Inter-Cloud Infrastructure

Cloud Infrastructure for Big Data Analysis

Inter-Cloud Infrastructure for Big Data Analysis

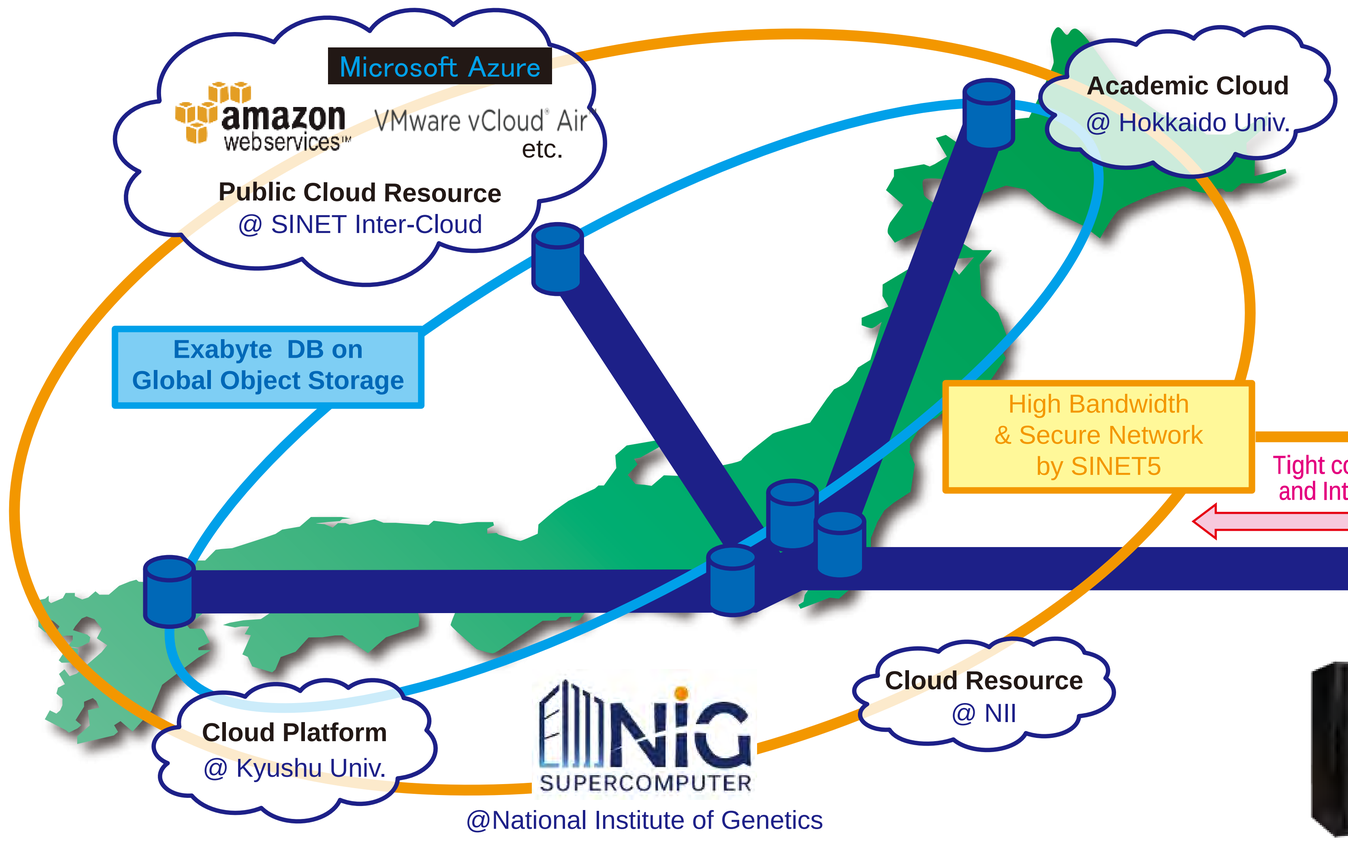

Building a Testbed Infrastructure on Overlay Cloud

- Using SINET5 network infrastructure (100Gbps Network)
- Cooperation with Colud and Supercomputer
- Providing the Science Data Repository
  - Testbed: Peta Bytes class object storage
  - Real System: Exa Bytes class object storage

Cloud and Supercomputer Federaton

- Development of collaboration technology between cloud and supercomputer using TSUBAME 3.0 as a test case

Tight connection of Supercomputer and Inter-Cloud resources via SINET L2VPN

Docker Container

Apply cloud container technology to Supercomputer for resource cooperation with the Inter-Cloud

TSUBAME3 Bigdata Storage

TSUBAME3.0
@Tokyo Tech.

HuronFS: New HPC Storage Hierachy for HPC & Cloud

Background

- HPC
  - HPC users can cause cngestion on PFS
  - Not all the HPC systems have fast PFS
- Cloud
  - Cloud storage cannot provide enough I/O throughput for data intensive applications.
  - Loose consistency model in cloud storage

System

- HuronFS (Hierarchical, UseR-level and ON-demand File system) as a new tier in torage hierarchy.
- Several dedicated nodes/instances as burst buffers to accelerate accesses by buffering intermediate data.
- Implemented on CCI : Utilizing the high performance network in HPC & High portability
- Implemented with FUSE : Supporting POSIX & No code modification required

