


Build a flexible and powerful High Performance Computing foundation with Google Cloud





Volker Eyrich

Customer Engineer,
Google Cloud



Joshua Fryer

Senior HPC Engineer,
Recursion

Google Cloud helps meet your HPC needs with the fastest time to insight and simplicity



Fast and Affordable

- Latest HPC technologies
- Flexible VM shapes
- Batch VM types
- MPI scalability



Easy and Compatible

- Simple and easy to use
- Turnkey HPC environments
- Compatible with leading apps and schedulers



Partner Preferred

- Broad ISV application support
- Broad network of Cloud HPC solution providers
- Broad ecosystem of HPC system integrators



Best of Google

- Industry expertise
- Leading ML and Data Analytics
- Hybrid & multi-cloud via Anthos/K8s
- Worldwide cloud network

Deploy and manage HPC environments easily with Google Cloud's comprehensive platform

Google Cloud's HPC Platform



Users

Command Line Interface | Remote Workstations | Workbenches
& Notebooks | HPC Platforms | API-Driven Workflows

Deployment

Scheduler, Deploy,
Integrations,
Monitoring,
Marketplace

Development

Built-in Tools,
Communication
Libraries,
Infrastructure Mgmt

Infrastructure

Compute,
Network,
Storage

Administrators


VM-based | Google Cloud, Multi-Cloud, On-Premises
Container-based | Google Cloud, Multi-Cloud, Private Cloud

Google Cloud HPC Customers

Enterprise HPC Customers



Academic HPC Customers



01

HPC on Google Cloud

Leverage Google Cloud's powerful HPC platform, tools, and partner solutions

Powerful Infrastructure

Compute Engine's VMs feature:

- Latest CPUs and GPUs
- High performance storage
- High throughput, low latency networking
- Live-migration, security built-in



HPC Tools and Services

Easily scale your workloads, run your containerized HPC workloads with GKE, and deploy auto-scaling HPC cluster environments with the Google Cloud HPC Toolkit



Broad Partner Ecosystem

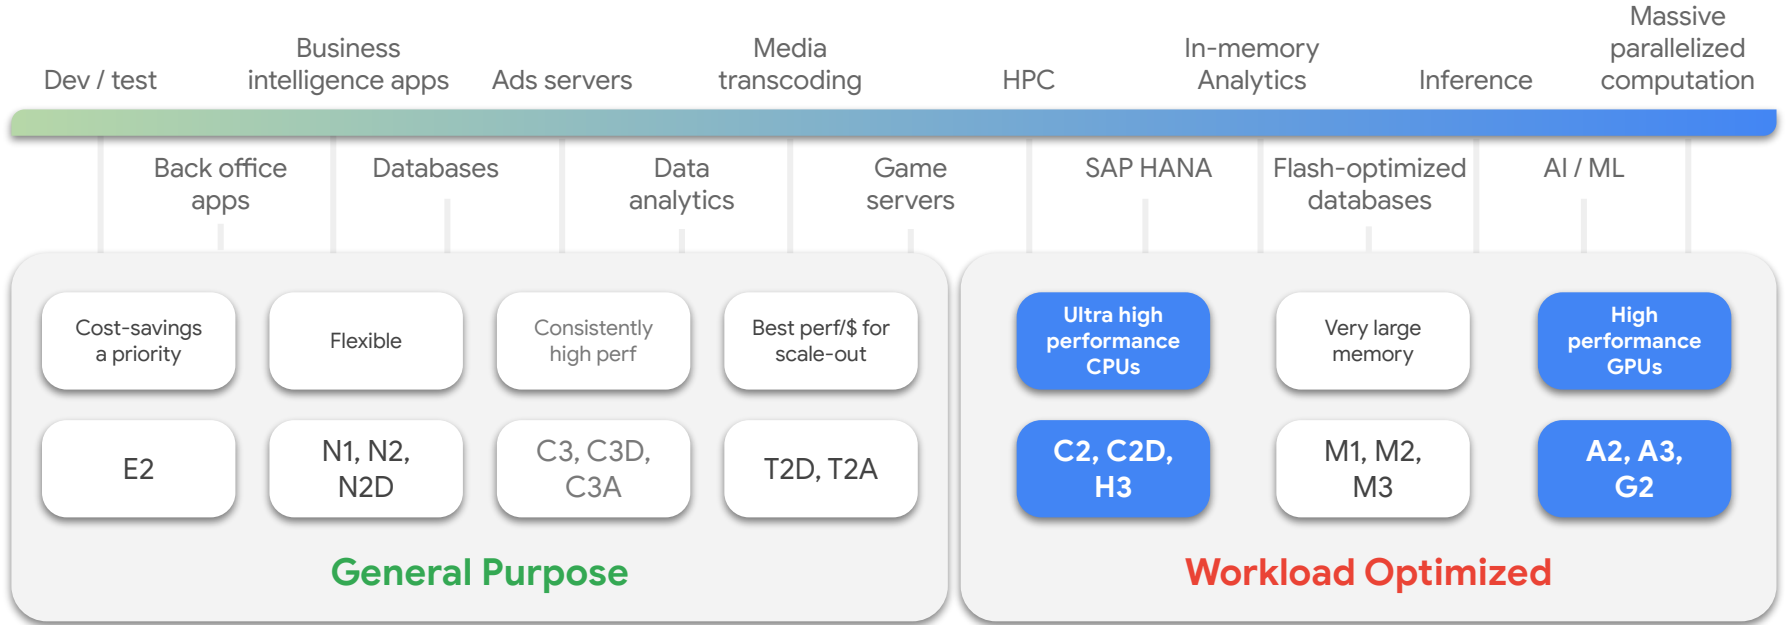
Google partners with a wide variety of application developers, workload managers, storage providers, and system integrators




