

Bare-Metal High Performance Computing in the Cloud

On June 8, 2018, the world's fastest supercomputer, the IBM/NVIDIA **Summit**, began final testing at the Oak Ridge National Laboratory in Tennessee. Peak performance is over 200 petaflops (2×10^{17} 64-bit floating point calculations) and 3.3 exaflops (3.3×10^{18}) for tensor operations. Previously untouchable problems will include deep learning for patterns in human proteins, analysis of the entire U.S. cancer population, supernova models running more than 100 times faster than previously, and materials simulations increasing from tens of atoms to hundreds using first-principles calculations at subatomic levels.¹

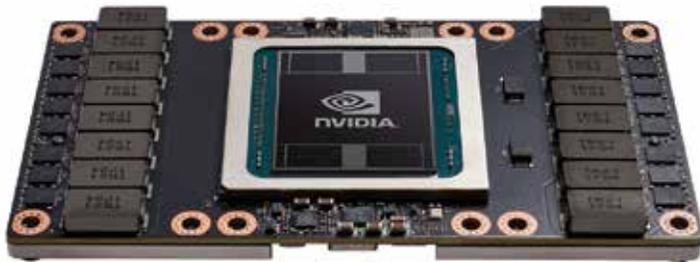


Figure 1. NVIDIA® Tesla® V100 Accelerator

Summit is bare-metal high performance computing (HPC) in the cloud. The US Department of Energy owns Summit, access will generally be remote, and applications will run directly on the entire “bare metal” machine under Linux (not on a virtual machine).²

Summit cost \$200 million to build and has an estimated operating cost of over \$20 million/year. 2018 is fully booked. Applications to use Summit in 2019 close the end of June 2018. To try for 2020, apply to the US Department of Energy in 2019, be one of a few dozen projects selected for all DOE HPC machines—awarded mostly to universities and the US Government, and wait.

Using the same NVIDIA technology as Summit, SkyScale makes bare-metal affordable commercial extreme performance computing in the cloud available now.

This white paper introduces different types of cloud computing, presents benefits and disadvantages and shows how “bare metal cloud” mitigates the disadvantages, and describes SkyScale’s Accelerated Cloud Solutions.

1. Oak Ridge National Laboratory. “ORNL Launches Summit Supercomputer.” ORNL News Desk. June 8, 2018. For more on Summit, see the NVIDIA infographic at <https://images.nvidia.com/content/pdf/worlds-fastest-supercomputer-summit-infographic.pdf>.

2. Wright, Chris. “How Red Hat helped to build Summit, America’s top science supercomputer.” Red Hat Blog. June 8, 2018.

Cloud Computing

Cloud computing generally means accessing computing servers, storage, and applications as a service over a public network—the Internet cloud. With traditional computing over a private network, an organization purchases its own IT equipment, pays employees to operate it, and pays network providers for external communications. With cloud computing, a cloud service provider purchases and operates the equipment, manages the computing resources, and provisions the resources on request from customers, generally using a “pay as you go” pricing model. **The customer application often runs on virtual machines, and the customer usually doesn't know what or where the actual equipment is.**

Vendors also offer on premise private clouds that are located at a customer site but are owned and operated by the vendor who guarantees a level of service; and also offer equipment and operations located at a vendor site but dedicated to a customer; and offer hybrids that combine public and private clouds to enable data and applications to use combined resources owned and operated by both cloud vendor and customer. A hybrid might be used to add resource flexibility by extending a private cloud, or to increase security by storing sensitive data on only the private part.

Cloud Benefits

Key benefits to a simple public cloud are trading off up-front capital cost to buy equipment and operating cost to run and keep it up-to-date and secure versus longer-term services cost; resource flexibility—large cloud companies can provide resources in minutes; for global customers, access to resources when and where needed; service reliability—downtime of an hour or less per year is common and downtime of a few minutes per year or better is available at higher cost; and business continuity, with cloud providers having high assurance of backup and recovery from geographically-dispersed locations. Security is mixed—public cloud service providers are attractive targets for attack, and in a typical multi-tenant architecture with multiple users on the same resources, there is potential for cross-user attack and leakage; however, IT security is a huge, ever increasing challenge, and the personnel and capital resources a large cloud service provider devotes to security are far greater than that of a typical cloud service customer.

