



Contribution ID: 130

Type: **not specified**

ATS vs IOMMU-regroup: A journey to optimize GDRDMA in cloud scenario

Wednesday 18 September 2024 10:40 (20 minutes)

We encountered a performance bottleneck while testing NCCL on a GPU cluster with 8x H100 GPUs and 8x 400G NIC nodes. Despite a theoretical capacity of 400 Gb/s, our system consistently reached only ~85 Gb/s. The primary issue was identified as communication constraints between GPUs and NICs under the same PCIe switch.

This session will concisely overview the challenges we experienced, such as the PCIe switch and NIC firmware issue, full test results, and solutions we explored to achieve ~390 Gb/s. Furthermore, we will focus on the root cause related to IOVA to HPA translation, evaluating the potential solutions we tried: Address Translation Services (ATS) and IOMMU regrouping.

We would like to hear comments on the pros and cons from the kernel and vendor experts and discuss further to find a better solution.

Primary author: YAN, Liang

Presenter: YAN, Liang

Session Classification: VFIO/IOMMU/PCI MC

Track Classification: VFIO/IOMMU/PCI MC