VM Families **Optimized for HPC Workloads**

Less demanding workloads

Performance intensive workloads




What makes an HPC VM?



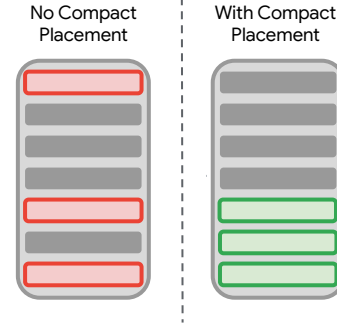
Purpose-built Infrastructure

Google Cloud's HPC VMs have the highest clock speeds and memory bandwidth of any Google VM type, as well as HPC features like disabling SMT. Choose from Intel or AMD CPUs.




Virtual NUMA (vNUMA)

vNUMA provides a direct mapping of Host CPU to Guest CPU and an accurate view of the NUMA layout of the guest on the host. Default on HPC VMs 2nd Gen and above.



Compact Placement

Compact placement policies put your VMs close together for low network latency between the VMs. Supports 100+ VMs per group. Supported by HPC VMs 2nd Gen and above.



Network Optimizations

Google's HPC VMs perform optimally with our MPI best practice tunings in the HPC VM Image combined with our hardware offloads, high bandwidth VMs, and gVNIC drivers.