Cloud Downsides

The key disadvantage of standard cloud computing for high performance computing users is that HPC resources are generally not available in the cloud! Cloud service providers are most proficient at offering resources for typical business and high-volume transaction applications, and specialized HPC resources are of limited availability.

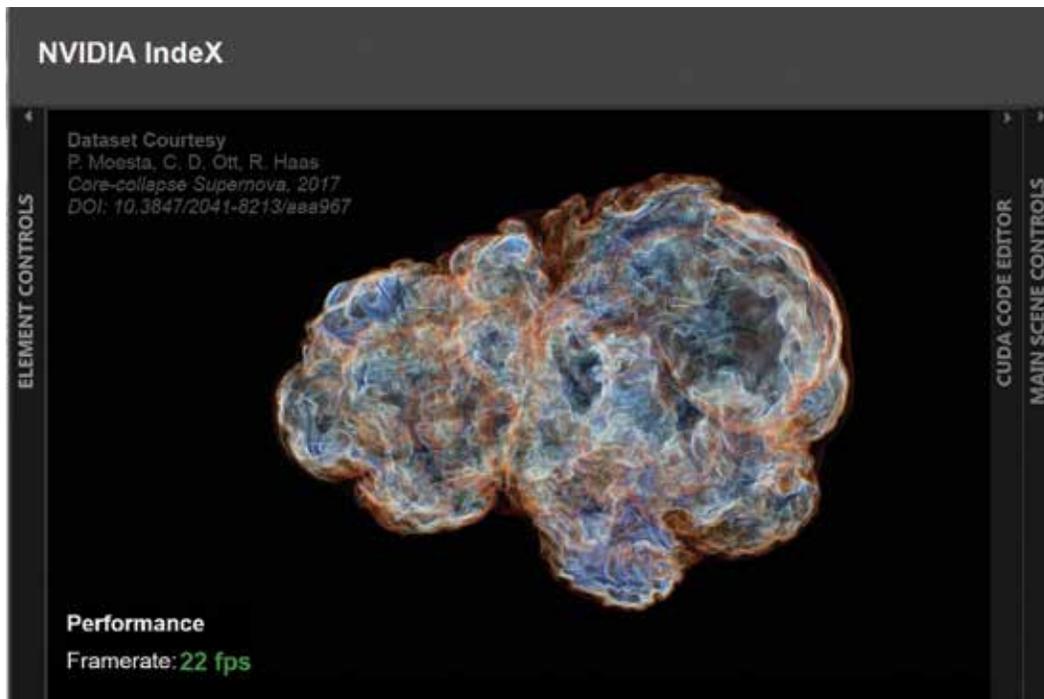


Figure 2. Core-collapse supernova image run on a SkyScale Volta™ V10016xP 16-GPU node. (Described at <https://developer.nvidia.com/index>.)

Other disadvantages to standard cloud computing, whether with larger cloud service providers or smaller vendors (there are dozens), include issues due to other users on the same hardware, limited control and flexibility—the vendor may not permit installation of particular software or desired configuration of hardware or operating system; unexpected startup cost—moving your applications to a particular provider may be expensive because the environment may be different than the prior environment; lock-in after taking advantage of a vendor's features may make it expensive to switch; unexpected operating cost as processor, storage, or data transfer needs increase; uneven support quality, especially for

large cloud service providers who necessarily have a range of support personnel, and even when available, the exceptional support required by high-performance computing users may require frequent time-consuming escalation. Finally, the claimed benefits may not match reality.

Bare-Metal Cloud

Dedicated, bare-metal cloud computing with the right provider mitigates many of the disadvantages above. With a bare-metal public cloud, the customer “rents” remote computing resources that are purchased, managed, and provisioned by the cloud service provider. There is no virtualized environment, no concern for multiple tenants on the same hardware fighting for resources and causing excessive latency or inconsistent performance (runtimes that double are reported), reducing communications bandwidth, or compromising security. Depending on the provider and needs of the application, data may reside permanently on storage equipment at the provider or may be loaded over the network to begin execution and unloaded when complete.

Because the cloud service provider is not concerned about the impact of one customer on another, the provider can offer the customer direct access to the hardware and ability to make operating system changes and install software, enabling customers to maintain their existing work flows.

