# The SPEC ACCEL Benchmark - Results and Lessons Learned

**Robert Henschel**

Chair, High Performance Group (HPG)
Standard Performance Evaluation Corporation (SPEC)

Treasurer, OpenACC.org

Director, Research Software and Solutions
Indiana University / Pervasive Technology Institute

SC19, WACCPD Workshop, November 2019

RESEARCH TECHNOLOGIES
INDIANA UNIVERSITY
University Information Technology Services

PERVASIVE TECHNOLOGY INSTITUTE
INDIANA UNIVERSITY

spec

# Thank you for your time today!

This work is support by NSF award 1842623 (Robert Henschel, IU, PI)

Jetstream is supported by NSF award 1445604 (David Hancock, IU, PI)

This research was supported in part by the Indiana University Pervasive Technology Institute (PTI), which was established with the assistance of a major award from the Lilly Endowment, Inc.

Opinions presented here are those of the author(s) and do not necessarily represent the views of the NSF, PTI, IU, or the Lilly Endowment, Inc.

# Content

- SPEC and SPEC HPG
  - SPEC Benchmark Philosophy
  - SPEC HPG Benchmarks
- Deep Dive: SPEC ACCEL
- Next Generation Benchmark

# SPEC and SPEC HPG


**SPEC** is a non-profit corporation formed in 1988 to establish, maintain and endorse standardized benchmarks and tools to evaluate performance and energy efficiency for the newest generation of computing systems.

- OSG: Open System Group
- **HPG: High Performance Group**
- GWPG: Graphics & Workstation Performance Group
- RG: Research Group

HPC benchmarks
- MPI
- OpenMP
- Accelerator
  - OpenCL 1.1
  - OpenACC 1.0
  - OpenMP 4.5

# SPEC Media Coverage 2017

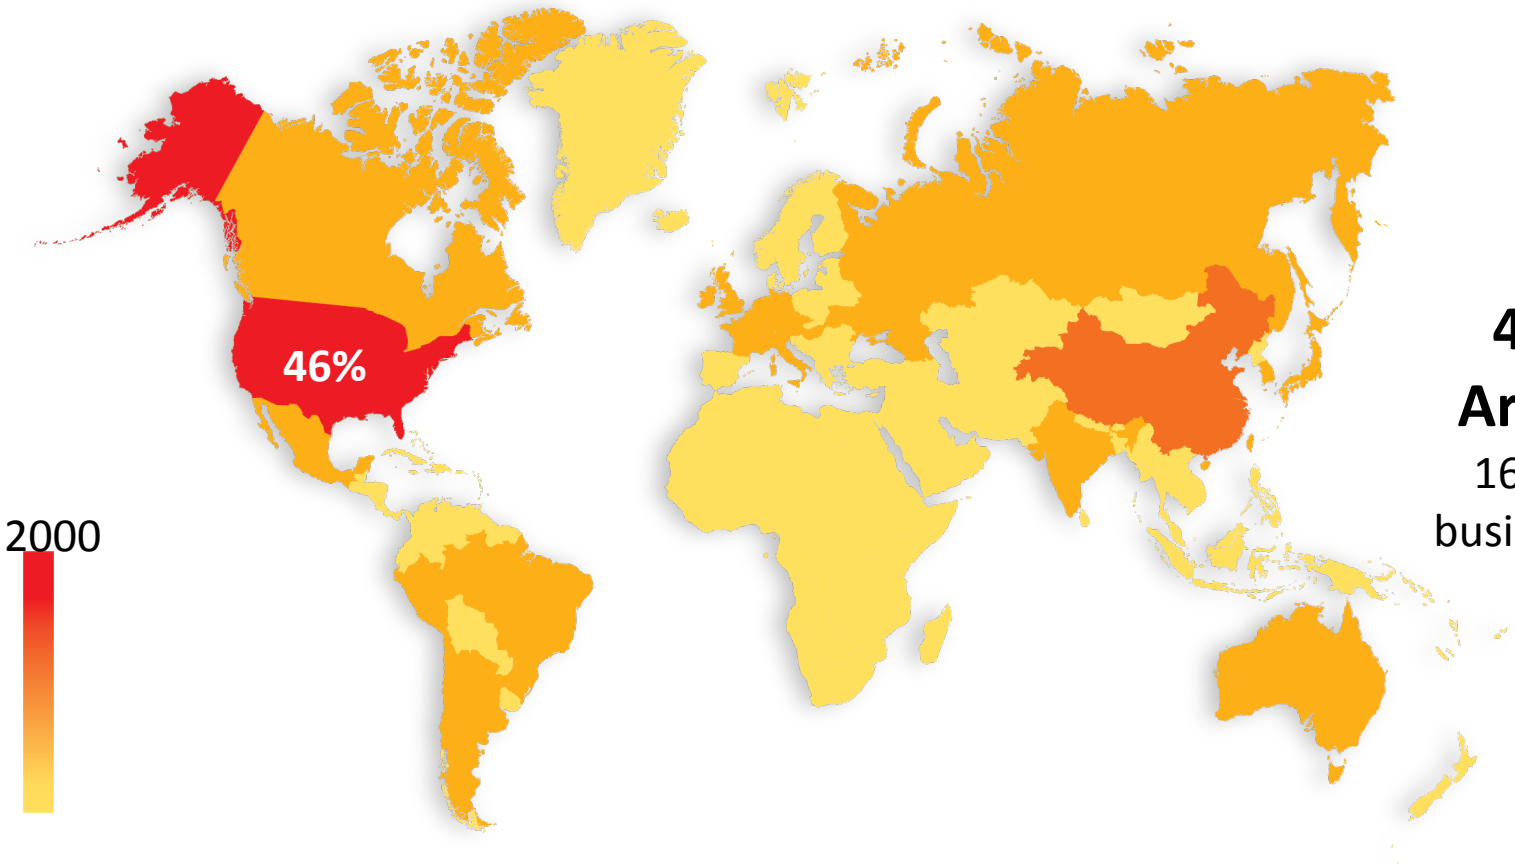


**4362 Articles**

16 every business day

# SPEC HPG


HPG develops benchmarks to represent high-performance computing applications for standardized, cross-platform performance evaluation.

33 Organizations
9 companies
24 academic




























# Content

- SPEC and SPEC HPG
  - SPEC Benchmark Philosophy
  - SPEC HPG Benchmarks
- Deep Dive: SPEC ACCEL
- Next Generation HPC Benchmark

# SPEC Benchmark Philosophy

- SPEC supports the full lifecycle of benchmark development!

- The result of a SPEC benchmark is one SPEC score.
  - Higher is better
  - Some benchmarks support power measurement
- This score is in relation to a reference machine.
  - Each benchmark has its own reference machine
- SPEC (HPG) benchmarks are “full” applications.
  - Including all the overhead of a real application
- SPEC harness ensures correctness of results.
  - To detect “overly aggressive optimization” and tampering
- Each benchmark suite has run rules and documentation requirements.

# SPEC Benchmark Philosophy

Hierarchy within benchmark suits

- Benchmark suite
- Benchmark
- Dataset size
- Component

SPEC ACCEL
↳ OpenACC/OpenMP/OpenCL
↳ Ref/Test
↳ 550.md

Benchmarks support “Base” and “Peak” configuration

- These yield separate SPEC scores, “Peak” runs allow for more freedom.

