

# Spectre and Meltdown Performance Documentation

The following patches can be applied according to [RHE's CVE](#):

- [kernel-3.10.0-514.36.5.el7.x86\\_64.rpm](#)
- [linux-firmware-20160830-51.git7534e19.el7\\_3.noarch.rpm](#)
- [microcode\\_ctl-2.1-16.5.el7\\_3.x86\\_64.rpm](#)

- [GPFS IO Performance](#)
  - [Comments](#)
  - [Impact of 'Page Table Isolation \(pti\)' fix only](#)
- [Application Performance](#)
  - [MiniFE](#)
  - [Gear](#)
  - [CPMD](#)
  - [SPEC benchmark \(CPU FP SPEED\)](#)
  - [Quantum Espresso](#)

In order to quantify the impact of the Spectre / Meltdown fixes a number of tests have been run. Unless otherwise stated all results are from the Fidis cluster.


If all the recommended fixes are applied then we see a large impact on IO. As computational kernels do not require frequent calls to the operating system the impact is far less significant.

## GPFS IO Performance

**IOR API=POSIX**

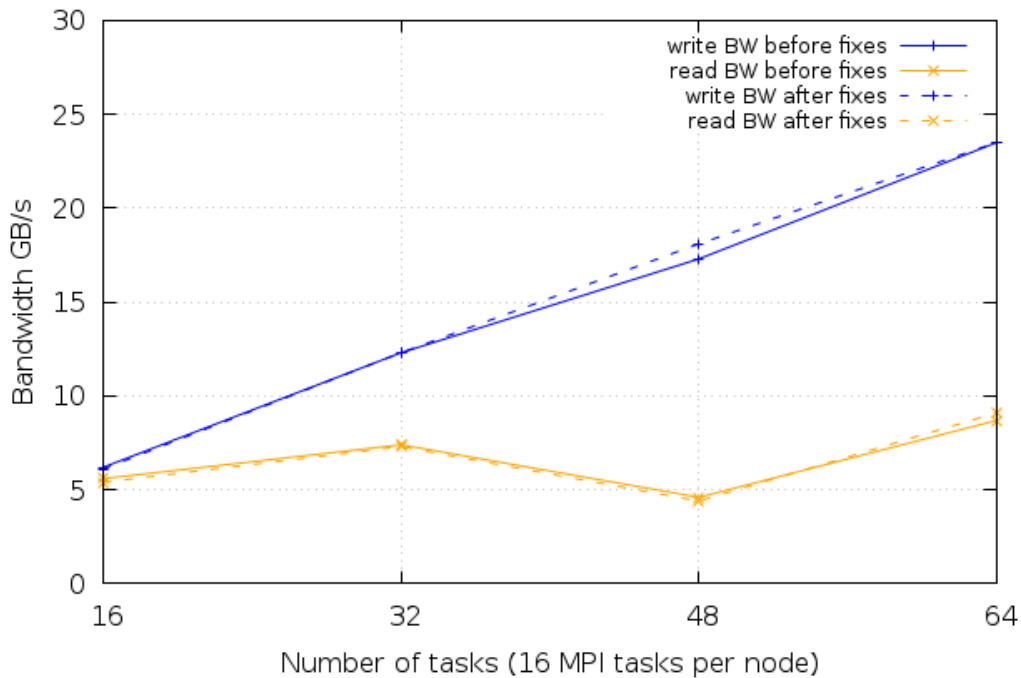
## IOR GPFS before and after Spectre and Meltdown fixes

16-Jan-2018, GPFS-4.2.2-2, Intel E5-2690-v4 2.60GHz, Linux 3.10.0-514.36.5  
IOR: api=POSIX, filePerProc=1, xfersize=16k, aggregate filesize=256GiB



## IOR before and after Spectre and Meltdown fixes


16-Jan-2018 IOR: api=POSIX, filePerProc=1, xfersize=16k, aggregate filesize=256GiB



IOR API=MPIO


## IOR before and after Spectre and Meltdown fixes

16-Jan-2018 IOR: api=MPIIO, filePerProc=0, xfersize=16k, aggregate filesize=256GiB



## IOR before and after Spectre and Meltdown fixes

16-Jan-2018 IOR: api=MPIIO, filePerProc=0, xfersize=16k, aggregate filesize=256GiB



### Comments

Impact is a loss of about 50% aggregate bandwidth when running IOR with one MPI task-per-node.

