - Initially only integrating Techlab hardware
  - No timelines for more resources or general availability