Base runs

- The same compiler optimization switches for all components of a language
- The same level of parallelism
- Only portability switches allowed

# SPEC Benchmark Philosophy

Result submission:

- Obtain and install the benchmark
- Perform a valid run and describe hardware and software configuration
- Submit result for review (and publication) to SPEC HPG – 2 week review process
- If needed, define embargo period
- Results are published on SPEC website

A curated result repository:

- Given appropriate hardware and software…. a published result should be reproducible with the information available in the submission.
- Peer reviewed results are so much better than “everyone can upload a result”!
- The value of a benchmark suite lies in public results, their correctness and the ability to compare them.

# SPEC Benchmark Philosophy

Demo: https://spec.org/accel/results/accel.html

Demo: https://spec.org/accel/results/res2014q1/accel-20140303-00018.html

# SPEC Benchmark Philosophy

**OpenACC (31):**

| Test Sponsor | System Name | Accelerator Name | Results | | Energy | |
|---|---|---|---|---|---|---|
| | | | Base | Peak | Base | Peak |
| Cirrascale Corporation | GIGABYTE MD70-HB0 Motherboard HTML \| CSV \| Text \| PDF \| PS \| Config | FirePro s9150 | 2.89 | 2.99 | -- | -- |
| Cirrascale Corporation | GIGABYTE MD70-HB0 Motherboard HTML \| CSV \| Text \| PDF \| PS \| Config | FirePro s9150 | 3.10 | 3.21 | -- | -- |
| Cirrascale Corporation | GIGABYTE MD70-HB0 Motherboard HTML \| CSV \| Text \| PDF \| PS \| Config | FirePro s9150 | 3.60 | Not Run | -- | -- |
| Indiana University | Cray XK7 HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K20 | 1.74 | Not Run | -- | -- |
| Indiana University | Cray XK7 HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K20 | 1.27 | Not Run | -- | -- |
| Indiana University | Cray XK7 HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K20 | 1.31 | Not Run | -- | -- |
| Indiana University | Cray XK7 HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K20 | 1.77 | Not Run | -- | -- |
| NVIDIA Corporation | ASUS P9X79 Motherboard HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K40c | 2.59 | 2.73 | 3.01 | 3.13 |
| NVIDIA Corporation | ASUS P9X79 Motherboard HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K40c | 2.59 | 2.72 | 3.35 | 3.49 |
| RWTH Aachen University | bullx R421-E3 HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Tesla K20Xm | 2.00 | Not Run | -- | -- |
| RWTH Aachen University | bullx R425-E2 HTML \| CSV \| Text \| PDF \| PS \| Config | NVIDIA Quadro 6000 | 1.05 | Not Run | -- | -- |

# SPEC® ACCEL™ OMP Result



Colfax International (Test Sponsor: Indiana University)
Xeon Phi 7210
Ninja Developer Platform Pedestal: Liquid Cooled

| | |
|---|---|
| SPECaccel_omp_base = | 3.40 |
| SPECaccel_omp_energy_base = | 4.54 |
| SPECaccel_omp_peak = | Not Run |
| SPECaccel_omp_energy_peak = | -- |

| | | | |
|---|---|---|---|
| **ACCEL license:** | 3440A | **Test date:** | May-2017 |
| **Test sponsor:** | Indiana University | **Hardware Availability:** | Aug-2016 |
| **Tested by:** | Indiana University | **Software Availability:** | Jan-2017 |

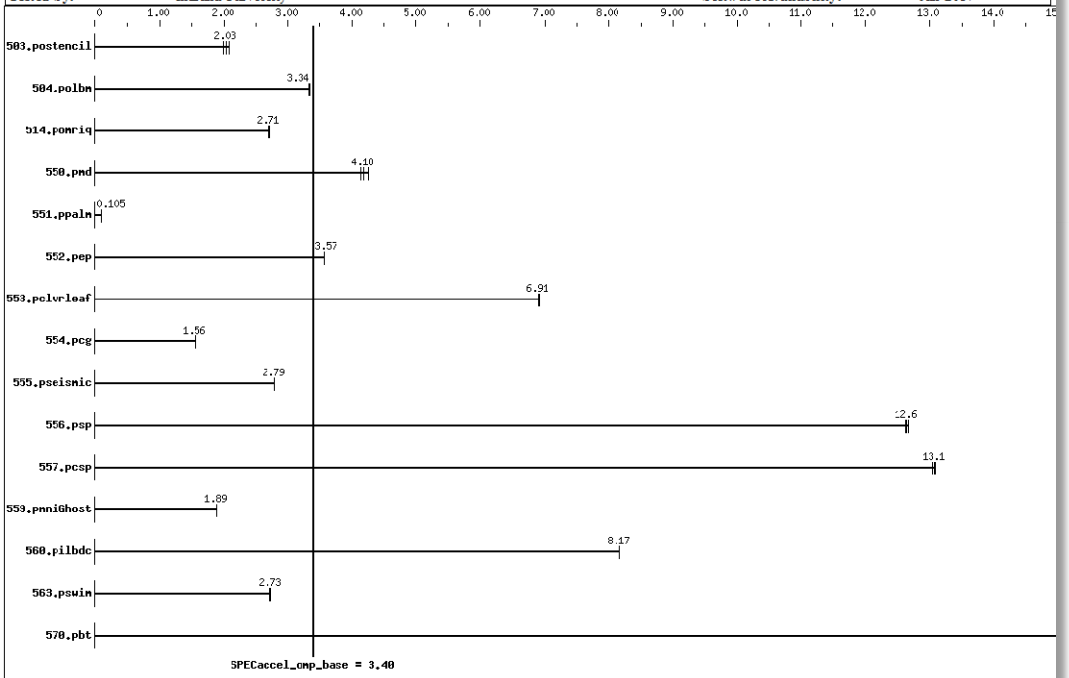


## Hardware

| | |
|---|---|
| **CPU Name:** | Intel Xeon Phi 7210 |
| **CPU Characteristics:** | Simultaneous multithreading (SMT) on, Turbo off. |
| **CPU MHz:** | 1300 |
| **CPU MHz Maximum:** | 1300 |
| **FPU:** | Integrated |
| **CPU(s) enabled:** | 64 cores, 1 chip, 64 cores/chip, 4 threads/core |
| **CPU(s) orderable:** | 1 to 1 chip |
| **Primary Cache:** | 32 KB I + 32 KB D on chip per core |
| **Secondary Cache:** | 1 MB I+D on chip per tile (2 cores) |
| **L3 Cache:** | None |
| **Other Cache:** | None |
| **Memory:** | 96 GB (6 x 16 GB 2Rx8 PC4-2400T-REB-11, ECC) + 16 GB MCDRAM |
| **Disk Subsystem:** | Intel S3510 SSD 800GB, SATA3 |
| **Other Hardware:** | None |

## Accelerator

| | |
|---|---|
| **Accel Model Name:** | Xeon Phi 7210 |
| **Accel Vendor:** | Intel |
| **Accel Name:** | Xeon Phi 7210 |
| **Type of Accel:** | CPU |
| **Accel Connection:** | N/A |
| **Does Accel Use ECC:** | Yes |
| **Accel Description:** | Second generation Xeon Phi self-bootable CPU, SMT on, Turbo off, flat DDR4+MCDRAM |
| **Accel Driver:** | N/A |