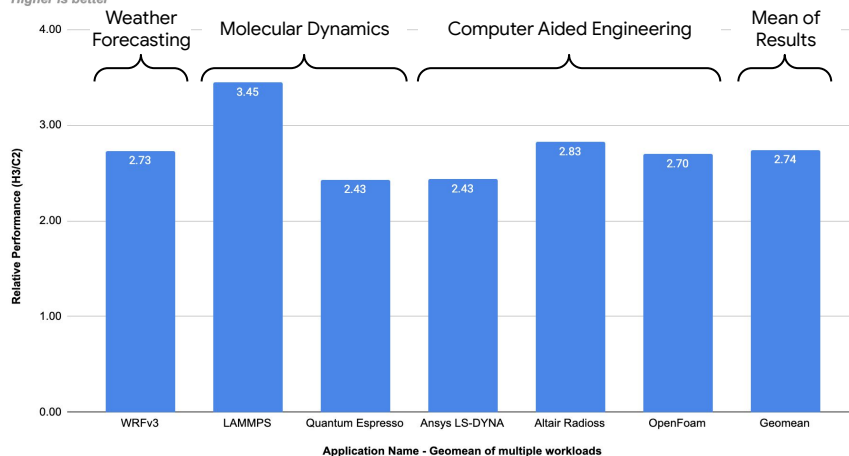
Introducing H3 VMs for HPC

Now in Private Preview!

- Whole machine VM, full-core vCPUs, SMT off, 200 Gbps
- Intel Sapphire Rapids CPU, Intel IPU Hardware Offload
- **Up to 3x better performance** compared to prior gen
- **30-40% off** the price of a similar C3 General Purpose VM

H3 vs. C2 (Single Node)

Higher is better



AirShaper



Parallel Works

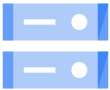








"With the new H3 machine series, we've noticed a **greater than 2.5 times performance increase** when run on multiple nodes, along with a significant **overall job cost savings of ~50-70%**. This makes H3's cost/performance level a necessity for leaders in the numerical weather prediction industry."

- **Matthew Shaxted**
President, Parallel Works

HPC Storage on Google Cloud

Object, Block, and File

Google Cloud Services			
Object	File	Parallel File	File
			
Cloud Storage	Filestore	Parallelstore	NetApp Volumes
JSON/S3 Unstructured data, objects	NFS Enterprise and High perf file storage	POSIX, LibFS, KV Next-gen DAOS-based HPC storage	NFS/SMB Fully managed and featured ONTAP
Persistent >> TB/s with high latency	Persistent 26 GB/s with low latency	Scratch >>TB/s with ultra low latency	Persistent Up to 4.5GBs

Google Cloud Partners		
Partner Parallel File Systems		
		
DDN EXAScaler	Sycomp Spectrum Scale	Weka.io
POSIX Enterprise Lustre	POSIX IBM Spectrum Scale with GCS auto-tiering	POSIX Weka Data Platform with GCS cold tier
Persistent and Scratch 100s of GB/s	Persistent and Scratch Up to 300GB/s	Scratch 10s of GB/s

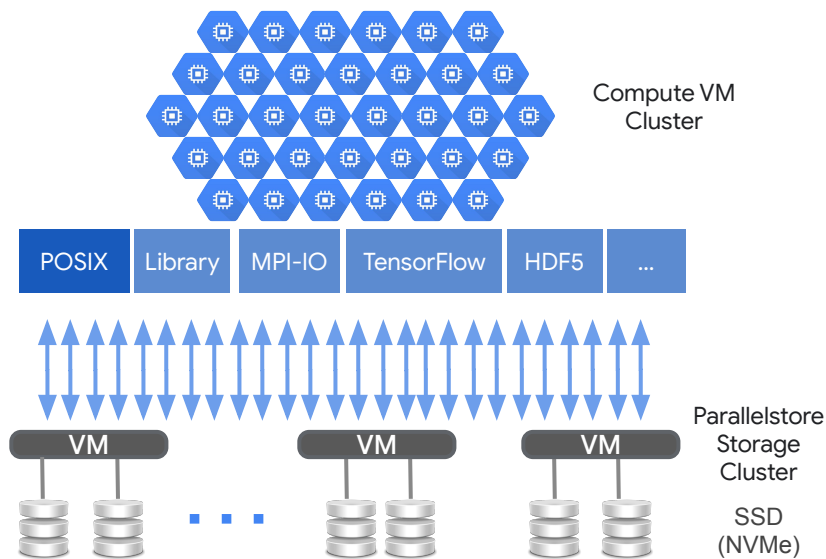
Parallelstore

Next Gen HPC Storage System - Now in Private Preview!

