

# Are We Witnessing the *Spectre* of an HPC Meltdown?

Verónica G. Vergara Larrea

Michael J. Brim

Wayne Joubert

Swen Boehm

Matthew Baker

Oscar Hernandez

Sarp Oral

James Simmons

Don Maxwell



Oak Ridge National Laboratory  
CUG 2018

ORNL is managed by UT-Battelle  
for the US Department of Energy

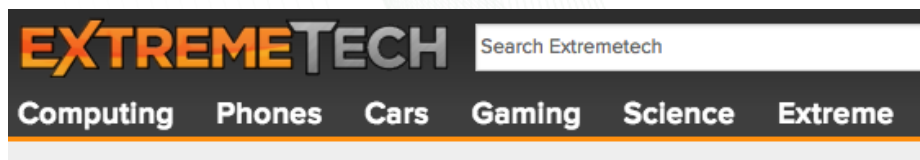


# Spectre and Meltdown: Background

The “Spectre” and “Meltdown” security vulnerabilities were discovered in 2017

These are flaws present in nearly all modern processors

There has been industry-wide concern that firmware patches required to fix these problems will significantly decrease processor performance



HOME > COMPUTING > RECENT INTEL CPUS TAKE PERFORMANCE HIT WITH SPECTRE, MELTDOWN PATCHES

## Recent Intel CPUs Take Performance Hit With Spectre, Meltdown Patches

By Joel Hruska on February 28, 2018 at 9:20 am | [43 Comments](#)

Forbes **CommunityVoice**™ Connecting expert communities to the Forbes audience. [What is this?](#)

MAR 21, 2018 @ 07:45 AM 769

[2 Free Issues of Forbes](#)

### 11 Potential Effects Of Meltdown And Spectre On The Tech Industry

# Are



HOME COMPUTE STORE CONNECT CONTROL CODE AI HPC

**Krebs**  
In-depth security news

## HOW SPECTRE AND MELTDOWN MITIGATION HITS XEON PERFORMANCE

March 16, 2018 Timothy Prickett Morgan

### 05 Scary Chip Flaws Raise Spectre of Meltdown

JAN 18

**Apple, Google, Microsoft** and other tech giants have released updates for a pair of serious security flaws present in most modern computers, smartphones, tablets and mobile devices. Here's a brief rundown on the threat and what you can do to protect your devices.

At issue are two different vulnerabilities, dubbed “Meltdown” and “Spectre,” that were

GE  
atory | LEADERSHIP  
COMPUTING  
FACILITY


## Why it matters to us

The Oak Ridge Leadership Computing Facility has commitments to programs such as INCITE, ALCC and ECP to deliver a quantity of compute cycles every year

Loss of system performance threatens our users' ability to deliver the promised science output for their projects

Also results in loss on our investments in HPC system purchase and maintenance

### Four Primary Ways for Access to LCF Distribution of allocable hours



## Overview of Spectre and Meltdown vulnerabilities

Flaws that enable potential side channel attacks by an adversary to access protected memory:

1. Spectre V1: use condition misprediction for speculative load into L3 cache, use this to indirectly infer a value in protected memory
2. Spectre V2: create branch misprediction condition that results in protected memory being loaded into L3 cache
3. Spectre V3 / Meltdown: thread executes illegal instruction, continues to execute speculatively, potentially loading data from protected memory

## Methodology of study

- Evaluate several of the top systems at the OLCF: Titan, Eos, Percival, Cumulus
- Use application codes from our system acceptance testing efforts, representing OLCF workloads, augmented with other benchmark codes
- Run codes before patches and after patches, compare performance
- Run each code at several node counts on each system and with 10 trials to measure runtime variability

# OLCF systems tested



## Titan

Cray XK7  
18,688 compute nodes:  
1 16-core AMD Interlagos CPU,  
1 NVIDIA K20X GPU  
Gemini interconnect



## Eos

Cray XC30  
736 compute nodes  
2 8-core Intel Xeon E5-2670 CPUs  
Aries interconnect



## Cumulus

Cray XC40  
112 compute nodes  
2 18-core Intel Broadwell CPUs  
Aries interconnect



## Percival

Cray XC40  
168 compute nodes  
1 64-core Intel KNL CPU  
Aries interconnect

⌘ Are We Witnessing the *Spectre* of an HPC Meltdown?