## Software

| | |
|---|---|
| **Operating System:** | CentOS Linux release 7.2.1511 (Core) 3.10.0-327.13.1.el7.xppsl_1.3.3.151.x86_64 |
| **Compiler:** | Intel Parallel Studio XE 2017 Update 1 for Linux, Version 17.0.1.132 Build 20161005 |
| **File System:** | ext4 |
| **System State:** | Run level 3 (multi-user with networking) |
| **Other Software:** | None |

## Power

| | |
|---|---|
| **Power Supply:** | 750W |

spec

Colfax International (Test Spons
Xeon Phi 7210
Ninja Developer Platform Pe

| | |
|---|---|
| **ACCEL license:** | 3440A |
| **Test sponsor:** | Indiana Univers |
| **Tested by:** | Indiana Univers |

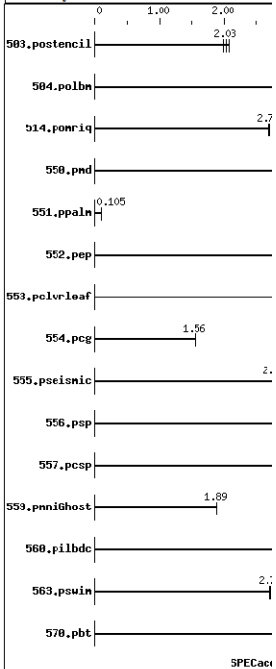


## Hard

| | |
|---|---|
| **CPU Name:** | Intel Xeon Phi 7 |
| **CPU Characteristics:** | Simultaneous mu |
| **CPU MHz:** | 1300 |
| **CPU MHz Maximum:** | 1300 |
| **FPU:** | Integrated |
| **CPU(s) enabled:** | 64 cores, 1 chip, |
| **CPU(s) orderable:** | 1 to 1 chip |
| **Primary Cache:** | 32 KB I + 32 KE |
| **Secondary Cache:** | 1 MB I+D on ch |
| **L3 Cache:** | None |
| **Other Cache:** | None |
| **Memory:** | 96 GB (6 x 16 G + 16 GB MCDR |
| **Disk Subsystem:** | Intel S3510 SSD |
| **Other Hardware:** | None |

**Power Supply:**

| | |
|---|---|
| **Power Supply Details:** | Seasonic SSR-750RM Active PFC F3 |
| **Max. Power (W):** | 286.39 |
| **Idle Power (W):** | 91.01 |
| **Min. Temperature (C):** | 21.69 |

## Power Analyzer

| | |
|---|---|
| **Power Analyzer:** | 156.56.179.146:8888 |
| **Hardware Vendor:** | ZES Zimmer |
| **Model:** | ZES LMG450:4-Channel |
| **Serial Number:** | 01001849 |
| **Input Connection:** | RS232 USB adapter |
| **Metrology Institute:** | NIST (National Institute of Standards and Technology) |
| **Calibration By:** | ZES Zimmer |
| **Calibration Label:** | 3783190001e |
| **Calibration Date:** | 02.20.2017 |
| **PTDaemon Version:** | 1.8.1 (a497ea15; 2016-12-20) |
| **Setup Description:** | connected to the single power supply that powers the system |
| **Current Ranges Used:** | 0.0A |
| **Voltage Range Used:** | 130V |

## Temperature Meter

| | |
|---|---|
| **Temperature Meter:** | 156.56.179.146:8889 |
| **Hardware Vendor:** | Digi |
| **Model:** | Watchport/H |
| **Serial Number:** | W40236768 |
| **Input Connection:** | USB |
| **PTDaemon Version:** | 1.8.1 (a497ea15; 2016-12-20) |
| **Setup Description:** | positioned in front of intake fan |

## Base Results Table

| Benchmark | Seconds | Ratio | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | | | | | | | | | | | | | | | | | | | |

spec

Colfax International (Test Spons
Xeon Phi 7210
Ninja Developer Platform Pe

| | |
|---|---|
| **ACCEL license:** | 3440A |
| **Test sponsor:** | Indiana Univers |
| **Tested by:** | Indiana Univers |

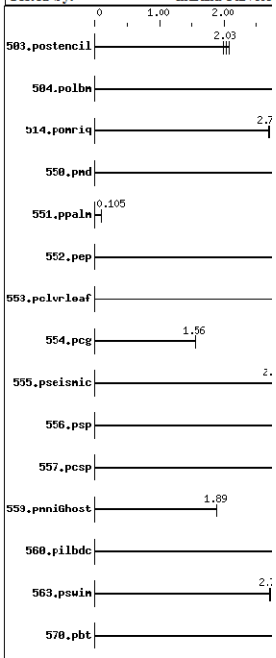


## Hard

| | |
|---|---|
| **CPU Name:** | Intel Xeon Phi 7 |
| **CPU Characteristics:** | Simultaneous mu |
| **CPU MHz:** | 1300 |
| **CPU MHz Maximum:** | 1300 |
| **FPU:** | Integrated |
| **CPU(s) enabled:** | 64 cores, 1 chip, |
| **CPU(s) orderable:** | 1 to 1 chip |
| **Primary Cache:** | 32 KB I + 32 KE |
| **Secondary Cache:** | 1 MB I+D on ch |
| **L3 Cache:** | None |
| **Other Cache:** | None |
| **Memory:** | 96 GB (6 x 16 G + 16 GB MCDR |
| **Disk Subsystem:** | Intel S3510 SSD |
| **Other Hardware:** | None |

**Power Supply:**

| | |
|---|---|
| **Power Supply Details:** | Sea |
| **Max. Power (W):** | 286 |
| **Idle Power (W):** | 91.( |
| **Min. Temperature (C):** | 21.( |

## Power Analyzer

| | |
|---|---|
| **Power Analyzer:** | 156.56.179.146:8888 |
| **Hardware Vendor:** | ZES Zimmer |
| **Model:** | ZES LMG450:4-Channel |
| **Serial Number:** | 01001849 |
| **Input Connection:** | RS232 USB adapter |
| **Metrology Institute:** | NIST (National Institute of Standards and Technology) |
| **Calibration By:** | ZES Zimmer |
| **Calibration Label:** | 3783190001e |
| **Calibration Date:** | 02.20.2017 |
| **PTDaemon Version:** | 1.8.1 (a497ea15; 2016-12-20) |
| **Setup Description:** | connected to the single power supply that po the system |
| **Current Ranges Used:** | 0.0A |
| **Voltage Range Used:** | 130V |

| Benchmark | Seconds | Ratio | | | | | |
|---|---|---|---|---|---|---|---|
| | | | | | | | |

```
      VERSION="7 (Core)"
      ID="centos"
      ID_LIKE="rhel fedora"
      VERSION_ID="7"
      PRETTY_NAME="CentOS Linux 7 (Core)"
      ANSI_COLOR="0;31"
      CPE_NAME="cpe:/o:centos:centos:7"
   redhat-release: CentOS Linux release 7.2.1511 (Core)
   system-release: CentOS Linux release 7.2.1511 (Core)
   system-release-cpe: cpe:/o:centos:centos:7

uname -a:
   Linux knl1.uits.indiana.edu 3.10.0-327.13.1.el7.xppsl_1.3.3.151.x86_64 #1 SMP
   Fri Jun 10 15:04:35 UTC 2016 x86_64 x86_64 x86_64 GNU/Linux

run-level 3 May 2 10:51

SPEC is set to: /home/lijunj/spec/accel-test/75
   Filesystem     Type  Size  Used Avail Use% Mounted on
   /dev/sda3      ext4  713G  174G  503G  26% /

Cannot run dmidecode; consider saying 'chmod +s /usr/sbin/dmidecode'

(End of data from sysinfo program)
```

