Solution Brief

Cloud Services
DevOps Automation

Maximize Your Cloud Investment with Intel and HashiCorp

HashiCorp Terraform and Intel® Cloud Optimization Modules for Terraform and Sentinel automate VM and platform provisioning, maximizing performance and delivering the best cost-performance on private and public cloud services on Intel architecture.

HashiCorp

Streamlining DevOps

Manually provisioning cloud resources detracts product and application teams from their prime focus of innovating and solving business challenges. The burden is on developers to research technology options and then choose, optimize, and deploy instances for their workloads. From the choices of numerous configurations to sizes and how to optimize an instance for a particular workload require specific knowledge about hardware and optimization requirements. It all takes time.

As organizations continue to move more workloads to the public cloud, developers turn to repeatable templates and code to provision the infrastructure needed to expedite development activities. This shift helps developers work faster; however, it also escalates cloud costs, security risks, and workstream inefficiencies due to suboptimal infrastructure choices. Policies are an important part to standardizing infrastructure automation for developers who are provisioning infrastructure to run their applications.

Consistency. Performance. Optimization. On-Demand.

HashiCorp's Terraform is an infrastructure as code software tool that enables teams to safely and predictably create, change and improve infrastructure providing end-to-end provisioning and governance platform for DevOps. Developers can quickly deploy the resources they need for their workloads while adhering to IT operators' policies. The Terraform platform streamlines provisioning, improves performance, and lowers costs of VM instances and platforms for on-premises or public cloud operations. Terraform's policy engine, Sentinel, helps automatically ensure continuous governance according to IT policies.

Intel has developed new Intel Cloud Optimization Modules for Terraform and Sentinel to seamlessly add optimizations that leverage built-in Intel® Xeon® Scalable processor features and accelerators to maximize performance and deliver the best cost-performance on Intel for your critical workload. Intel provides these optimizations in the form of Intel Tuning Guides requiring manual configuration and adoption. Together Intel and HashiCorp have automated adoption of these optimizations with Terraform and policy enforcement with Sentinel. For example, for MySQL, Intel optimizations for Terraform include the necessary options to enhance performance of InnoDB on Intel processors. Intel also provides policy libraries for the Sentinel policy engine. Terraform and Intel Cloud Optimization Modules enable application and product teams to focus on solving business problems with the best performance instead of solving provisioning challenges.

Terraform Pipeline Automates Provisioning

Terraform provides the workspaces from which developers automatically provision instances and platforms for their needs. It communicates with the developers version control system (VCS) of choice, creating a deployment plan that is checked by the Sentinel policy engine for compliance (Figure 1).

The developer can then automatically apply provisioning to the cloud environment of choice.

As application code is changed or released in the VCS, Terraform automatically builds a new plan and checks it against appropriate policies, allowing the developer to redeploy as needed.

Figure 1. DevOps IaC/PaC pipeline with HashiCorp

Intel Cloud Optimization Modules Build Better Instances

Terraform includes many Intel Cloud Optimization Modules and policy libraries that can be quickly integrated into developer code with simple snippets copied from the Terraform platform site (Figure 2). Terraform automatically integrates these optimizations into the provisioning plan.

Intel policies for Sentinel make it easy to configure provisioning to meet the developer's needs, such as defaulting to the latest Intel architecture or a specific Intel processor or instance type, while being compliant with IT policies. Intel Cloud Optimization Modules are also available directly through GitHub.

Intel continues to expand the number of optimization modules available. Modules coming in 2023 and beyond include the following:

- GKE, CloudSQL for Google Cloud
- Databricks for Azure

With Terraform and Intel Cloud Optimization Modules and policies, provisioning happens seamlessly, automatically, and optimized for the workload.

