High Performance Computing on Google Cloud

Nir Oz
Customer Engineer
Google Cloud

Google Cloud

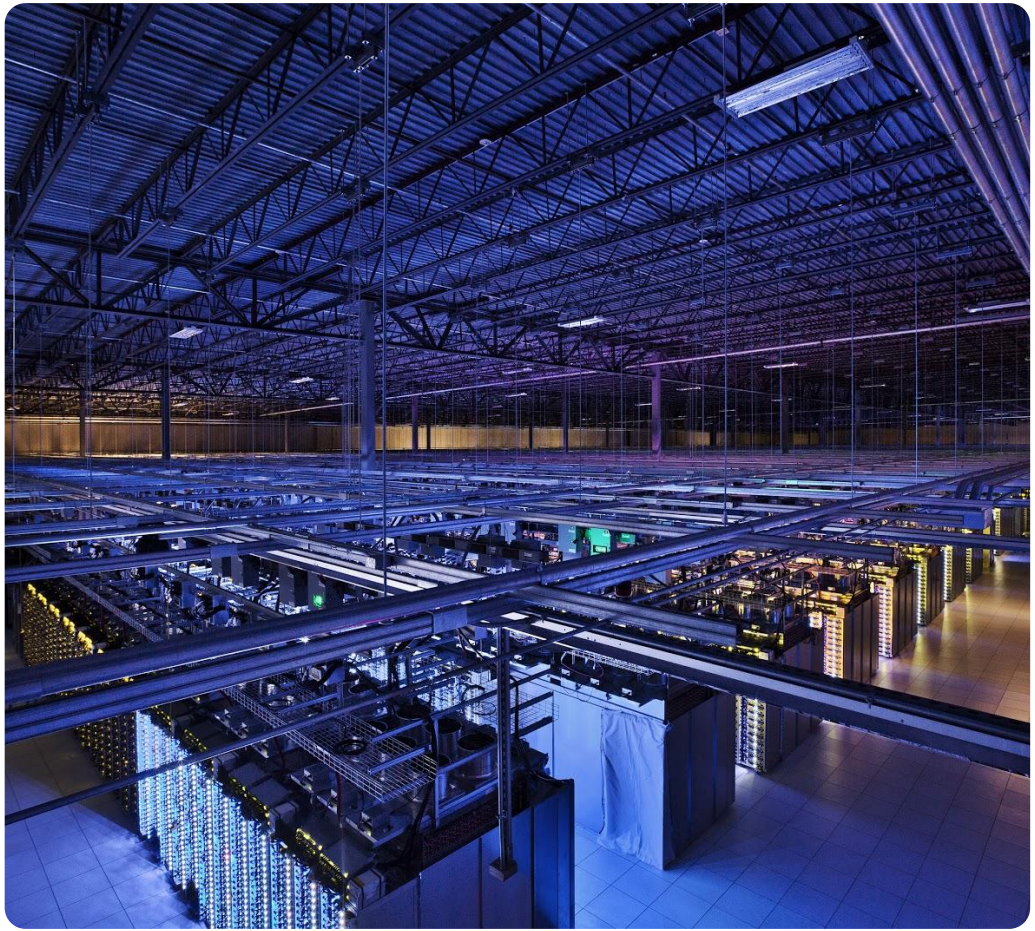

HPC on Google Cloud

Meeting HPC needs with fastest time to insight and simplicity

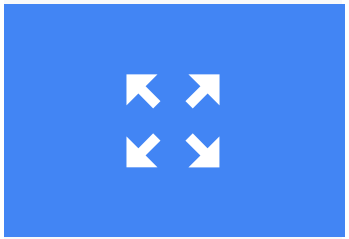

Fast and Affordable

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

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

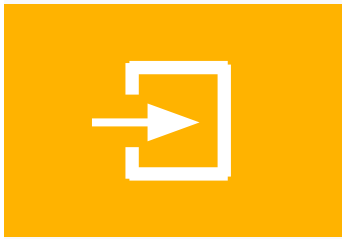

Best of Google

- Leading ML and Analytics
- IoT <-> R&D HPC integration
- Hybrid & multi-cloud via Anthos/K8s
- Largest worldwide network