## General Notes

```
BIOS settings:
  Intel Simultaneous Multithreading (SMT): on
  Intel Turbo Boost Technology (Turbo)  : off
  Cluster Mode: quadrant
  Memory Mode: flat
  (MCDRAM is partitioned to the second NUMA node)

Current range for power measurement is 2.5A.
```

## Base Compiler Invocation

**C benchmarks:**

icc

**Fortran benchmarks:**

ifort

**Benchmarks using both Fortran and C:**

icc ifort

## Base Portability Flags

503.postencil: -DSPEC_USE_INNER_SIMD
504.polbm: -DSPEC_USE_INNER_SIMD
514.pomriq: -DSPEC_USE_INNER_SIMD
550.pmd: -DSPEC_USE_INNER_SIMD -80
551.ppalm: -DSPEC_USE_INNER_SIMD
552.pep: -DSPEC_USE_INNER_SIMD
553.pclvrleaf: -DSPEC_USE_INNER_SIMD
554.pcg: -DSPEC_USE_INNER_SIMD
555.pseismic: -DSPEC_USE_INNER_SIMD
556.psp: -DSPEC_USE_INNER_SIMD
557.pcsp: -DSPEC_USE_INNER_SIMD
559.pmniGhost: -DSPEC_USE_INNER_SIMD -nofor-main
560.pilbdc: -DSPEC_USE_INNER_SIMD
563.pswim: -DSPEC_USE_INNER_SIMD
570.pbt: -DSPEC_USE_INNER_SIMD

## Base Optimization Flags

**C benchmarks:**

-O3 -qopenmp -qopenmp-offload=host -xMIC-AVX512

**Fortran benchmarks:**

-O3 -qopenmp -qopenmp-offload=host -xMIC-AVX512

**Benchmarks using both Fortran and C:**

-O3 -qopenmp -qopenmp-offload=host -xMIC-AVX512

The flags files that were used to format this result can be browsed at
https://www.spec.org/accel/flags/Intel-ic17.0-linux64.html,
https://www.spec.org/accel/flags/colfax-knl.html.

You can also download the XML flags sources by saving the following links:
https://www.spec.org/accel/flags/Intel-ic17.0-linux64.xml,
https://www.spec.org/accel/flags/colfax-knl.xml.

# Content

- SPEC and SPEC HPG
  - SPEC Benchmark Philosophy
  - SPEC HPG Benchmarks
- Deep Dive: SPEC ACCEL
- Next Generation HPC Benchmark

# SPEC HPG Benchmarks - Pricing

- Different groups in SPEC have different policies on the sale of benchmarks.
- **Since March 2018, SPEC HPG benchmarks are available free of charge to non-profit organizations, including universities and research labs.**
- SPEC HPG hopes that this will encourage even more organizations to actively participate.

# Published Results

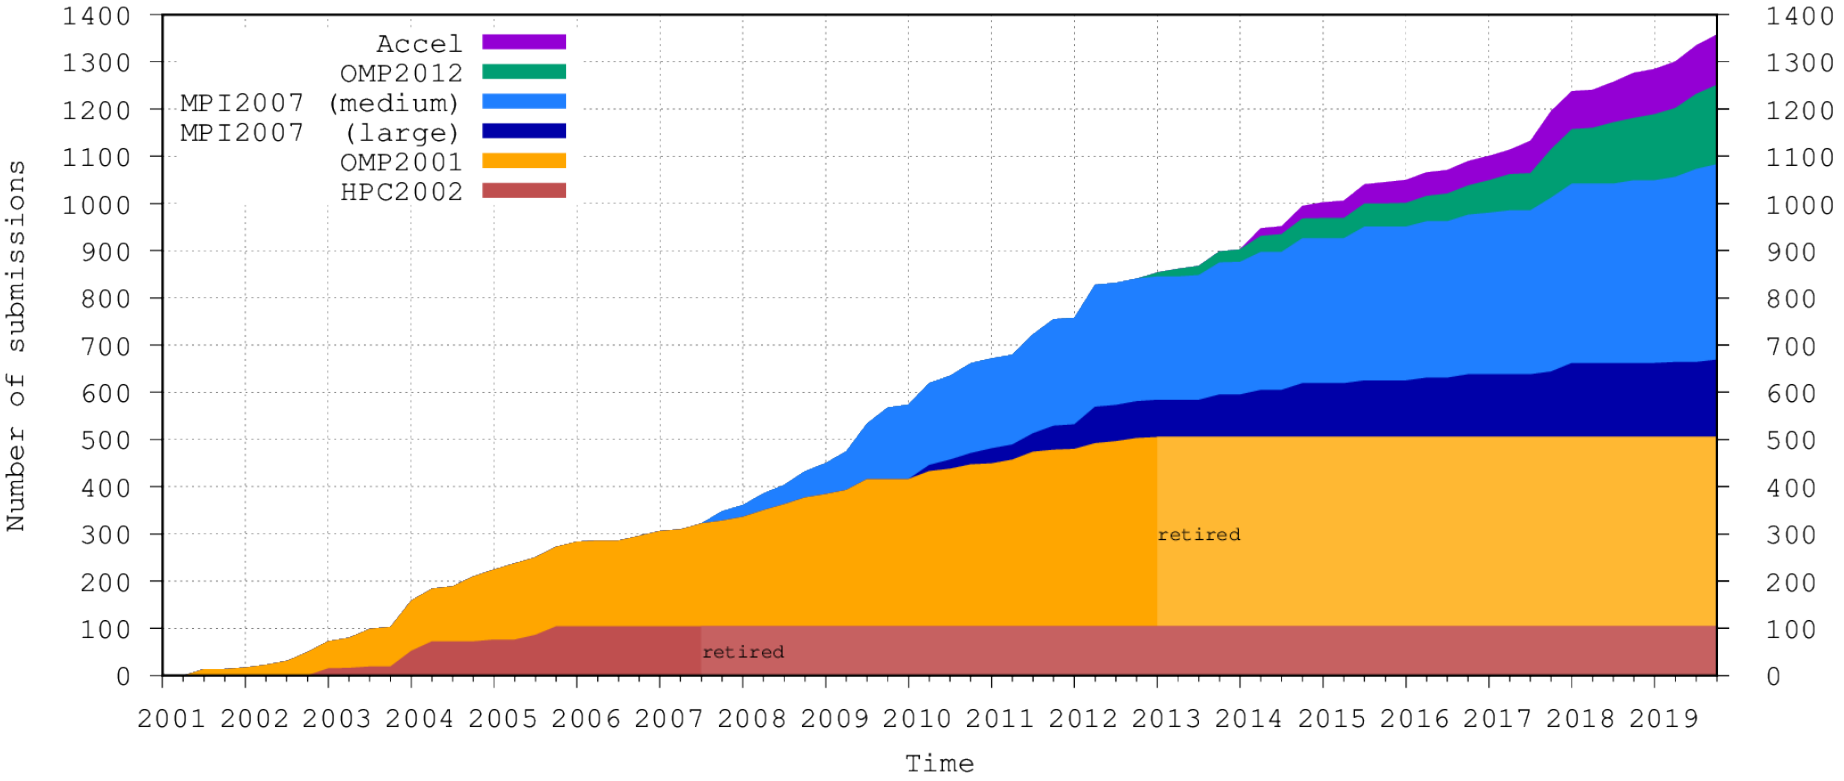

# SPEC HPG Benchmarks – OMP2012

- Follow on to SPEC OMP2001
- 14 applications Fortran/C
- Scales up to 512 threads
- Support for power measurement

