

A PERFORMANCE COMPARISON USING HPC BENCHMARKS: WINDOWS HPC SERVER 2008 AND RED HAT ENTERPRISE LINUX 5

R. Henschel, S. Teige, H. Li, J. Doleschal , M. S. Mueller

October 2010



INDIANA UNIVERSITY

Contents

- HPC at Indiana University
- FutureGrid
- Benchmarking Setup
 - Hardware
- Results
 - SPEC MPI2007 medium
 - NSF Acceptance Test Suite
 - HPCC
- Conclusion

HPC at Indiana University

- Indiana University
 - Founded in 1820
 - 8 campuses
 - Bloomington and Indianapolis
 - 80,000 students, 7,000 faculty
- HPC systems are operated by Research Technologies of University Information Technology Services
- Open access to graduate students, faculty and staff



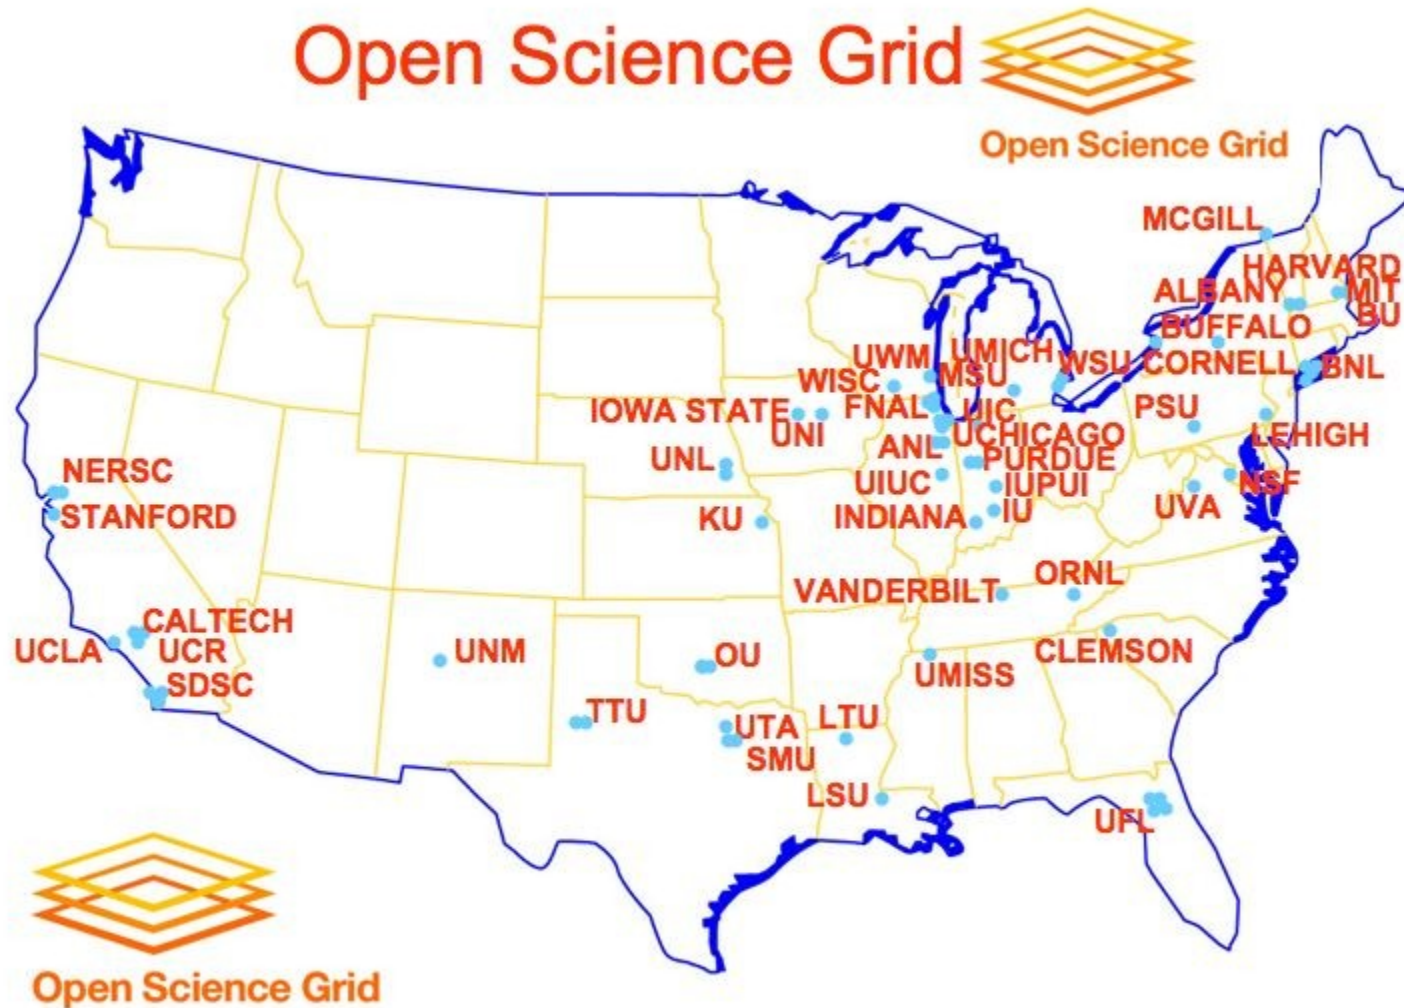
HPC at Indiana University

- IU is one of 11 TeraGrid resource providers



HPC at Indiana University

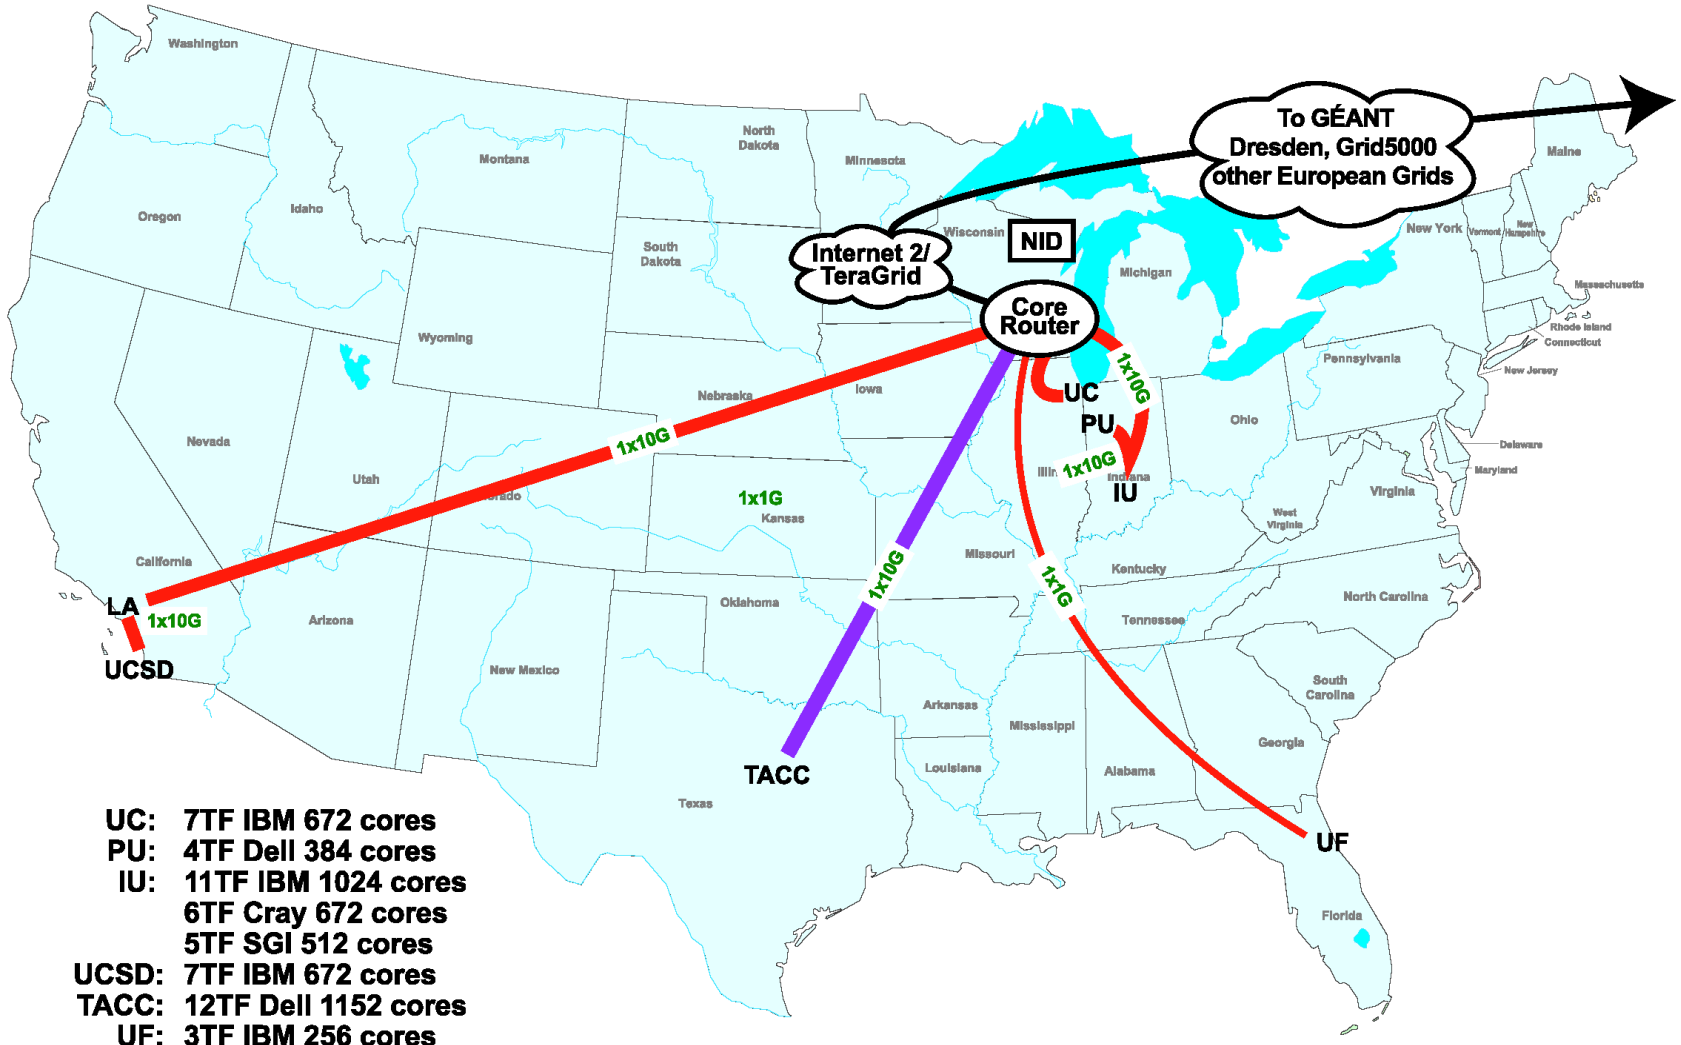
- IU hosts the Open Science Grid (OSG) operations center