Lastly, bare-metal cloud service providers are generally highly focused on the needs of their particular customers and the exact resources and configurations those customers require, and can offer high levels of support consistent with that customer focus.

SkyScale High Performance Computing in the Cloud

Many bare-metal providers offer just that—access to generic bare-metal CPUs and storage with an unmanaged, complicated interface for setup and operation. SkyScale offers high-performance computing equipment, Linux or Windows-based operating environment for easy remote access, extensive HPC software libraries (see Software Support next page), and support from experienced engineers to help its customers succeed.

Figure 3 shows the main elements of the most powerful of SkyScale’s Accelerated Cloud Platforms. The heart of the system is the **same multi-core Graphics Processing Unit (GPU) accelerators used in the Summit supercomputer, the NVIDIA Tesla V100** pictured in Figure 1. Originally developed for graphics processing in 3D game rendering, GPUs, with

their high-density, many-core, single-instruction multiple-data architecture (SIMD, extended to SIMT: Single Instruction Multiple Thread by NVIDIA) were found to be highly suitable for other problems subject to massive parallelism.

With multiple GPUs, high-bandwidth efficient GPU-GPU and GPU-CPU interconnect is critical to performance for some problems. SkyScale platforms use one of two options: traditional PCIe or NVIDIA’s new NVLink high-speed interconnect, with up to six NVLink links and total bandwidth of 300 GB/sec in a V100.

For customers requiring more than a single 16-GPU node, SkyScale can interconnect nodes in a cluster using InfiniBand® and Remote Direct Memory Access (RDMA) Ethernet technology.

Metric	Tesla V100	16-GPU Node
NVIDIA Tesla V100 Module with Volta GPU	1	16
Volta Streaming Multiprocessors (SM) / GPU	80	1,280
FP32 Cores / SM	64	1,024
FP32 Cores / GPU	5,120	81,920
FP64 Cores / SM	32	512
FP64 Cores / GPU	2,560	40,960
Tensor Cores / SM	8	128
Tensor Cores / GPU	640	10,240
Peak FP32 Teraflops	16 Tflops	224 Tflops
Peak FP64 Teraflops	8 Tflops	112 Tflops
Peak Tensor Teraflops	125 Tflops	1,792 Tflops
Memory Interface	4096-bit HMB2 ¹	
Memory Size	16 GB	256 GB
Control processor (CPU)	Dual Intel Xeon (3.2 GHz)	
System Memory	512 GB	
System Storage	1.6 TB NVMe SSD ²	
1. HMB2 = High Bandwidth Memory 2nd Generation		
2. NVMe = Non-Volatile Memory Express Solid State Storage		

Figure 3. SkyScale NVIDIA Volta V10016xP 16-GPU Accelerated Cloud Platform specification

Results

The performance of the SkyScale V10016xP exceeds everything else available commercially. *Figure 4* graphs images per second processed by a TensorFlow™ application against number of GPUs. The top line shows the SkyScale Volta system with PCIe interconnect at 8, 12, and 16 GPUs. The bottom line shows the lower performance of a comparable system from a major cloud service provider at 8 GPUs; the plot for 12 and 16 GPUs is “Not/Available”—no other HPC cloud service provider offers more than 8 GPUs on a node, giving SkyScale the fastest machine learning performance available in a single node in the cloud.

Figure 2 is a screen grab of an image of a core-collapse supernova rendered by the NVIDIA IndeX™ visualization application running on a SkyScale V10016xP 16-GPU node. Moving the “timestep” slider in the application regenerates the image continuously on the SkyScale node. The figure shows performance of 22 frames-per-second and the subjective impression is that there is no perceptible delay.

Why SkyScale

The bare-metal cloud market is growing rapidly. From under \$1B in 2016, it is forecast to approach \$5B by 2021.³ For most users, the reasons are need for flexible control over equipment, operating environment, and workflow; performance issues with multi-tenant virtualized environments (CPU and bandwidth consistency); concerns about privacy and security; and—for those requiring high performance computing, the lack of HPC equipment in the cloud and effective support.

In mid 2018, the system described in *Figure 3* is available only from SkyScale. Major cloud providers are beginning to consider or offer the ability to run HPC workloads, but none offer this level of performance, based on 16 NVIDIA Tesla V100 accelerators with a total of 81,920 cores ready for immediate use, and none offer the experienced deep support that flows from SkyScale’s exclusive commitment to HPC.