# SPEC HPG Benchmarks – MPI2007

- Large and medium data set
- 13 applications in Fortran/C/C++
- Scales to 2048 MPI processes
- Power not supported

# Content

- SPEC and SPEC HPG
  - SPEC Benchmark Philosophy
  - SPEC HPG Benchmarks
- Deep Dive: SPEC ACCEL
- Next Generation HPC Benchmark

# SPEC HPG Benchmarks - ACCEL

- SPEC Accel provides a comparative performance measure of
  - Hardware accelerator devices (GPU, Co-processors, etc.)
  - Supporting software tool chains (Compilers, Drivers, etc.)
  - Host systems and accelerator interface (CPU, PCIe, etc.)
- Computationally-intensive parallel HPC applications and mini-apps
- Portable across accelerator platforms
- Three distinct benchmarks, initially released in 2014, updated in 2017 and 2019:
  - OpenCL 1.1 — 19 C/C++ applications
  - OpenACC 1.0 — 15 Fortran/C applications
  - OpenMP 4.5 — 15 Fortran/C applications
- Support for power measurement

# Timeline of SPEC ACCEL

- First release in 2014
  - Support for OpenCL 1.1 and OpenACC 1.0
- Second release in 2017
  - Support for OpenMP 4.5
- Minor updates in between
- Latest update in August 2019

# Creating the OpenACC Benchmark

- Sources of applications:
  - Parboil, University of Illinois
  - Rodinia, University of Virginia
  - NAS Parallel Benchmarks (NPB)
  - Other SPEC benchmarks
- Cray and PGI compiler (CPU and GPU)
- NVIDIA Tesla C2070 GPU as reference machine

# Mini Applications

| Benchmarks | Language | Origin | Domain |
|---|---|---|---|
| 503.ostencil | C | Parboil, University of Illinois | Thermodynamics |
| 504.olbm | C | Parboil, University of Illinois | CFDm Lattice Boltzmann |
| 514.omriq | C | Rodinia, University of Virginia | Medicine |
| 550.md | Fortran | Indiana University | Molecular Dyn. |
| 551.palm | Fortran | Leibniz University of Hannover | Large-eddy sim. |
| 552.ep | C | NAS Parallel Benchmarks (NPB) | Embarrassing P. |
| 553.clvrleaf | C, Fortran | Atomic Weapons Establishments | Hydrodynamics |
| 554.cg | C | NPB | Conjugate Grad. |
| 555.seismic | Fortran | GeoDynamics.org | Seismic Wave Modeling (PDE) |
| 556.sp | Fortran | NPB | Scalar Peta-d solv |
| 557.csp | C | NPB | Scalar Peta-d solv |
| 559.miniGhost | C, Fortran | Sandia National Laboratory | Finite difference |
| 560.ilbdc | Fortran | SPEC OMP2012 | Fluid Mechanics |
| 563.swim | Fortran | SPEC OMP2012 | Weather |
| 570.bt | C | NPB | BTS 3D PDE |

# Converting OpenACC to OpenMP 4.5

- We started with 15 OpenACC applications of SPEC ACCEL.
- The Intel Compiler for XEON/Phi was used as reference.
  - Reference machine is dual Intel SandyBridge E5-2650, 8C, 2Ghz, with an intel XEON Phi 5110.
- We ported to OpenMP 4.0, but then 4.5 came out.
- The group agreed on guidelines how to turn OpenACC code into OpenMP 4.5.
- The applications were ported twice, first by PathScale, then by ZIH/TU-Dresden and then a consensus was used.

# How to write OpenMP 4.5 Code

- Rely on compilers to generate implementation specific values for a given architecture:
  - # of teams
  - # thread_limit,
  - # of threads – in parallel regions
  - SIMD length
  - dist_schedule – in distribute
  - loop schedules – in parallel do
- Compiler implementers pick these values to enable performance portability and generate platform specific optimizations.