Accelerate HPC and AI workloads that require extreme scale and/or low latency I/O operations

Key Advantages

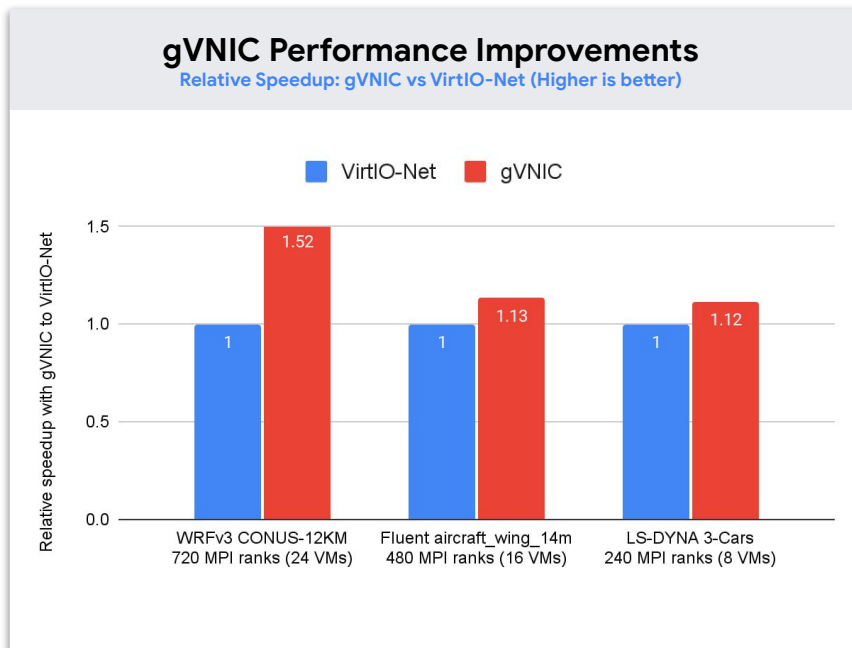
- Unique open-source DAOS storage architecture improves performance over existing POSIX storage options
- Standalone storage, or accelerate access to Cloud Storage
- Well-aligned to emerging patterns in AI workloads with distributed metadata, extreme IOPS, and K/V architecture
- Demonstrated >1GB/s per TiB, >1.5M create/s with open-source on GCP
 - **Up to 6.3x of FSx Lustre** read throughput of 200MB/s per TiB (Burst to 1,300)



HPC Networking on Google Cloud

Scalable, high-bandwidth, low-latency VM networking

- **Scalable VM Bandwidth**
 - Up to **32 Gbps** by default, 2 Gbps per vCPU
 - Up to **200 Gbps** with **3rd Gen VM types**
- **Latency Tuning & Optimization**
 - Predictable, low latency (<10 μ s average)
 - **Google's HPC VM Image** implements our MPI Best Practices and tunings
 - **gVNIC** - Open Source Linux Kernel module built for GCE networking
 - **Placement Policies** allow compact colocation, reducing VM to VM latency
 - **9K MTUs** (Jumbo Frames) increase packet size
- Up to 15,000 VMs per Virtual Private Cloud Network



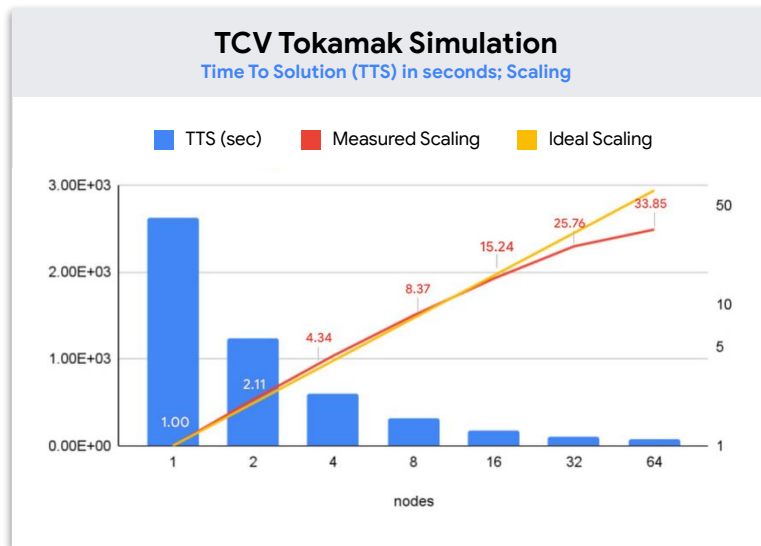
<https://cloud.google.com/blog/topics/hpc/running-mpi-workloads-efficiently-on-google-cloud-using-gvnic>