Software Support

High performance computing applications on many-core system require parallel programming, a discipline new to many engineers. SkyScale supports both new and

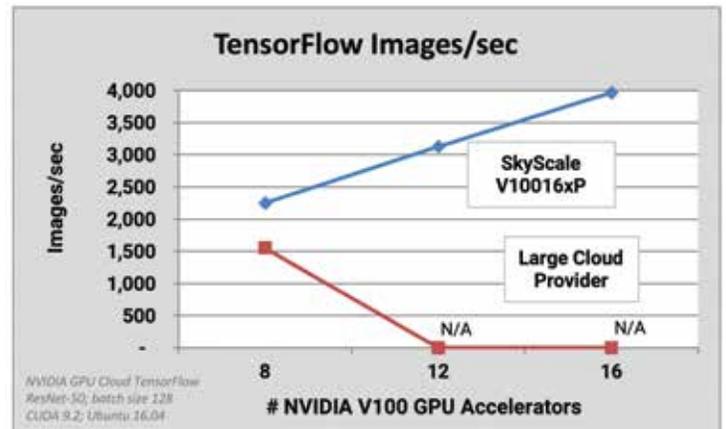


Figure 4. SkyScale V10016xP execution compared to a major cloud services provider

experienced HPC developers and users with machine learning framework options such as Caffe, TensorFlow, Theano, and Torche, plus preinstalled machine learning libraries that include MLPython, cuDNN™, DIGITS, Caffe on Spark and more.

With the V100, NVIDIA has released Volta-optimized versions of GPU accelerated libraries such as cuDNN, cuBLAS™, and TensorRT™ that leverage the new architecture to deliver higher performance for both deep learning training and HPC applications, and the NVIDIA CUDA® Toolkit version 9.0 includes new APIs and support for Volta features to provide even easier programmability.

NVIDIA’s *GPU-Accelerated Applications Catalog* includes over five hundred applications across nearly two-dozen industries. Many are available at no cost. Search by industry, category, and keyword or download the full 48 page catalog.⁴

Customer Support

SkyScale “bare metal” does not mean customers are on their own. SkyScale includes, at no additional cost, direct support from engineers, both pre and post-sale, including help with configuration and with tuning to maximize performance. SkyScale’s goal is to reduce the complexity of HPC so that customers succeed.

3. Patrizio, Andy. “Why a bare-metal cloud provider might be just what you need.” *Network World*. March 8, 2018.

4. Search at <https://www.nvidia.com/en-us/data-center/gpu-accelerated-applications/catalog/> and download at <https://www.nvidia.com/content/gpu-applications/PDF/gpu-applications-catalog.pdf>.

SkyScale Security

Security—cyber and physical—must be a core competence for a cloud service provider. Cybersecurity is an absolute requirement of remote customers for data in transit and for data at rest on provider resources, and customers must be able to trust the physical security of provider data centers while their applications and data are resident there.

SkyScale's deploys enterprise-grade intrusion prevention, detection, and recovery systems and monitors them 24x7.

Its datacenters have manned security 24x7, with biometric identity verification and HD camera coverage.

SkyScale Partners

SkyScale partners expand its usefulness to its customers.

- One Stop Systems develops the computing and flash storage systems that SkyScale deploys to its customers through the cloud.
- Rescale has incorporated SkyScale Accelerated Cloud Platforms into its massive cloud-based simulation platform. This enables SkyScale/Rescale customers to access SkyScale systems by the hour.

Easy to Use, Flexible Provisioning

SkyScale requires no complex setup, no challenging configuration across multiple locations—log in and go—by the hour (with a SkyScale partner), week, month, or year.

Try SkyScale Now at No Cost

The Summit supercomputer is extraordinary—but will be available to very few. For users who want to experience the flexibility, customizability, performance, and affordability of 16 NVIDIA Tesla V100 GPUs, with 81,920 cores delivering 224 teraflops of 32-bit floating-point, 40,960 cores at 112 teraflops of 64-bit floating point, and 10,240 Tensor cores at **1.8 petaflops** now, SkyScale offers a no-cost trial with engineering support. For information, contact <http://www.skyscale.com/contact/#freetrial>.