# Compiler Performance Over Time

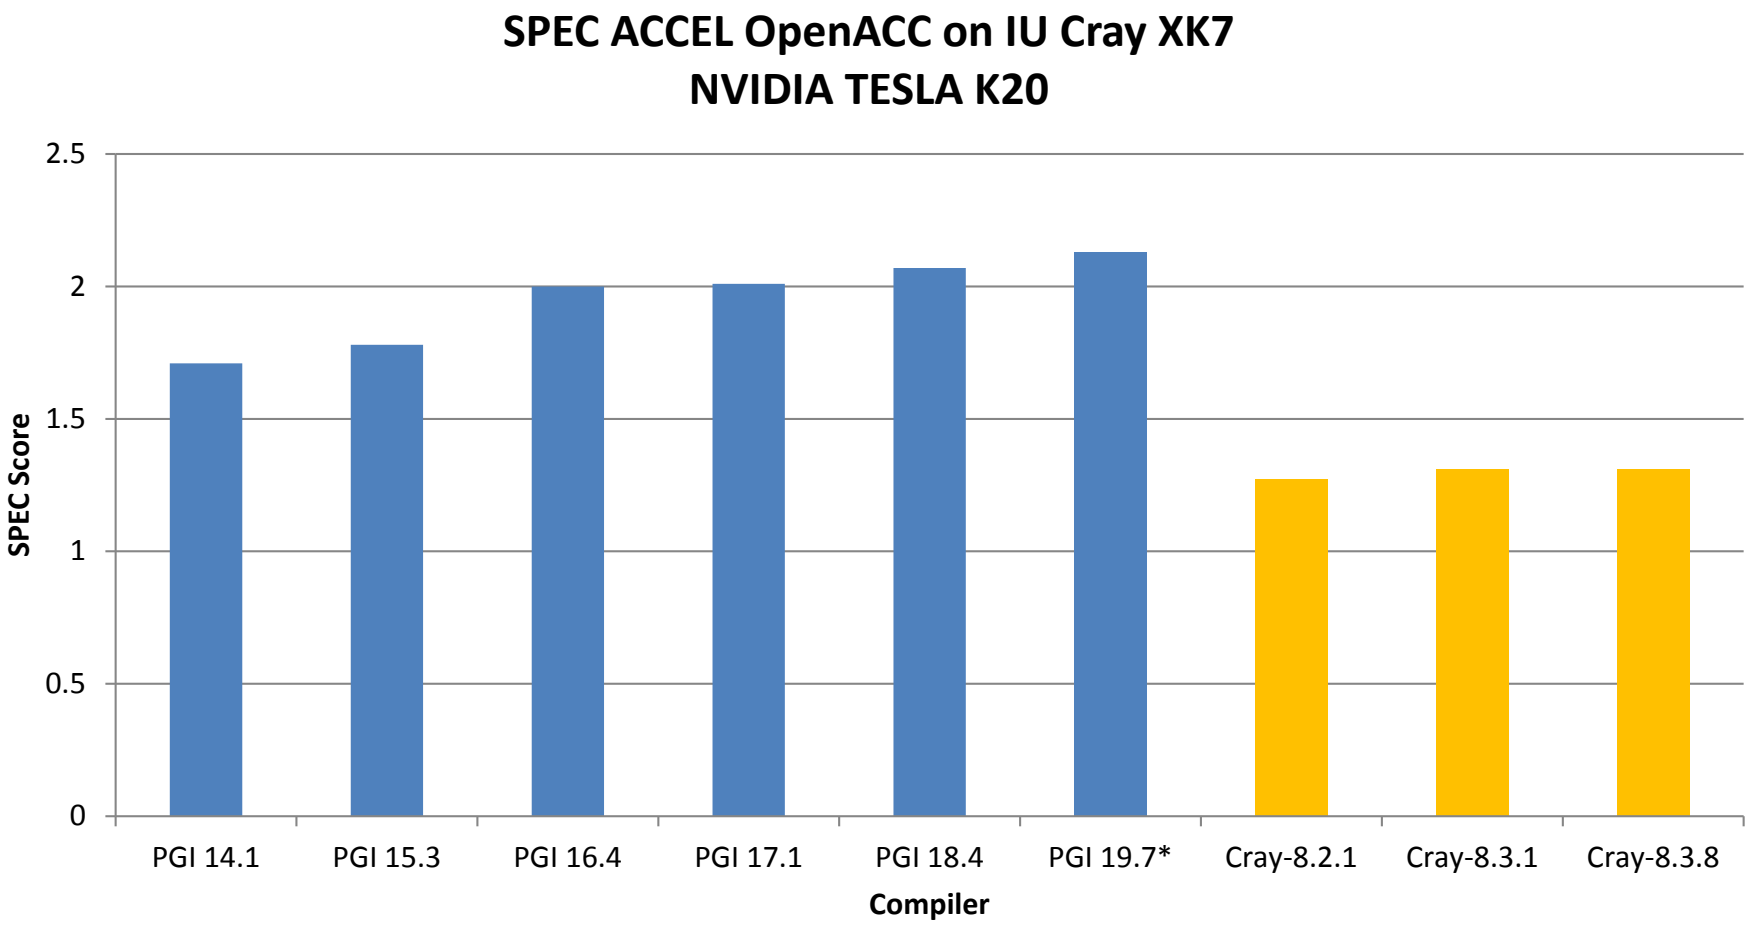


* Estimate (hardware was too old to publish on spec.org)