Investing in HPC Infrastructure and Software

1st = First in cloud

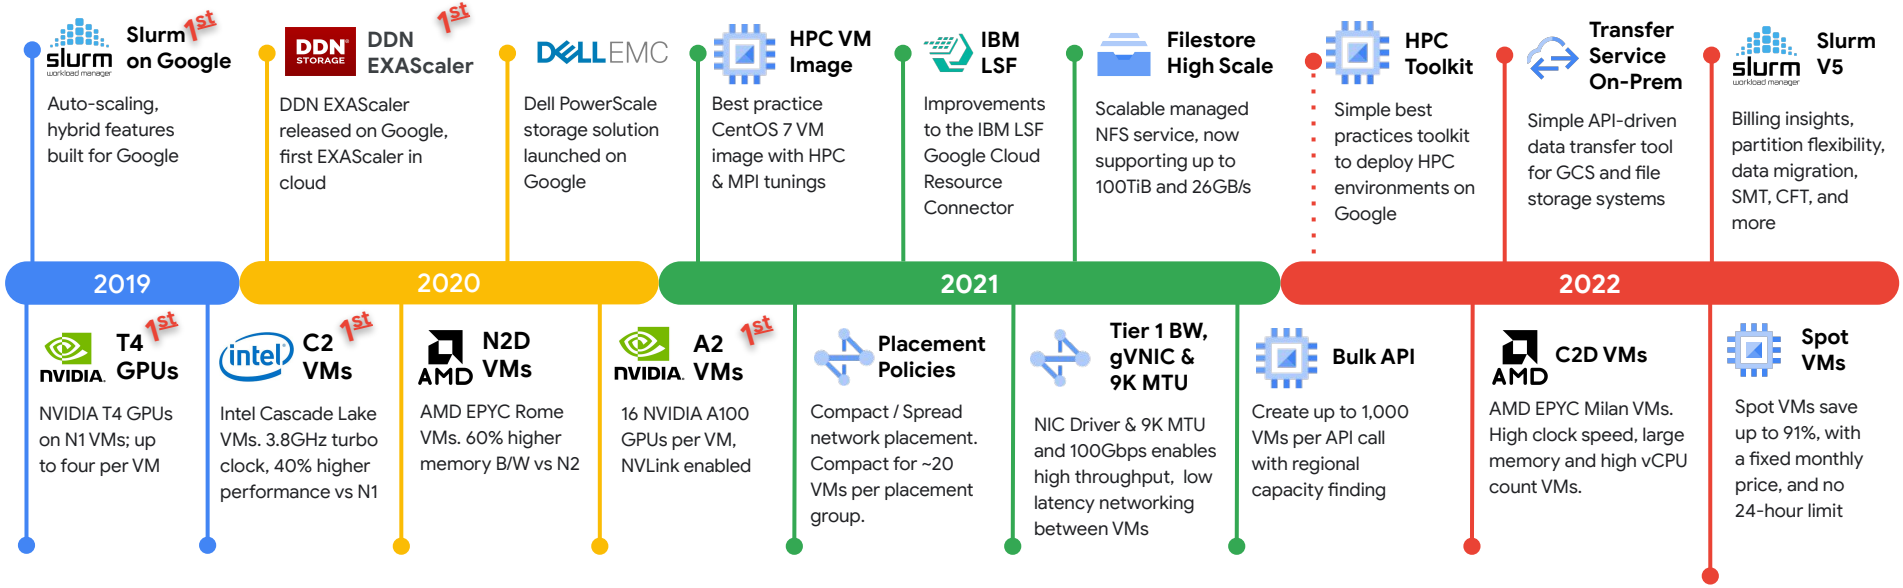

Legend

- Generally Available
- In Preview

HPC Simulation Platform

Google's HPC Simulation Platform

HPC Users | Remote Workstations (VDI) | Jupyter Notebooks | Batch Schedulers | API-Driven Workflows

Deployment

- **Integrations**: ALTAIR, slurm, rescale, IBM Spectrum LSF, HTCondor, Parallel Works, CloudyCluster, NIMBIX
- **Purpose-Built**: Life Sciences
- **Monitoring**: Operations Suite
- **Marketplace**: Cloud Marketplace

Development

- **Built-In Tools**: HPC VM Images, HPC Toolkit
- **Message Passing Libraries**: Intel MPI, Open MPI
- **Infrastructure Management**: Terraform, Deployment Manager

Infrastructure

- **Compute**: CPU, GPU, TPU, Spot
- **Network**: Advanced Networking, Placement Policies, 9K MTU
- **Storage**: Filestore (NFS), Object Storage, Transfer Services, Storage Partners

HPC Admins

- **VM-Based**: Google Cloud, On-Premises
- **Container-Based**: Google Cloud, Multi-Cloud, Private Cloud

Urgent HPC

Most researchers do not have access to the HPC resources needed to study complex systems

Traffic Management is a major hurdle to efficient emergency evacuations

Models of evacuation traffic patterns during hurricanes are valuable for evacuation planning

Reducing traffic congestion saves lives, time, money, and the environment

Google Cloud

Urgent HPC: 2.1M vCPUs

Clemson University broke the world record for running an HPC workload using the most compute cores on any commercial cloud.

2.138M vCPU | 133,573 instances | 200+ TB processed

Average cost of **$0.008 USD per vCPU hour**

Published **“On-Demand Disaster Management using High Performance Computing in the Commercial Cloud,"** in The Journal of Supercomputing

A corpus of [traffic data for public use] is available via the Public Dataset Program

[Google Cloud Blog Post]

[10 Minute Lightning Talk - CloudHub]

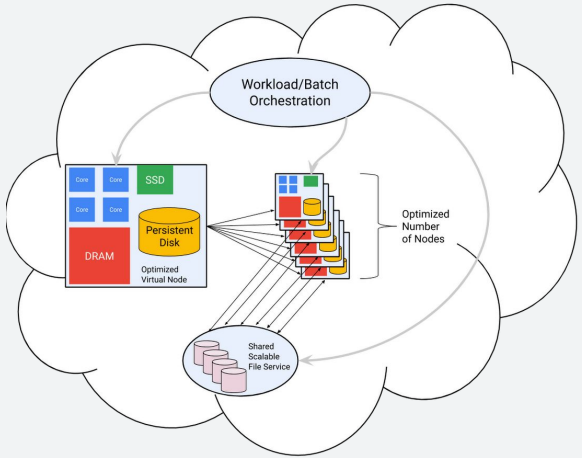

High Performance Computing
on Google Cloud

Compute

Google Compute Engine's VMs boot in seconds, are built for consistently high performance, and have security built in.