## Codes used

- **HPL** – *dense linear solve, compute intensive*
  - **HPCG** – *sparse linear algebra, multigrid, halo exchange, global reductions*
  - **OSU Microbenchmarks (OMB)** – *MPI tests*
  - **SPEC OMP2012** – *OpenMP application benchmarks*
- **GTC** – *fusion code, particles, sparse Poisson field solve*
  - **S3D** – *explicit combustion code, structured 3D grid*
  - **LAMMPS** – *molecular dynamics, particles*
  - **NWCHEM** – *chemistry code, network intensive*
  - **LSMS** – *electronic structure code, dense linear algebra, Monte Carlo*
  - **MILC** – *QCD code; optimized KNL version used*
- **IOR** – *I/O benchmark, POSIX and MPI-IO tests*
  - **mdtest** – *I/O metadata operations test*
  - **simul** – *performance of POSIX I/O operations*

## Results

We received Spectre, Meltdown patches from Cray for these systems over the past several months

These were not installed in isolation but were part of larger patch sets delivered by Cray

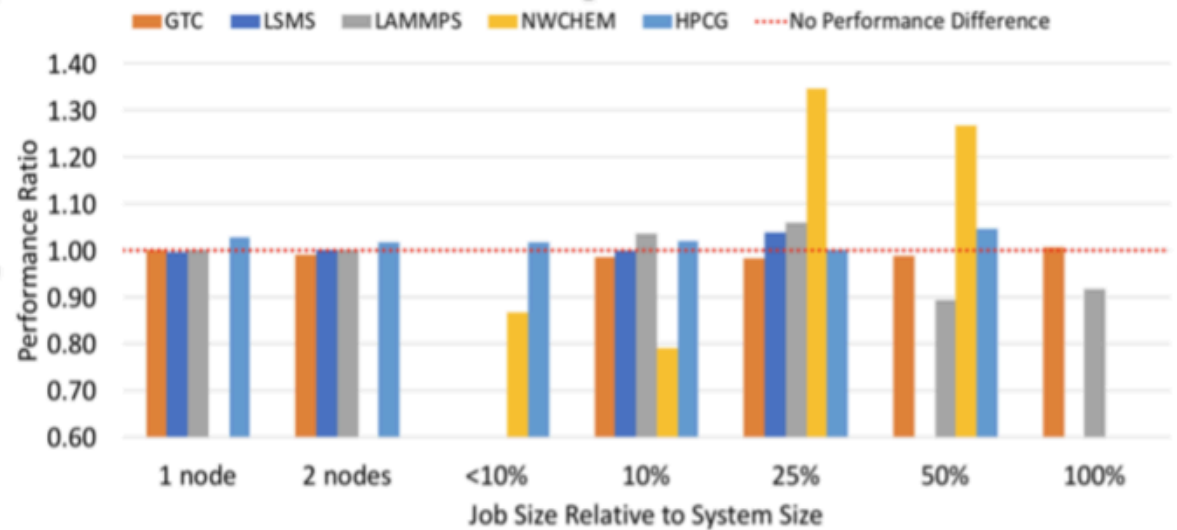
In the following we show performance after patches compared to before patches