Swiss Plasma Center on Google Cloud

“We can now deploy a flexible and powerful HPC infrastructure that is virtually identical to the one we maintain at EPFL in **less than 15 minutes** and dynamically offload on-prem workloads in times of high demand.”

“Using Tokamak Configuration Variable (TCV) geometry, our results show excellent scalability: we managed to get a **33X speedup for the TCV tokamak simulation, with a near-perfect scale up to 32 nodes.**”

- Swiss Plasma Center



<https://cloud.google.com/blog/topics/hpc/swiss-plasma-center-uses-google-cloud-to-simulate-fusion>

How to run HPC on Google Cloud



Do It Yourself

Combine Google Cloud's HPC software, infrastructure, and partner solutions on your own to create your specific HPC environment. Configure, deploy, and maintain your HPC system yourself.



Cloud HPC Toolkit

Use Google's powerful, flexible Cloud HPC Toolkit to design and deploy a bespoke HPC environment leveraging Google's partner solutions, with Google's HPC best practices built in.



HPC Platform Partner

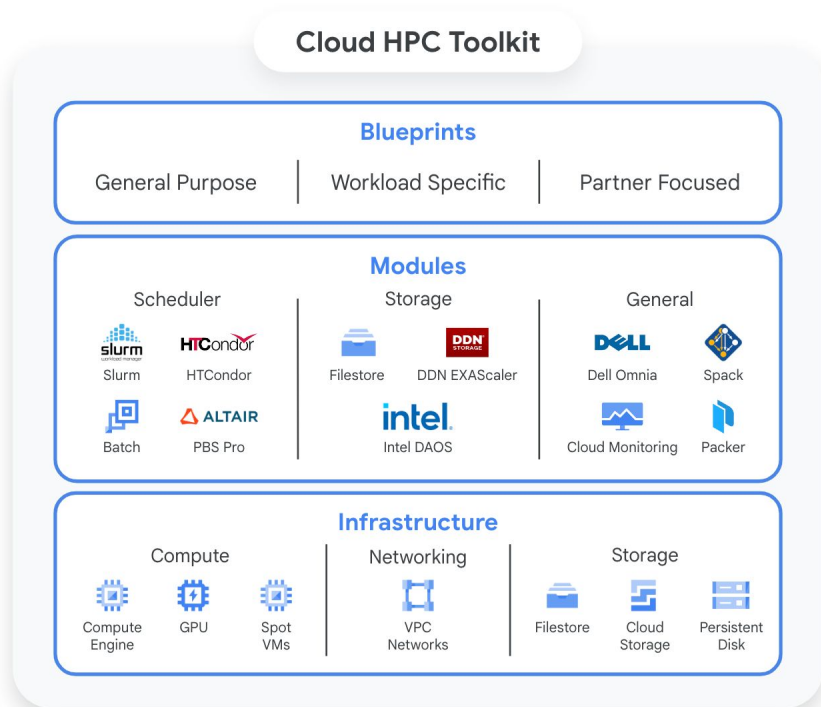
Leverage an HPC Platform provider that partners with Google Cloud to deliver out-of-the-box performance and reliability, combined with an easy-to-use web interface and application menu.