Storage

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

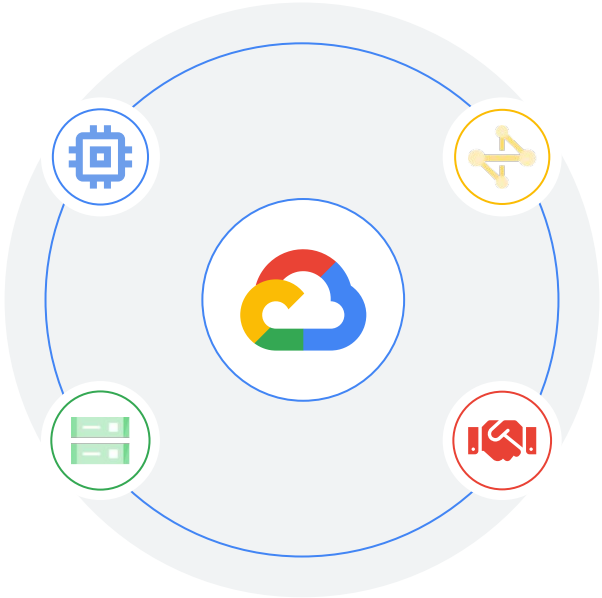

Network

Google's high performance private network connects VMs with high throughput, low latency interconnects.

HPC Software

Google Cloud offers native HPC tooling, and supports a broad portfolio of HPC software from our HPC partners and open source projects.

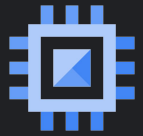

Compute

Google Cloud

VM families for all workloads

HPC Compute Building Blocks

HPC optimized, customizable, scalable, reliable instances

Purpose	General Purpose			Workload-Optimized		
Capability	Flexibility, GPUs	Balanced Performance	High Memory BW	Fastest Clock Speeds	Large, Fast Clocks	Large, A100 GPUs
Hardware Support (Up To)	• 48 cores (96 vCPU) • Intel Skylake thru Broadwell • 624GB RAM • 32Gbps • 9TB Local SSD • NVIDIA K80, T4, P4, P100, V100 • 4 or 8 GPUs • Supports TPUs	• 64 cores (128 vCPU) • Intel Cascade Lake or Ice Lake • 2.6GHz Base, 3.4GHz All-Core • 864GB RAM • 100Gbps • 9TB Local SSD	• 112 cores (224 vCPU) • AMD EPYC 2nd or 3rd Gen • 2.7GHz All-Core • 896GB RAM • 100Gbps • 9TB Local SSD	• 30 cores (60 vCPU) • Intel Cascade Lake • 3.8 GHz All-Core • 240GB RAM • 100Gbps • 3TB Local SSD • Compact Placement • vNUMA	• 56 cores (112 vCPU) • AMD EPYC 3rd Gen • 3.3GHz All-Core • 896GB RAM • 100Gbps • 3TB Local SSD • Compact Placement • vNUMA	• 48 cores (96 vCPU) • Intel Cascade Lake • 3.8 GHz All-Core • 1,360GB RAM • 100Gbps • 3TB Local SSD • Compact Placement • NVIDIA A100 (40GB) GPUs • 16 GPUs • NVSwitch @ 600GB/s • vNUMA
VM Series	**General Purpose (N1)**	**General Purpose (N2)**	**General Purpose (N2D)**	**Compute-Optimized (C2)**	**Compute-Optimized (C2D)**	**Accelerator-Optimized (A2)**

Google Cloud

What makes an HPC VM?

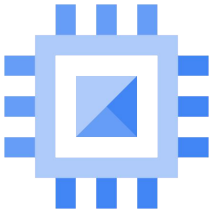

High Performance Compute

Google Cloud's HPC VMs have the highest clock speeds and highest memory bandwidth of any Google VM type. Choose Intel or AMD CPUs. VMs support up to 16 NVIDIA GPUs.

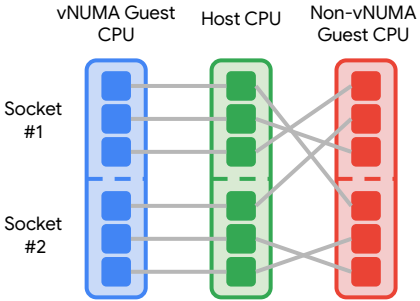

Virtual NUMA (vNUMA)

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. vNUMA is default on HPC VMs (C2*, A2).

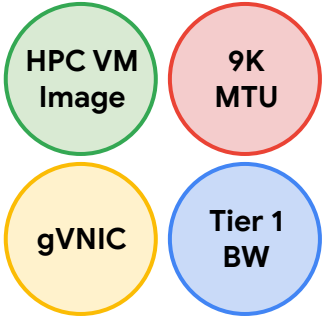

Network Optimizations

Google's HPC VMs perform optimally with the best practice tunings in the HPC VM Image, as well as features like Tier 1 Bandwidth for 100Gbps, 9K MTU support, and gVNIC drivers.