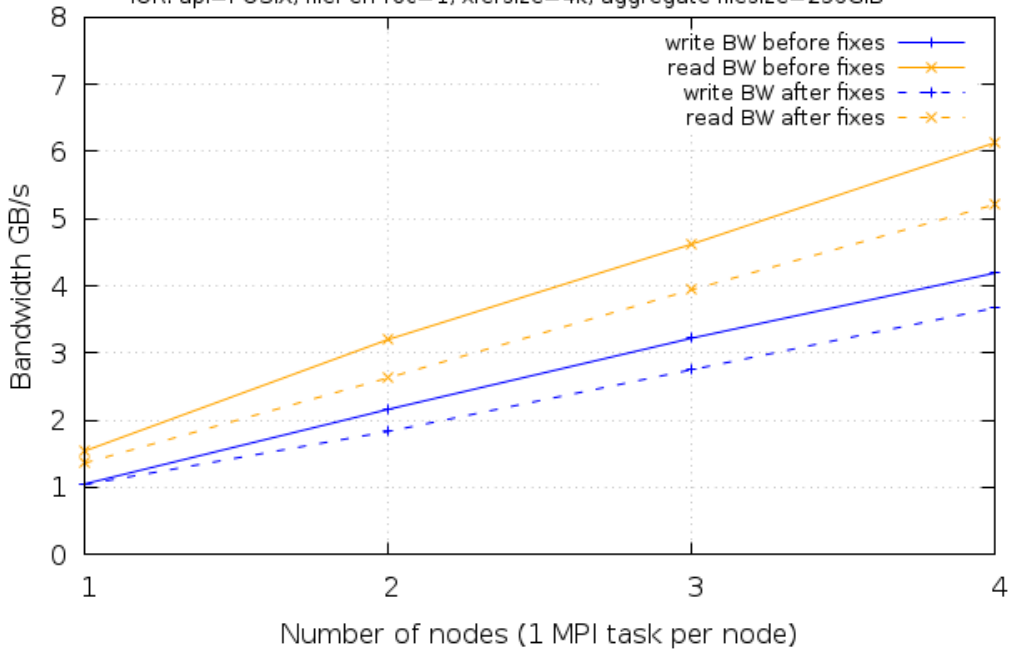
Impact is much less visible when running IOR with 16 MPI tasks-per-node.

### Impact of 'Page Table Isolation (pti)' fix only

Situation of Fidis compute nodes after maintenance of January 24th, 2018:

### IOR GPFS Fidis - Impact of Page Table Isolation (pti) fix

24-Jan-2018, GPFS-4.2.2-2, Intel E5-2690-v4 2.60GHz, Linux 3.10.0-514.36.5  
 IOR: api=POSIX, filePerProc=1, xfersize=4k, aggregate filesize=256GiB



## Application Performance

### MiniFE

Total running time in seconds on Fidis (lower is better)

nodes	size	before	after
4	1000x500x600	34.9368	33.9214
2	1000x500x600	67.2132	67.2067

### Gear

Large test case. Steps per unit of time (higher is better)

nodes	before	after
1	0.225932	0.224453
2	0.397746	0.395822
4	0.626741	0.622084

### CPMD

Average time per iteration in seconds (lower is better)

nodes	before	after
-------	--------	-------

1	3.49835	3.53886
2	6.13975	6.20212
4	3.18266	3.22177

## SPEC benchmark (CPU FP SPEED)

SPEC points (higher is better)

Before	After
67.9	66.6

## Quantum Espresso

```

espresso: version: 6.2.0
arch: platform: linux
platform_os: rhel7
target: x86_E5v4_Mellanox
compiler:
name: intel
version: 17.0.2
namespace: builtin
parameters:
elpa: true
hdf5: false
mpi: true
openmp: false
scalapack: true

```

Before	After
1m22.35s CPU 1m27.18s WALL	1m26.21s CPU 1m37.75s WALL