We do this across the multiple codes tested, for each system



## Results: Titan Cray XK7

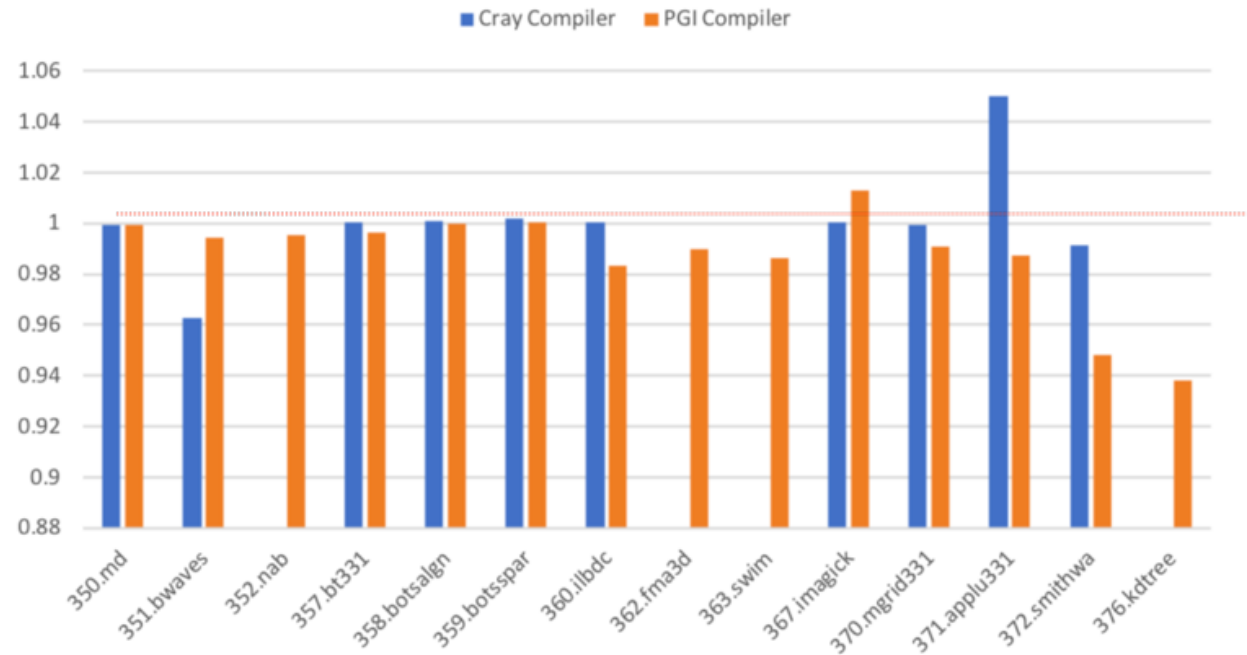
- **GTC**: differences are within run-to-run variability
- **LSMS**: minimal differences, slight improvement
- **LAMMPS**: lower at high node counts, but high run-to-run variability
- **NWCHEM**: higher at high node counts, but very high run-to-run variability
- **HPCG**: no degradation, slight improvement
- **OSB**: high variability, inconclusive
- **I/O benchmarks**: no significant performance difference



Performance after patch. Higher is better.

## Results: Titan Cray XK7

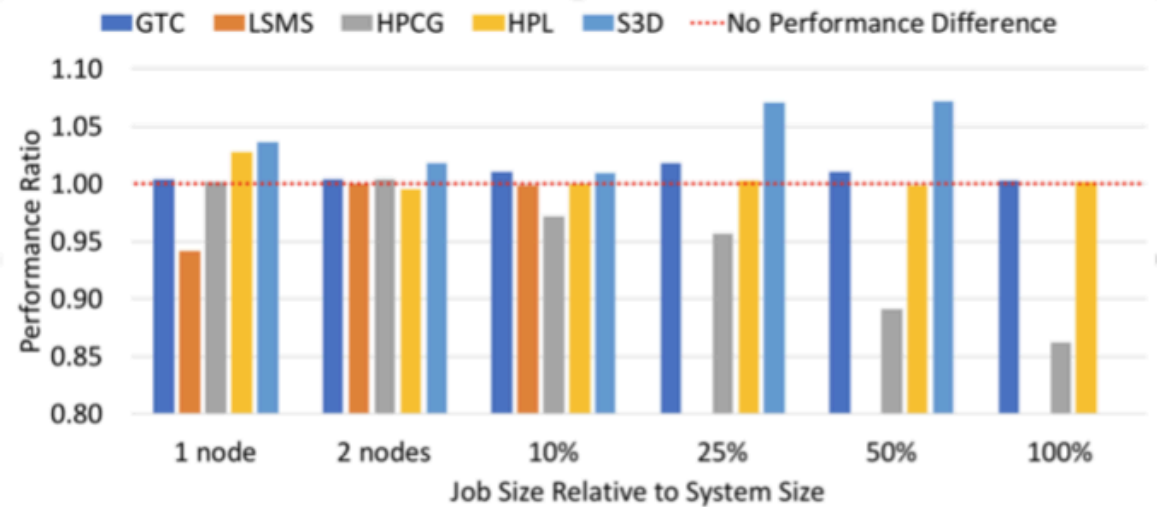
- **SPEC OMP2012:** mixed; for most cases minimal impact



Performance after patch. Lower is better.