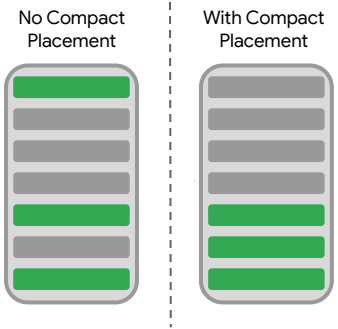

Compact Placements

Compact placement policies put your VMs close together for low network latency between the VMs. Supports 100+ VMs per group. Supported by HPC VMs (C2*, N2*, A2).

Google Cloud

Google Compute Engine - GPUs

- Attached directly to the VM via PCIe x16, with NVLink and NVSwitch to achieve the best possible performance
- Per-second billing, Preemptible support (~70% off)
- First to market with NVIDIA A100, T4 GPUS

Google Cloud

GCP A2 VMs INCLUDE UP TO 16 GPUs

Introducing the *A2-MegaGPU-16g* VM

- **2 HGX Boards**
- **2.4TB/s total NVLINK BW** w/ 16 A100 GPUs
- ***ONLY* CSP** offering up to 16 GPUs in a single node
- Effective performance up to 10 petaflops FP16 or 20 PetaOps of int8 in a single VM *(using the new sparsity feature)*

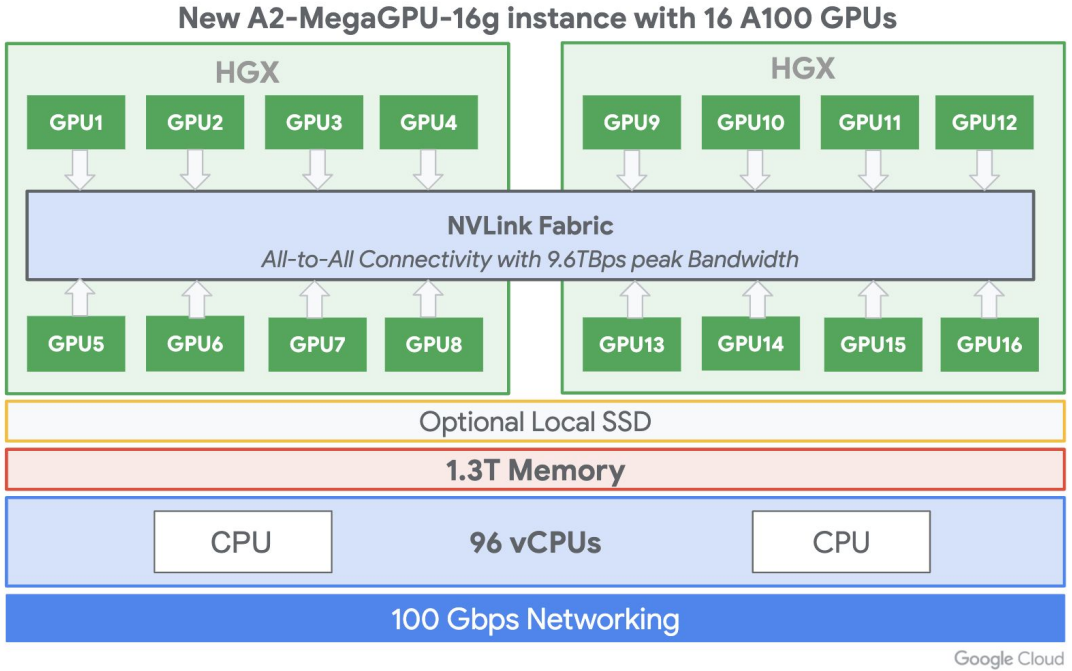

Google Cloud

Confidential Compute

- Available on N2D and upcoming C2D
- Just like a regular GCE Virtual Machine
  - Anything that runs on a VM runs on CVM
- Data encrypted while in-use
  - Memory encrypted, decrypted only on CPU chip
  - A key per Virtual Machine
    - Random, ephemeral, generated by HW
    - Not extractable from HW
- Leverages the AMD SEV powered by 2nd Gen AMD EPYC™ processors
- Scale up to 224 vCPUs and 896 GiB memory