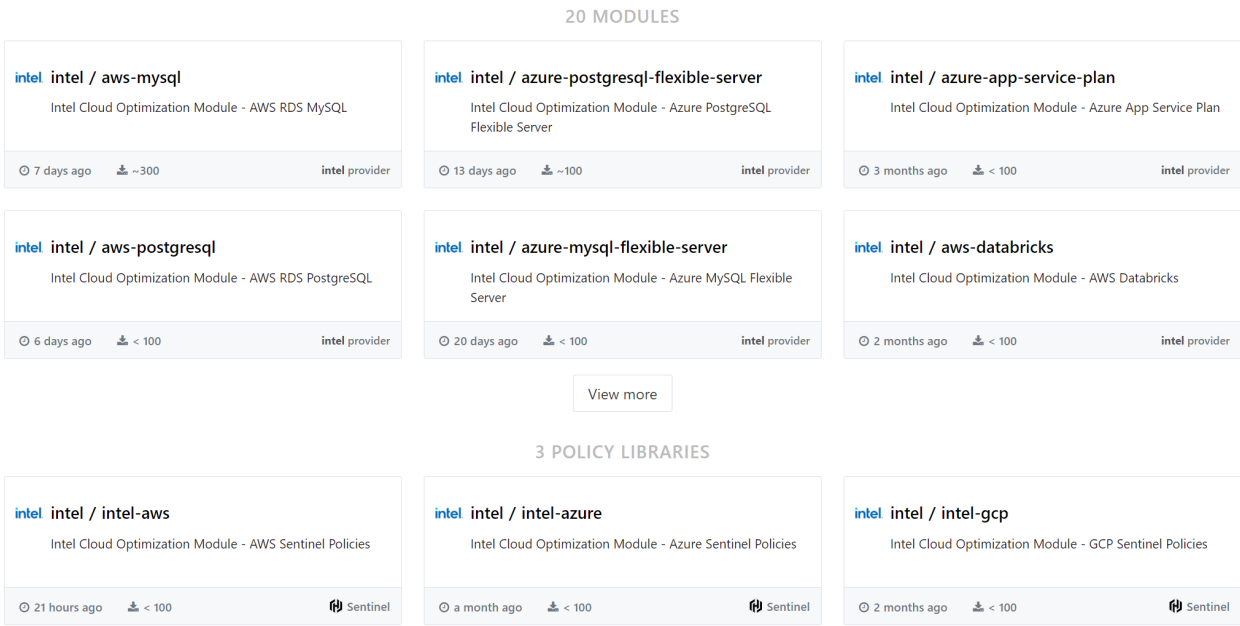

Figure 2. Currently there are over 20 Intel Cloud Optimization Modules and Policies for Terraform at https://registry.terraform.io/namespaces/intel

Optimizations Deliver Performance and Help Reduce Costs

Optimizing for Intel architecture delivers the most performance and often best cost-performance for workloads. For example, as shown in Figure 3, instances with 3rd Generation Intel Xeon processors on Azure can deliver up to 35 percent more performance for MySQL.

Normalized General-Purpose MySQL Performance

Normalized NOPM | Higher is better

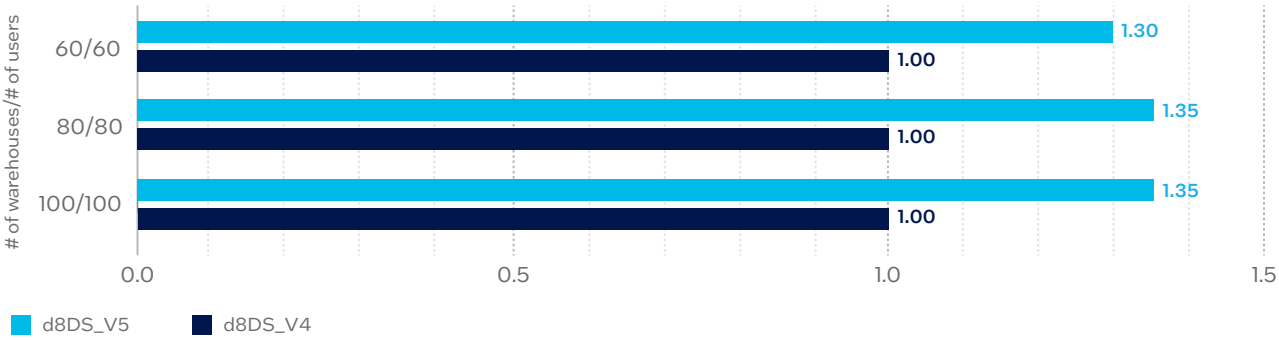

Figure 3. Handle up to 35% More MySQL New Orders per Minute on Microsoft Azure VMs Featuring 3rd Gen Intel Xeon Processors[1]