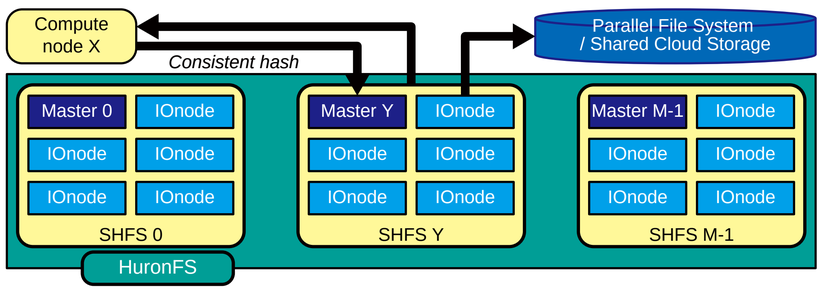

Multiple Client Sequential I/O

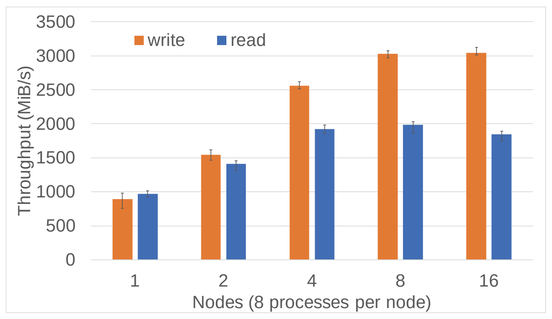

Metadata Performance

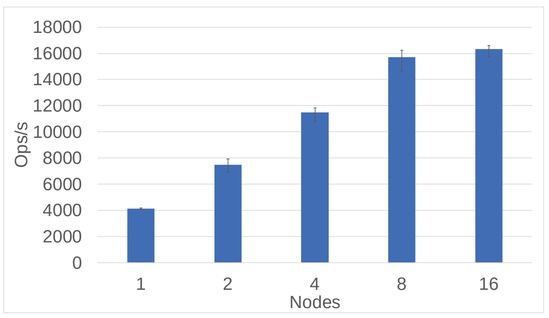

Code available at:
https://github.com/EBD-CREST/HuronFS

Evaluation of HPC-Big Data Applications

Backgroud

Cloud platforms exhibit elasticity, flexibility, usability and scalability, which have been attracting users to use these environments as a cost effective measure to run their applications or businesses. However, the feasibility of running high performance computing applications on clouds has always been a concern mainly due to virtualization overheads and high-latency interconnection network.

Goals

- To investigate the potential role of these virtual machines in addressing the needs of HPC and data-intensive workloads

Experimentation

- Performance evaluation of applications on AWS C4 instances against the baseline results of a supercomputer, TSUABME-KFC

NICAM-DC-MINI: A compute-intensive miniapp from Fiber miniapp suite

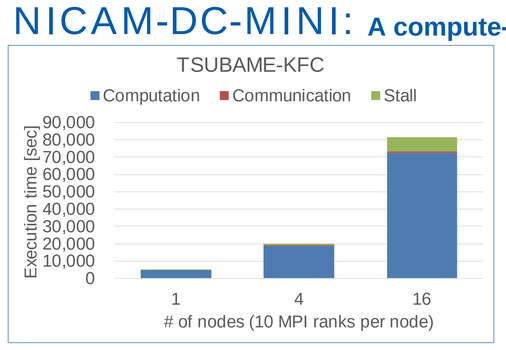

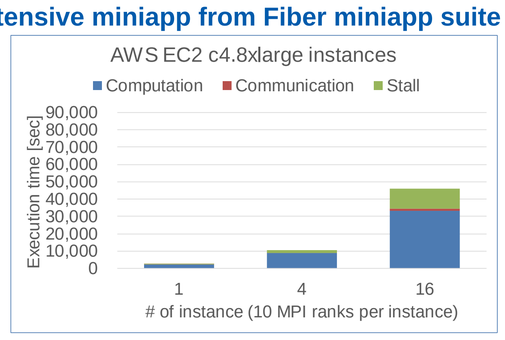

Graph500 Benchmark at Scale 26

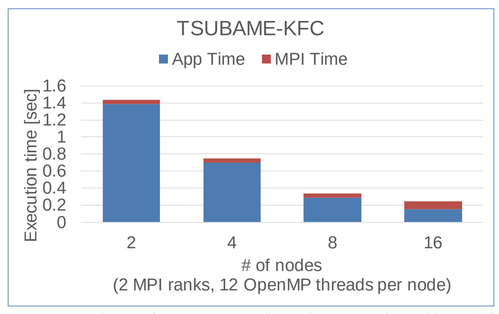

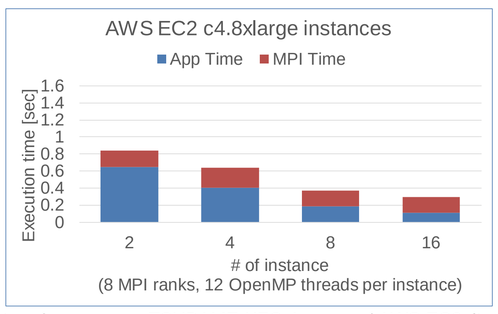

* 2 MPI ranks and 12 openmp threads per node achieved the best performance on TSUBAME-KFC. In case of AWS EC2, it was 8 MPI ranks and 12 openmp threads per instance.

TSUBAME-KFC: Intel Xeon E5-2620v2 x2, InfiniBand FDR
Amazon EC2 c4.8xlarge instance: Intel Xeon E5-2666v3, 10Gbps Ethernet

Acknowledgments. These researches are supported by CREST, JST *CREST Grant Numbers JPMJCR1303 and JPMJCR 1501.*
(Research Area: Advanced Core Technologies for Big Data Integration).

http://www.gsic.titech.ac.jp/sc17