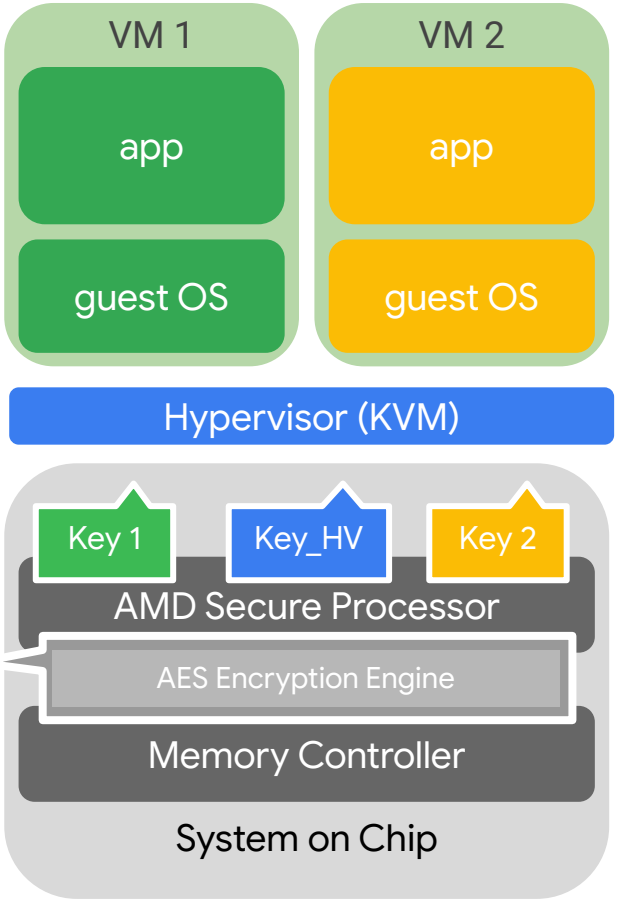

Google Cloud

Harvard Medical School conducts ultra-large virtual drug screenings on Google Cloud to speed up SARS-CoV-2 research

Google Cloud and Harvard Medical School won the 2020 HPCwire Reader's Choice Award for "Best Use of HPC in Cloud"!

- Developed an open source, scalable **virtual drug screening platform** called VirtualFlow that runs on Google Cloud
- Tested over **one billion compounds in five days**, using **80k vCPUs** for **75 million CPU hours**
- Team **targeted 16 proteins**, covering **40 target sites**, and **performed multiple screenings** to improve accuracy
- **Released the 1,000 most promising compounds** for researchers to explore and to begin lab testing

> Bigger is better in virtual screening. We've taken this to the next magnitude with 1.4 billion compounds now. Hopefully in the near future we'll go to about twenty billion compounds. This will revolutionize drug discovery.

Haribabu Arthanari
Assistant Professor, Department of Biological Chemistry and Molecular Pharmacology and creator of Virtualflow

HARVARD GLOBAL HEALTH INSTITUTE

https://edu.google.com/why-google/case-studies/harvard-covid19-gc
https://www.hpcwire.com/20-hpcwire-awards-readers-editors-choice/
https://vf4covid19.hms.harvard.edu/

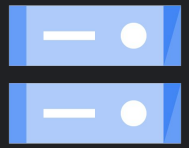

Storage

Google Cloud

Which storage type?

In Memory	Relational		NoSQL		Analytical	Object	Block		File
						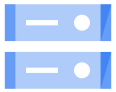			
Memorystore	Cloud SQL	Cloud Spanner	Firestore	Cloud Bigtable	BigQuery	Cloud Storage	Persistent Disk	Local SSD	Filestore & Partners
Managed Redis & Memcached	Managed MySQL and PostgreSQL, and SQL Server	Scalable relational database	Serverless, scalable, document store	Low-latency, scalable key-value and wide-column store	Enterprise DW	Unstructured data, objects or blobs	Flexible VM block storage	High performance NVMe-based block storage	HPC workloads requiring file and POSIX

Network

Google Cloud

Global network infrastructure

The largest cloud network: 100,000s of miles of fiber optic cable, 8 subsea cables
More edge and peering points than any public cloud