# SPEC ACCEL OpenACC on CPUs and GPUs

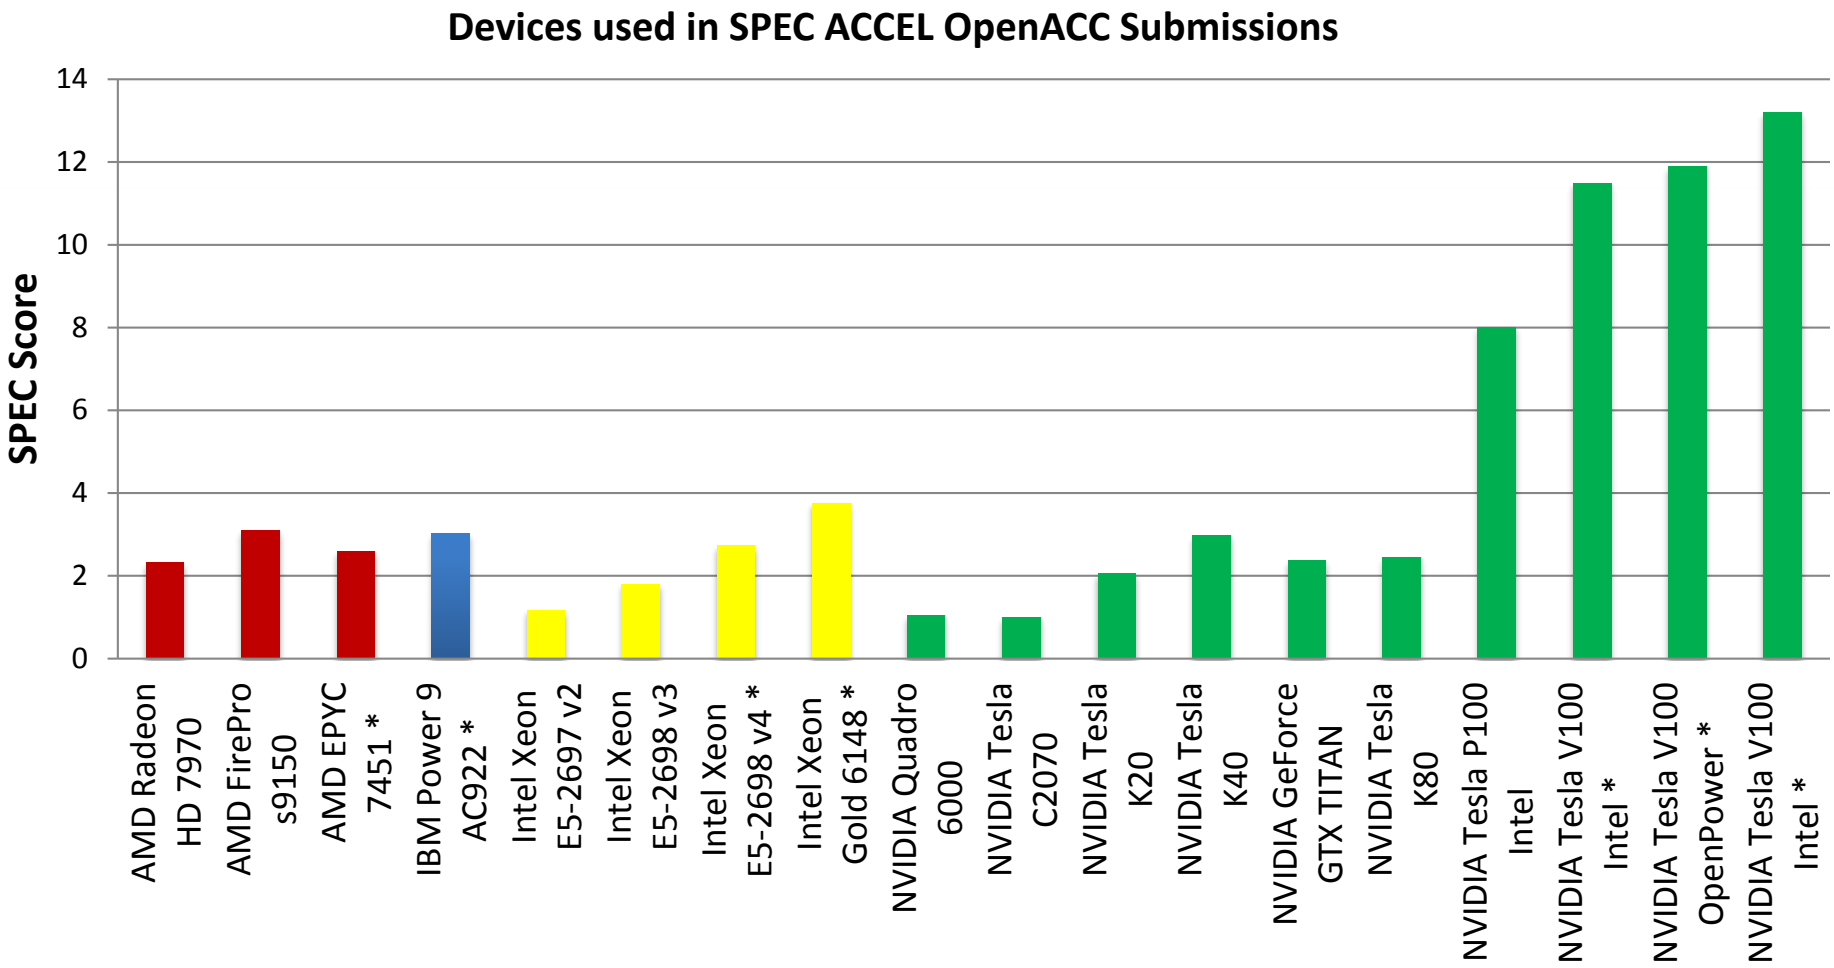




* Results from Version 1.2 of the SPEC ACCEL benchmark while all other results are from version 1.1.

# Experimental Results OpenMP Offload

- Cray and IBM compilers support OpenMP 4.5 offload to GPUs. We only had access to the Cray compiler and currently only 6 of 15 benchmarks work!
- RPeak: KNL-7210 2.60 TFlops
  K20 1.17 TFlops Ratio: 2.2x

| | SPEC Score (Estimate) | | | Speedup |
|---|---|---|---|---|
| Benchmark | KNL (MCDRAM) | KNL (DDR4) | K20 | KNL vs. K20 |
| 503.postencil | 1.99 | 0.70 | 1.26 | 1.6x |
| 504.polbm | 3.42 | 0.75 | 0.90 | 3.8x |
| 514.pomriq | 2.71 | 2.72 | 1.11 | 2.4x |
| 555.pseismic | 2.83 | 1.06 | 1.43 | 2.0x |
| 560.pilbdc | 8.43 | 1.97 | 4.61 | 1.8x |
| 570.pbt | 27.4 | 20.2 | 18.2 | 1.5x |
| Geometric Mean | | | | **2.1X** |

# SPEC ACCEL on Jetstream Virtual GPUs

- KVM with NVIDIA's Virtual Data Center Workstation Software (vDWS)
  - Based on the Linux kernel's Virtual Function I/O (VFIO)
  - Virtualized device functions are passed through by the hypervisor to guest VM kernel drivers
- GPUs can be "partitioned" using a fixed portion of the GPUs memory, but with access to all CUDA cores on a time division multiplexing basis.

# SPEC ACCEL on Jetstream Virtual GPUs

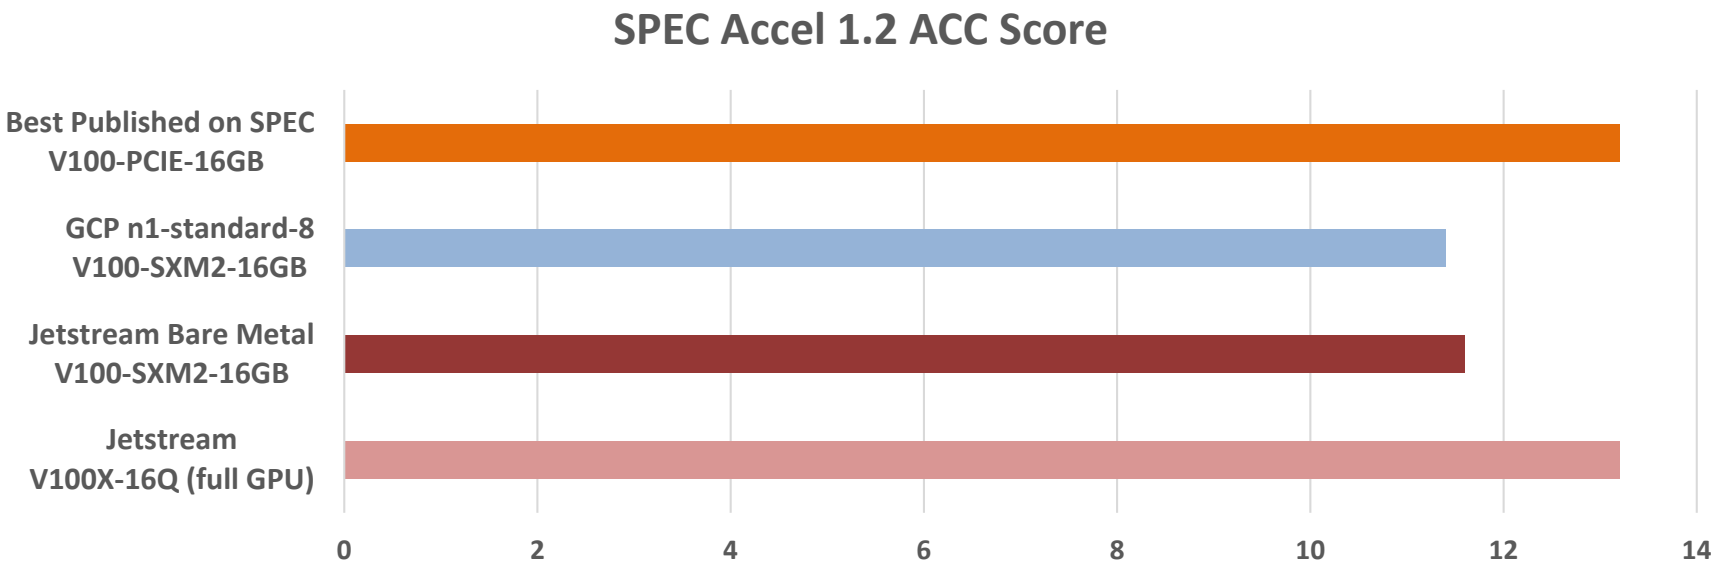


Can a VM be this much faster than bare metal?!