Automatically Keep Up With Technology Advancements

4th Gen Intel Xeon processors and Intel Xeon CPU Max Series are designed to accelerate a wide range of workloads with new accelerator engines. As Cloud Service Providers begin offering instances based on these processors, developers can expect new levels of performance for their applications. Figure 3 charts performance of common cloud applications of Google Cloud C3 instances with 4th Gen Intel Xeon processors compared to Google c2 compute-optimized instances with 2nd Gen Intel Xeon processors.

The C3 instances offer:

- More than double the performance for NGINX
- Around 40% performance improvement for MySQL, WordPress, and media transcoding

Intel Cloud Optimization Modules for Terraform and Sentinel policies can inform and automate for developers new technologies as they are released in the cloud, so applications can take advantage of the latest emerging technologies and highest performance—automatically and seamlessly.

Google C3 (4th Gen Intel Xeon Processors) vs. Google C2 (2nd Gen Intel Xeon Processors)[1]

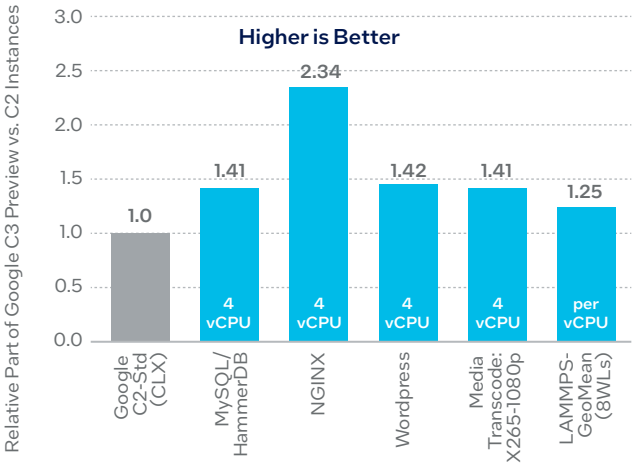

Google Cloud new C3 Machine Series[2]

Observing a **20% increase** in performance of C3 over C2 in testing with one of our key workloads.

As much as **10x quicker** for about the same computational cost for weather research and forecasting on C3 clusters.

Seeing up to **3x performance** gains over C2 VMs due to higher memory bandwidth and lower network latency.

Figure 4. Cloud instance performance gains with 4th Generation Intel Xeon Processors on Google Cloud[2]

Conclusion

Terraform streamlines cloud resource provisioning for product and application teams so they can focus on business problems instead of deployment challenges. Optimizing these instances with Intel Cloud Optimization Modules for Terraform and Sentinel policy libraries allows developers to easily specify and optimize instances that are deployed automatically to the cloud platform of their choice. Terraform Sentinel enforces Operators' policies to help IT departments effectively manage costs and efficiencies in their operations. Together, development teams are able to seamlessly optimize their operations and application performance while reducing cost.

For more information about HashiCorp and Terraform, visit **hashicorp.com** and **terraform.io**

To learn more go to **https://registry.terraform.io/namespaces/intel**

Download Intel Cloud Optimization Modules on **GitHub**

1. https://www.intel.com/content/www/us/en/partner/workload/microsoft/mysql-on-azure-8vcpu-benchmark.html
2. The ratio for LAMMPS assumes that the LAMMPS workloads scale linearly with the number of cores for both the platforms compared.

* Results measured by Intel using preview version of Google C3-highcpu, including pre-production processors, systems, and software, vs Google C2-standard.