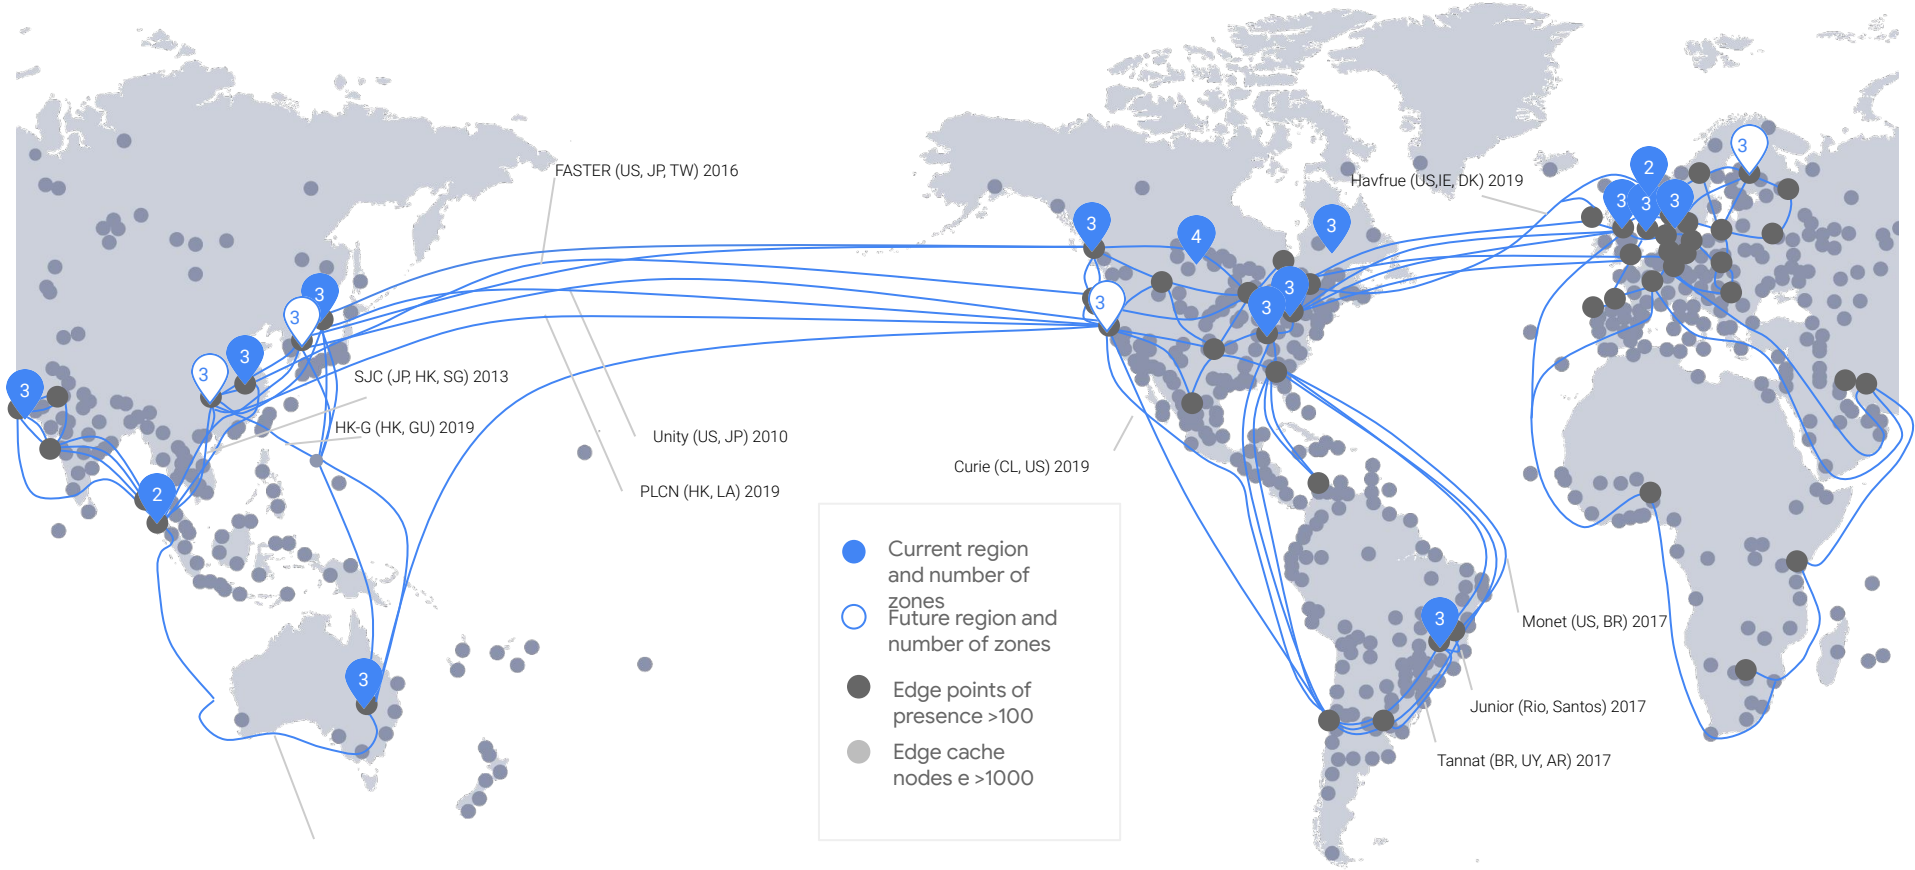

Global network infrastructure

The largest cloud network: 100,000s of miles of fiber optic cable, 8 subsea cables
More edge and peering points than any public cloud