Cloud HPC Toolkit

The Cloud HPC Toolkit is a modular, composable, terraform-based toolkit designed to make it easy to deploy repeatable, turnkey HPC environments that follow Google Cloud's HPC best practices.

Key components:

- **Blueprints** defines an HPC environment by importing and configuring individual modules which the toolkit assembles to build the specified environment
- **Modules** are pieces of code (imported from Github or locally) which deploy specific components of an HPC system (i.e. a project, network, VDI, Slurm controller)
- **Infrastructure** will host the HPC system that is built. The Cloud HPC Toolkit supports Google Cloud's latest and greatest HPC services and features out of the box



Cloud HPC Toolkit **Blueprint**

```
blueprint_name: hpc-cluster-small
```

```
vars:
```

```
  project_id:  ## Set GCP Project ID Here ##  
  deployment_name: hpc-small  
  region: us-central1  
  zone: us-central1-c
```

```
deployment_groups:
```

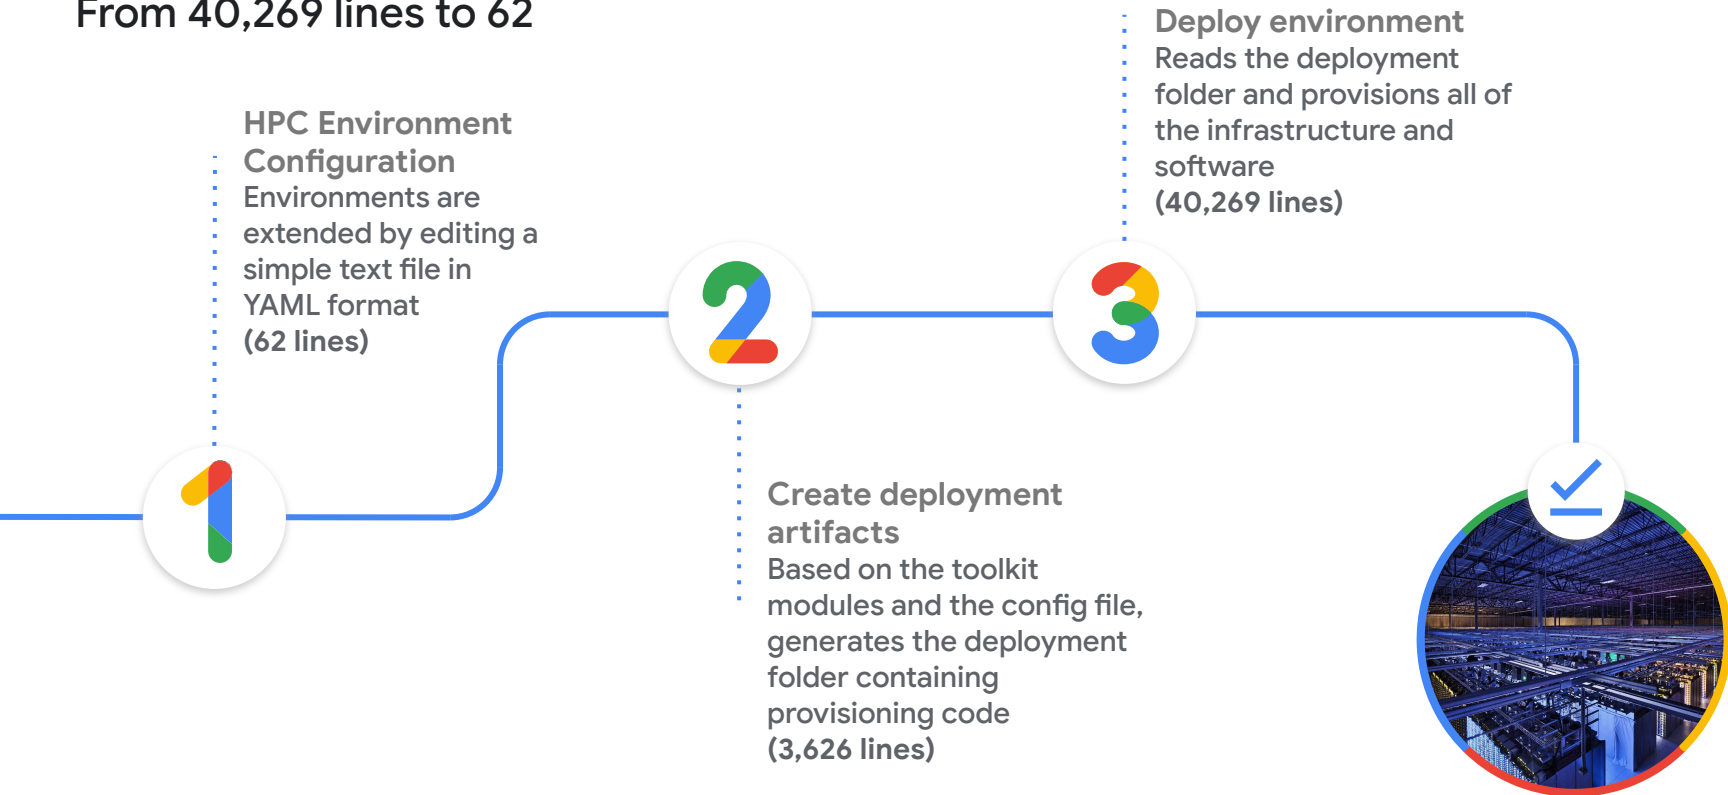
```
- group: primary  
  modules:  
    - id: network1  
      source: modules/network/pre-existing-vpc
```