** Source: Google Cloud blog post October 11, 2022, cloud.google.com/blog/products/compute/introducing-c3-machines-with-googles-custom-intel-ipu.

MySQL / HammerDB:
Google c2-standard-4 instance based on 2nd Generation Intel Xeon Scalable processor (Cascade Lake), Ubuntu 20.04.4 LTS, Kernel 5.15.0-1017-gcp, HammerDB-v4.3, MySQL 8.0.30, VU=10, Warehouses=100, Measured by Intel September 22, 2022

Google c3-highcpu-4 private preview instance based on 4th Generation Intel Xeon Scalable processor (Sapphire Rapids), Ubuntu 20.04.4 LTS, Kernel 5.15.0-1017-gcp, HammerDB-v4.3, MySQL 8.0.30, VU=10, Warehouses=100, Measured by Intel September 22, 2022 using preview version of Google C3-highcpu, including pre-production processors, systems, and software.

NGINX:
Google c2-standard-4 instance based on 2nd Generation Intel Xeon Scalable processor (Cascade Lake), Ubuntu 20.04.5 LTS, Kernel 5.15.0-1017-gcp, NGINX 1.18.0 QAT cipher AES128-GCM-SHA256, NGINX_QAT_clients=2000, Measured by Intel September 16, 2022

Google c3-highcpu-4 private preview instance based on 4th Generation Intel Xeon Scalable processor (Sapphire Rapids), Ubuntu 20.04.5 LTS, Kernel 5.15.0-1017-gcp, NGINX 1.18.0 QAT cipher AES128-GCM-SHA256, NGINX_QAT_clients=2000, Measured by Intel September 21, 2022 using preview version of Google C3-highcpu, including pre-production processors, systems, and software.

Wordpress:
Google c2-standard-4 instance based on 2nd Generation Intel Xeon Scalable processor (Cascade Lake), WordPress Single-Tier GCP: TLSv1.3, WPv5.6.5, PHPv8.0, 10 server threads, 300 client threads, Measured by Intel September 22, 2022

Google c3-highcpu-4 private preview instance based on 4th Generation Intel Xeon Scalable processor (Sapphire Rapids), WordPress Single-Tier GCP: TLSv1.3, WPv5.6.5, PHPv8.0, 10 server threads, 450 client threads, Measured by Intel September 22, 2022 using preview version of Google C3-highcpu, including pre-production processors, systems, and software.

Media Transcode X.265-1080P

Google c2-standard-4 instance based on 2nd Generation Intel Xeon Scalable processor (Cascade Lake), Ubuntu 20.04.5 LTS, Kernel 5.15.0-1017-gcp, FFMPEG v2.3.1-2-g45492669, X265 Release_3.5, Measured by Intel September 27, 2022

Google c3-highcpu-4 private preview instance based on 4th Generation Intel Xeon Scalable processor (Sapphire Rapids), Ubuntu 20.04.5 LTS, Kernel 5.15.0-1017-gcp, FFMPEG v2.3.2-beta, X265 Release_3.5, Measured by Intel September 23, 2022 using preview version of Google C3-highcpu, including pre-production processors, systems, and software.

LAMMPS-Geomean (8WLs):

Google c2-standard-60 instance based on 2nd Generation Intel Xeon Scalable processor (Cascade Lake), Centos7, 3.10.0-1160.62.1.el7.x86_64, version 29 Oct 2020, Intel C Compiler icc (ICC) 2021.5.0, Intel(R) MPI Library, Version 2018 Update 4, Measured by Intel August 26, 2022

Google c3-highcpu-176 private preview instance based on 4th Generation Intel Xeon Scalable processor (Sapphire Rapids), Centos7, 3.10.0-1160.76.1.el7.x86_64, version 29 Oct 2020, Intel C Compiler icc (ICC) 2021.5.0, Intel(R) MPI Library, Version 2018 Update 4, Measured by Intel September 13, 2022 using preview version of Google C3-highcpu, including pre-production processors, systems, and software.

0623/SMR/HBD/PDF ♻PleaseRecycle 356070-001US