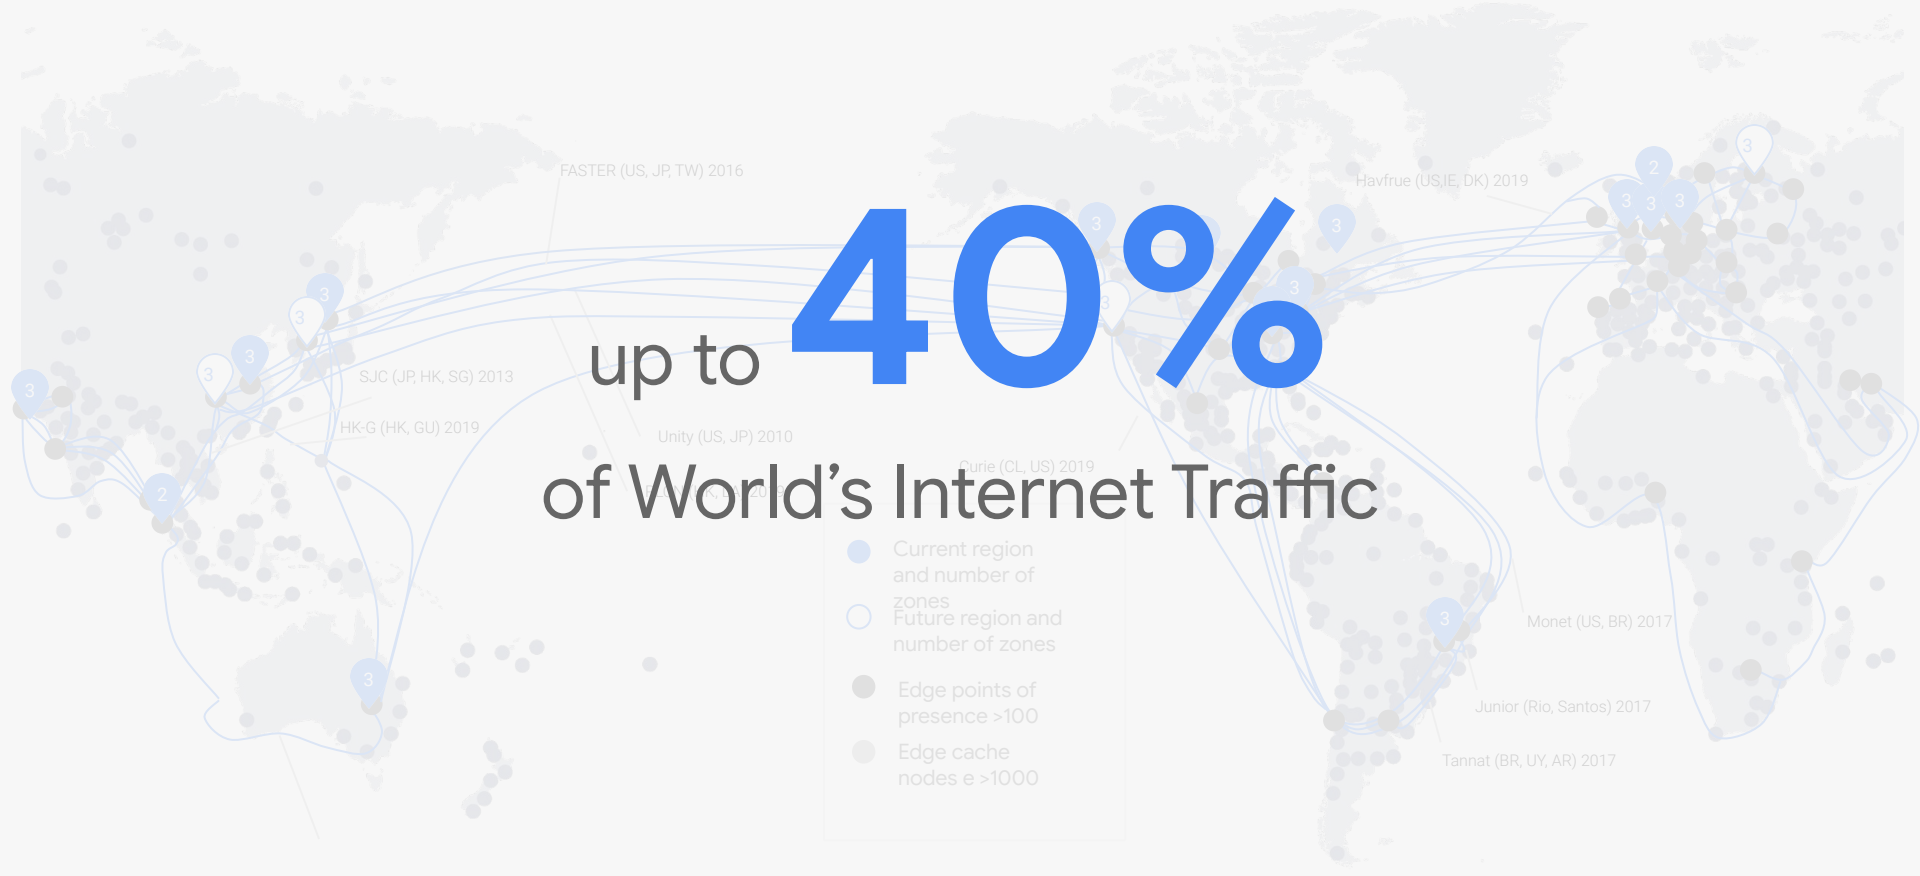

up to **40%**

of World's Internet Traffic

The network matters.

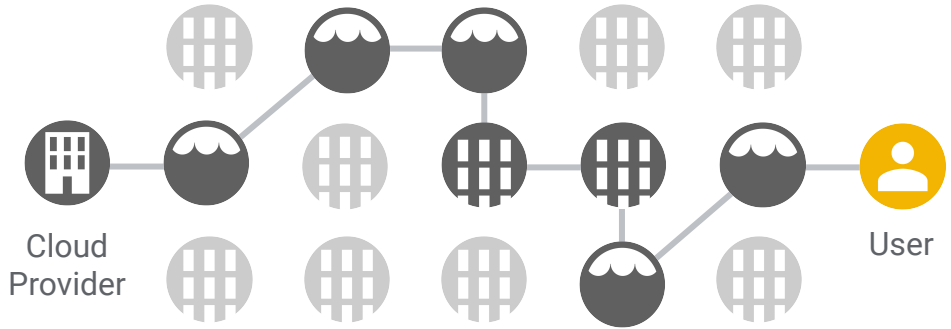

Typical cloud provider

Google Cloud

Google Cloud

HPC Networking on Google Cloud

Scalable, high-bandwidth, low-latency VM networking

- **Scalable Bandwidth**
  - 2 Gbps per vCPU
  - Up to **32 Gbps** per VM by default
  - Up to **100 Gbps** with **Tier 1 Bandwidth**
- **Low Latency**
  - Predictable, low latency (~10 µs average)
- **Tuning & Optimization**
  - **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 MTU** (Jumbo Frames) in preview
- 15,000 VMs per Virtual Private Cloud Network