- id: **homefs**
 source: (...)/file-system/DDN-EXAScaler
 use: [**network1**]
- id: **compute_partition**
 source: (...)/SchedMD-slurm-on-gcp-partition
 use: [**network1**, **homefs**]
 settings:
 partition_name: compute
 max_node_count: 20
- id: **slurm_controller**
 source: (...)/SchedMD-slurm-on-gcp-controller
 use: [**network1**, **homefs**, **compute_partition**]
- id: **slurm_login**
 source: (...)/SchedMD-slurm-on-gcp-login-node
 use: [**network1**, **homefs**, **slurm_controller**]
- id: **hpc_dash**
 source: modules/monitoring/dashboard



Journey to a working environment

From 40,269 lines to 62



Cloud HPC Toolkit Blueprint for GROMACS

What was deployed?

- APIs enabled, VPC network created
- Storage
 - GCS Buckets, Filestore (NFS)
- Spack
 - Builder installs Intel MPI, GCC, and GROMACS to Filestore
- Slurm Cluster
 - VMs: Login Node, Controller
 - Auto-scaling Partitions: CPU, GPU (A2 + NVIDIA A100)
 - Storage Mounted
- Remote Desktop VMs with GPU Acceleration
 - Chrome Remote Desktop, VMD
 - Storage Mounted
- HPC Monitoring Dashboard

