BRIDGES

PITTSBURGH SUPERCOMPUTING CENTER

Comparing Bridges and Cloud Infrastructure

Building value from data analytics, artificial intelligence, and simulation requires advanced computational and data management infrastructure. Choosing the right infrastructure can dramatically increase efficiency and productivity.

This summary compares today's two most powerful approaches to computational and data management infrastructure:

- ***Bridges*** integrates performance computing (HPC) and high performance data analytics (HPDA), combining decades of experience of developing high-end servers and data storage systems with recent, innovative approaches that emphasize large-scale data. Bridges, a uniquely capable, data- and memory-intensive system designed to integrate HPC with Big Data, deliver HPC-as-a-Service (HPCaaS), and amplify human productivity through intuitive interfaces.
- ***Cloud computing*** emphasizes virtualization of commodity hardware to deliver on-demand access, providing great elasticity to cope with periods of high demand. Examples of cloud computing are Amazon Web Services (AWS), Google Cloud Platform, and Microsoft Azure.

Both HPC/HPDA and cloud computing can be integrated with on-premises infrastructure (such as analysts' laptops or engineers' workstations), and some workloads may warrant using both HPC/HPDA and cloud.

Whether HPC/HPDA or cloud is best for a given job depends on many factors. The following table summarizes some strengths of each approach.

Property	***Bridges*** (HPC/HPDA)	Cloud
Performance	• Maximum, predictable performance due to "bare metal" access. Allocate actual CPUs, not vCPUs.	• Virtualization shares hardware across unknown numbers of client jobs, decreasing performance and performance predictability. Dedicated servers (foregoing elasticity) available at higher cost.
Server (instance) types	• High-end servers provide features such as very large memory (up to 12TB RAM), up to 352 cores per servers, and NVIDIA Tesla P100 GPUs.	• A wide range of server types provide from 1 to 128 vCPUs and up to 2TB of RAM, and NVIDIA Tesla K80 GPUs (AWS, as of 3/29/2017).
Cost	• Access for open research is available at no charge. • Other access is available to industrial affiliates, government agencies, and other entities on a cost recovery basis to foster discovery and innovation. • No charge for data transfers.	• Price schedules and cost calculators are available at each cloud's website. • Cost based on instance type, number of instances, on-demand vs. reserved, payment schedule, storage type and amount, and data transfers.
Elasticity	• Limited. Well-suited to applications that need tens to thousands of CPU cores and up to approximately 1000 TB of storage.	• Effectively limitless. Pricing for on-demand instances is typically higher than for reserved instances.
Software	• Large base of pre-installed applications & libraries.	• Installed by user after provisioning servers.
Security	• Secure solutions available.	• Secure solutions available.
Simple provisioning	• 4 server types/rates, 1 storage rate, no data transfer charge. • Talk to PSC for flexible configuration and a quote.	• Up to 74 server types/rates, 5 storage types/rates, 6 data transfer rates. • Cost calculator and web application.
Solutions planning	• Highly experienced PSC staff available for analysis, planning, and implementation at all levels. • Connections to leading research faculty at Carnegie Mellon, Pitt, and elsewhere.	• Basic support for IaaS or PaaS.

For additional information or to talk to PSC about solutions, please contact Cheryl Begandy, Director of External Relations and Outreach, at begandy@psc.edu or 412-268-5129