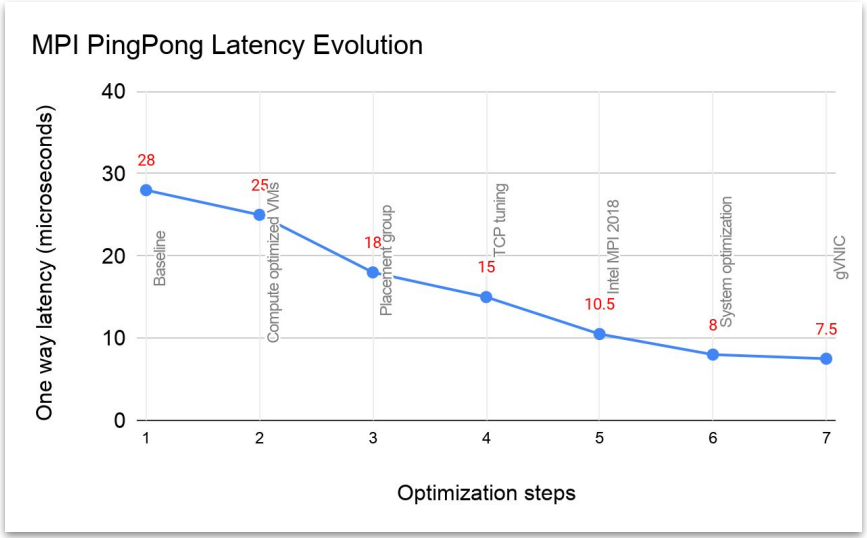

Example MPI Scalability

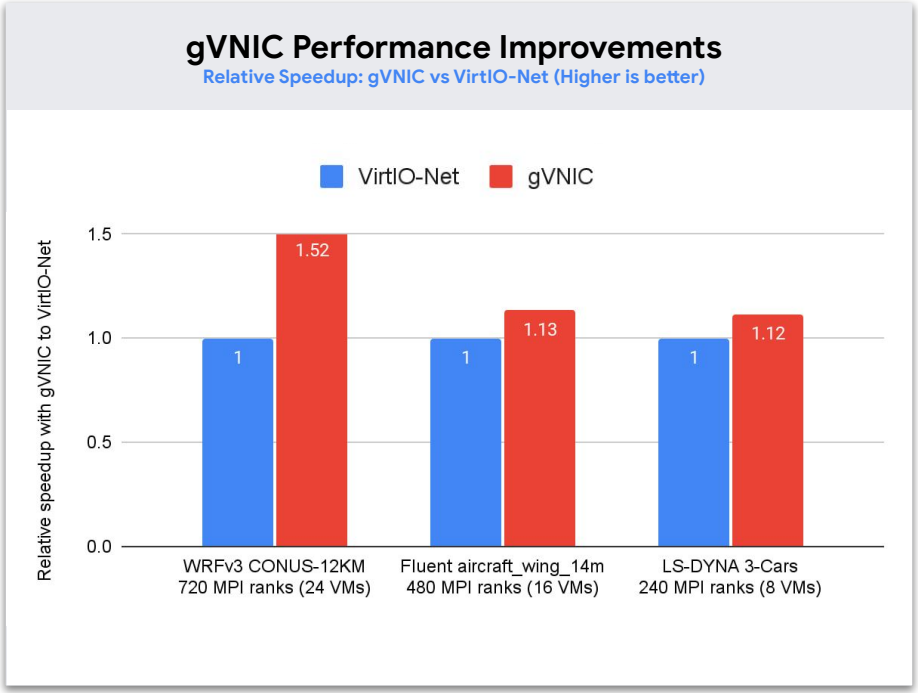

https://cloud.google.com/blog/topics/hpc/running-mpi-workloads-efficiently-on-google-cloud-using-gvnic

Take advantage of configurable, short-lived instances, and discounts to save on cost

Custom Machine Types

Choose your exact CPU/RAM ratio to match your workload

Committed use discounts

(up to 57% for most resource types) on longer term contracts

Spot VMs

Up to 91% cheaper, no bidding, no set time limit, all zones & regions

Google's Unique HPC Capabilities

Hybrid, Standards-based HPC Environment
Cloud-capable, hybrid support for all major HPC schedulers and HPC platforms. Intel HPC standards-based, open source supporters.

High Performance Cloud Platform
Fastest VM creation in cloud. HPC-tuned VM types. Bulk API. 100Gbps, Placement Policies, and gVNIC for network-intensive workloads. High performance file, block, and object storage.

Unique Accelerator Types
Market-leading selection of NVIDIA GPUs, and Google's unique Cloud TPUs for purpose-built Machine Learning acceleration.

Scale and Resilience
Software-defined global network (lowest latency CSP), Cloud Storage with 11 nines of durability, VMs with seamless Live Migration.

Secure by Default
End-to-end encryption by default. Fully private environment built around Google's Zero Trust approach. Confidential Compute VMs.

Thank you.

https://cloud.google.com/hpc

Google Cloud