## Results: Eos Cray XC30 (Xeon)


- **GTC**: up to 2% faster, higher than run-to-run variability
- **LSMS**: at higher node counts no significant difference
- **HPCG**: lower, but best trial for before and after patch nearly the same
- **HPL**: at higher node counts no significant difference
- **S3D**: up to 7% improvement at large node counts
- **OSB**: high variability, inconclusive



Performance after patch. Higher is better.

## Results: Eos Cray XC30

- **SPEC OMP2012:** in most cases small improvement in performance



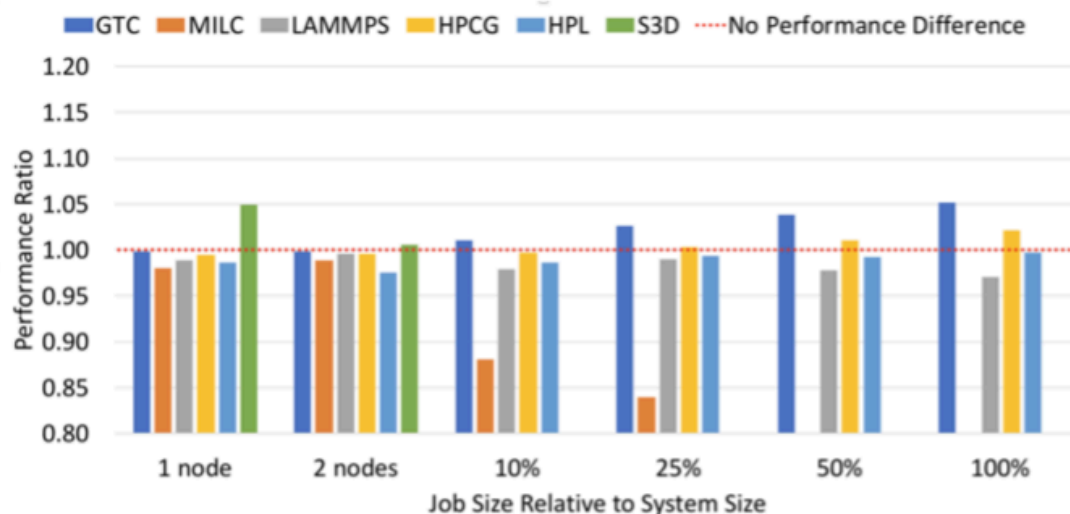
Performance after patch. Lower is better.

## Results: Cumulus Cray XC40 (Xeon)

- (patches not available at press time)

## Results: Percival Cray XC40 (Phi)

- **GTC**: up to 5% faster, higher than run-to-run variability
- **MILC**: modest degradation from 2% to 16%
- **LAMMPS**: 1-2% performance degradation
- **HPCG**: within 2% difference, faster for higher node counts
- **HPL**: less than 2% difference
- **S3D**: small performance improvement
- **OSB**: high variability, inconclusive



Performance after patch. Higher is better.

## Conclusions

- Results were mixed, but overall trend in performance seems to be that Spectre / Meltdown patches impact when considered in aggregate is minimal – no large systemic slowdown across codes
- Some apps on some systems are somewhat faster or slower after the patches, this may be of interest to the specific app users
- Also some differences accompanied by performance variability that are difficult to interpret, merit further investigation
- However, the overall impact to the OLCF scientific user community seems to be minor, containable
- The effect in most cases seems similar to some other software upgrades we have seen in the past – can potentially cause small unexpected increases or decreases in system performance
- We will continue to monitor performance changes as future patches may be released

**Questions?**  
**Wayne Joubert**  
**joubert@ornl.gov**

This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