GCP VM: driver 418.67, cuda 10.1
Jetstream BM: driver 418.67, cuda 10.1
Jetstream VM: driver 418.70, cuda 10.1

# SPEC ACCEL on Jetstream Virtual GPUs

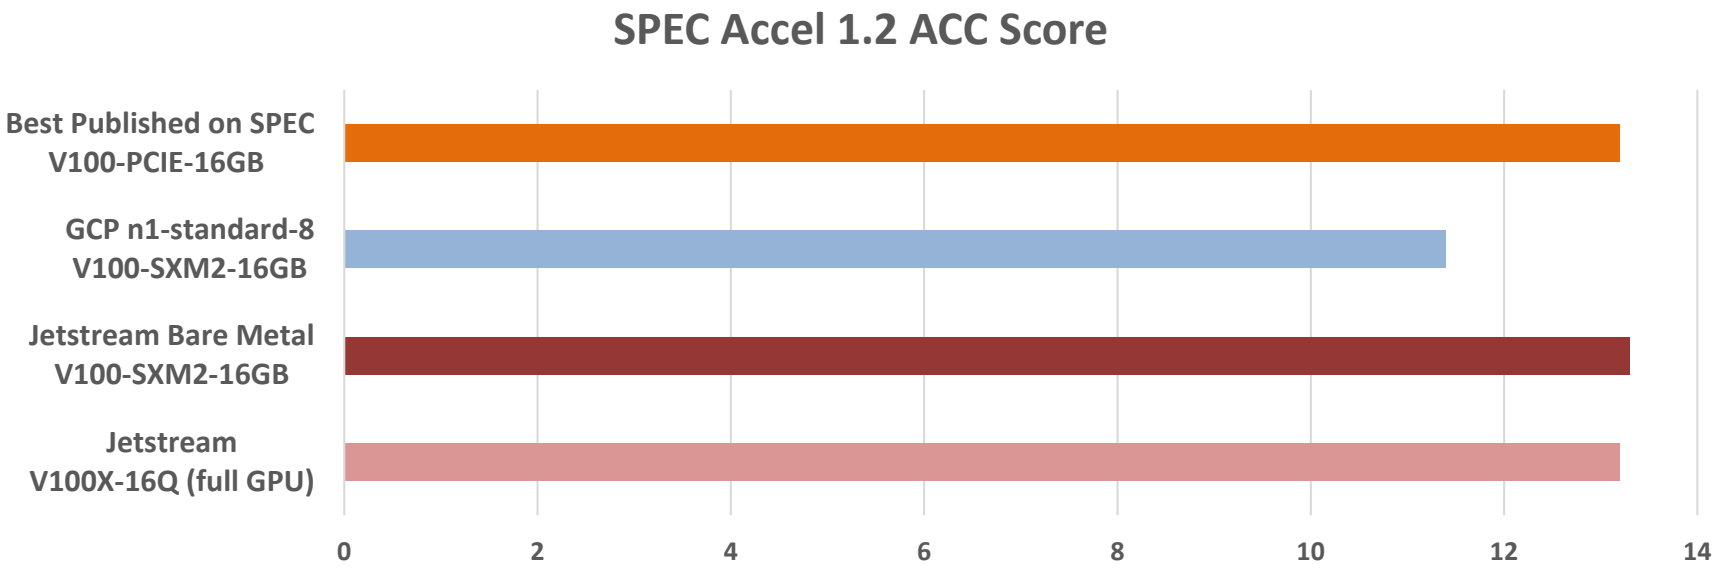


Finding on Jetstream bare metal:

When on driver 418.67, all 4x V100 need to have persistent mode (PM) on.

Turning PM on for each additional card, all cards get about 4% increase in speed.

# SPEC ACCEL on Jetstream Virtual GPUs

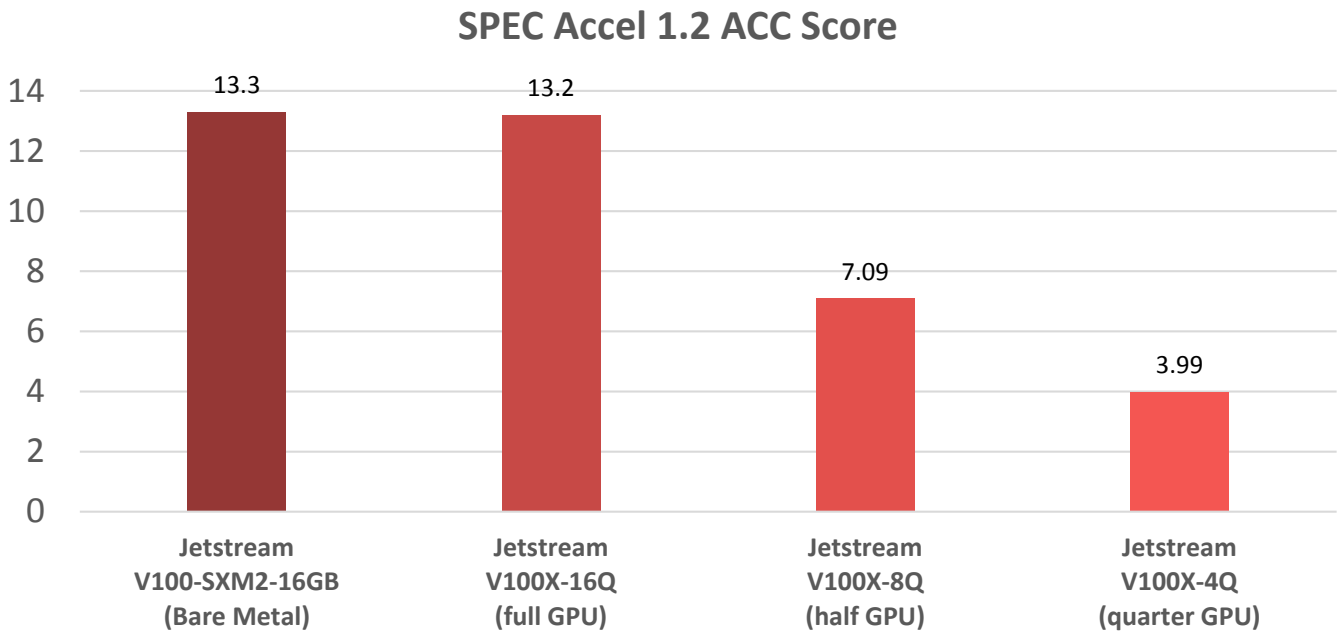


Results published on: https://spec.org/accel/results/res2019q4/

# Content

- SPEC and SPEC HPG
  - SPEC Benchmark Philosophy
  - SPEC HPG Benchmarks
- Deep Dive: SPEC ACCEL
- Next Generation HPC Benchmark

# Next Generation HPC Benchmark

- First hybrid benchmark for SPEC HPG, posing lots of challenges for run rules and metrics.
- A strong scaling benchmark, with 3 workload sizes.
  - Small (single node), Medium (2048 MPI Ranks), Large (8192 MPI Ranks)
  - Support at least 3 different node level parallel models (OpenACC, OpenMP 5, OpenMP "classic")
- Not directly replacing other HPG benchmarks, but lots of possible downstream uses.

# Conclusion

- SPEC High Performance Group has been around since 1994
  - Track record of creating and maintaining HPC benchmarks.
- Benchmarks are created and maintained by a broad mix of industry and academia.
  - Consider joining SPEC HPG to connect with compiler teams from AMD, IBM, Intel, PGI/NVIDIA.
- There is great value in peer reviewed results and a public searchable result repository.
- SPEC ACCEL has proven to be portable across many architectures and compilers.
- SPEC HPG is working on the next generation HPC benchmark, targeting hybrid parallelism.

# Thank You!

Questions?