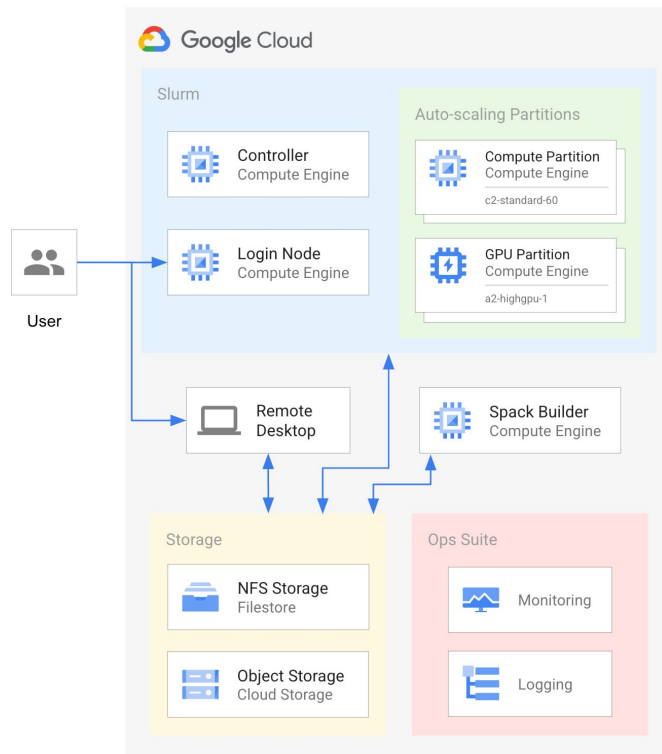


Cloud HPC Toolkit
GROMACS Blueprint



Cloud HPC Toolkit
GROMACS Demo Video

Cloud HPC Toolkit Blueprint for GROMACS



02

Google & SchedMD

Slurm on Google Cloud

Since 2017, Google has partnered with SchedMD to integrate the Slurm Workload Manager to harness the power of Google Cloud

Ways to use Slurm:

- **Cloud Auto-Scaling:** Automatic elastic scaling of instances, on demand, according to queue depth and job requirements. Spins resources down once idle timeout is reached.
- **Hybrid Clusters:** Extend your Slurm cluster to the cloud to dynamically offload jobs to Google Cloud using Cloud Auto-Scaling functionality. Accounting passed through from existing cluster.
- **Enterprise support** available directly from SchedMD
- **Open Source from SchedMD:**
<https://github.com/schedmd/slurm-gcp>



Slurm on Google Cloud v6

Improved feature functionality for Slurm on Google Cloud

- Support for the latest 3rd Gen and Arm VM families (C3, T2A)
- Full integration with the Google Cloud HPC Toolkit
- Enhanced error messages, actionable troubleshooting steps
- Improved job-level monitoring and reporting capabilities
- Reliability & update path enhancements
- Support for Shielded VMs for highly secure environments
- Expanded configuration options for auto-scaled VMs
- Enhancements to Slurm Hybrid capabilities
- Performance tunings for specific Slurm usage patterns
- TPU on Slurm support GA



TPUs on Slurm v1


Support for Google Cloud TPUs in Slurm on GCP

Overview

- Define TPU-accelerated partitions like any other resource type
- Host TPU partitions alongside GPU partitions for AI workloads!
- Developed with input from major AI developers
 - In use today by AI industry leaders, including Midjourney

Features

- Schedule and fair share across large TPU estates
- Deploy individual TPUs or TPU Pods
 - Supports TPU v2 and v3 today, v4+ coming soon!
- Supports PyTorch, JAX, and Tensorflow





RECURSION

Integrating an On-Prem HPC cluster with Google Cloud: Power of Slurm

Why Slurm for our HPC needs?

Easy Integration

Slurm WLM provides **seamless integrations** with Prometheus, Python, Open OnDemand and many more tools and applications

SchedMD Support

Various **storage service offerings** remove much of the burden of building and managing storage and infrastructure.




A Thriving Community

Widespread community knowledge with new Data Scientists and ML engineers already acquainted.

Freedom & Flexibility

Enables **streamlined pre-production testings**, hassle-free licensing, and lightweight operation.



RECURSION • BioHive 1



RECURSION • BioHive 1



RECURSION • BioHive 1

Expanding Horizons with Google Cloud

Google Cloud's HPC Platform



Scalability

GCP Offers real-time resource adjustment, seamless load distribution, and a global infrastructure for consistent performance worldwide. No more problematic nodes and hardware.

Services and Tools


GCP provides powerful ML/AI tools, real-time big data analytics, and integrated cloud-based development environments for enhanced data processing and streamlined workflows. Supports tools like TF.

Cost Optimization


GCP offers a cost-effective model with pay-for-what-you-use flexibility, long-term usage discounts, and customizable resource options to match precise needs.

Harnessing Best of Both Worlds

Cloud bursting at high load



At low load



Thank you.