FutureGrid

- \$15 Mio. NSF Track 2 D award (2009)
- Support the research on **the future of distributed, grid, and cloud computing.**
- Build a robustly managed **simulation environment and testbed** to support the development and early use in science of new technologies at all levels of the software stack: from networking to middleware to scientific applications.
- **Mimic TeraGrid** and/or general parallel and distributed systems.
- FutureGrid is a (small 5600 core) Science Cloud but it is more accurately a **virtual machine based simulation environment.**

FutureGrid



- UC:** 7TF IBM 672 cores
- PU:** 4TF Dell 384 cores
- IU:** 11TF IBM 1024 cores
- 6TF Cray 672 cores
- 5TF SGI 512 cores
- UCSD:** 7TF IBM 672 cores
- TACC:** 12TF Dell 1152 cores
- UF:** 3TF IBM 256 cores

Benchmarking Setup

- SPEC MPI2007 medium
- NSF Acceptance Test Suite
- HPCC

- Windows HPC Server 2008
- RedHat Enterprise Linux 5.4

- IBM System x iDataPlex dx340 cluster

- Intel Compiler Version 11.1 (incl. MKL)
- OpenMPI 1.3.1 / MS MPI 2008

Hardware

- IBM System x iDataPlex dx340
 - 84 nodes
 - 64 used for benchmarking
- Intel Xeon L5420 at 2.5 GHz
- 32 GByte of memory per node
- Mellanox MT26418 DDR Infiniband
- Cisco SFS 7024D switch



Benchmarks


- SPEC MPI2007 (medium)
- NSF Acceptance Test Suite
- HPCC

SPEC MPIM2007

- Developed by the SPEC High Performance Group
- Includes 13 MPI parallel applications
 - Computational fluid dynamics, molecular dynamics
Electromagnetism, geophysics, ray tracing, and
hydrodynamics
- Version 1.1, released October 2008
- Results were published on the SPEC website, after
having been reviewed by the HPG
 - <http://www.spec.org/mpi2007/results/res2010q1/>
- First published results running SPEC MPIM2007 on
Windows (HPC Server 2008)


SPEC MPIM2007

- Overall performance




SPEC MPIM2007

- No Difference in performance, 6 applications




SPEC MPIM2007

- RHEL5 outperforms WinHPC in 5 applications



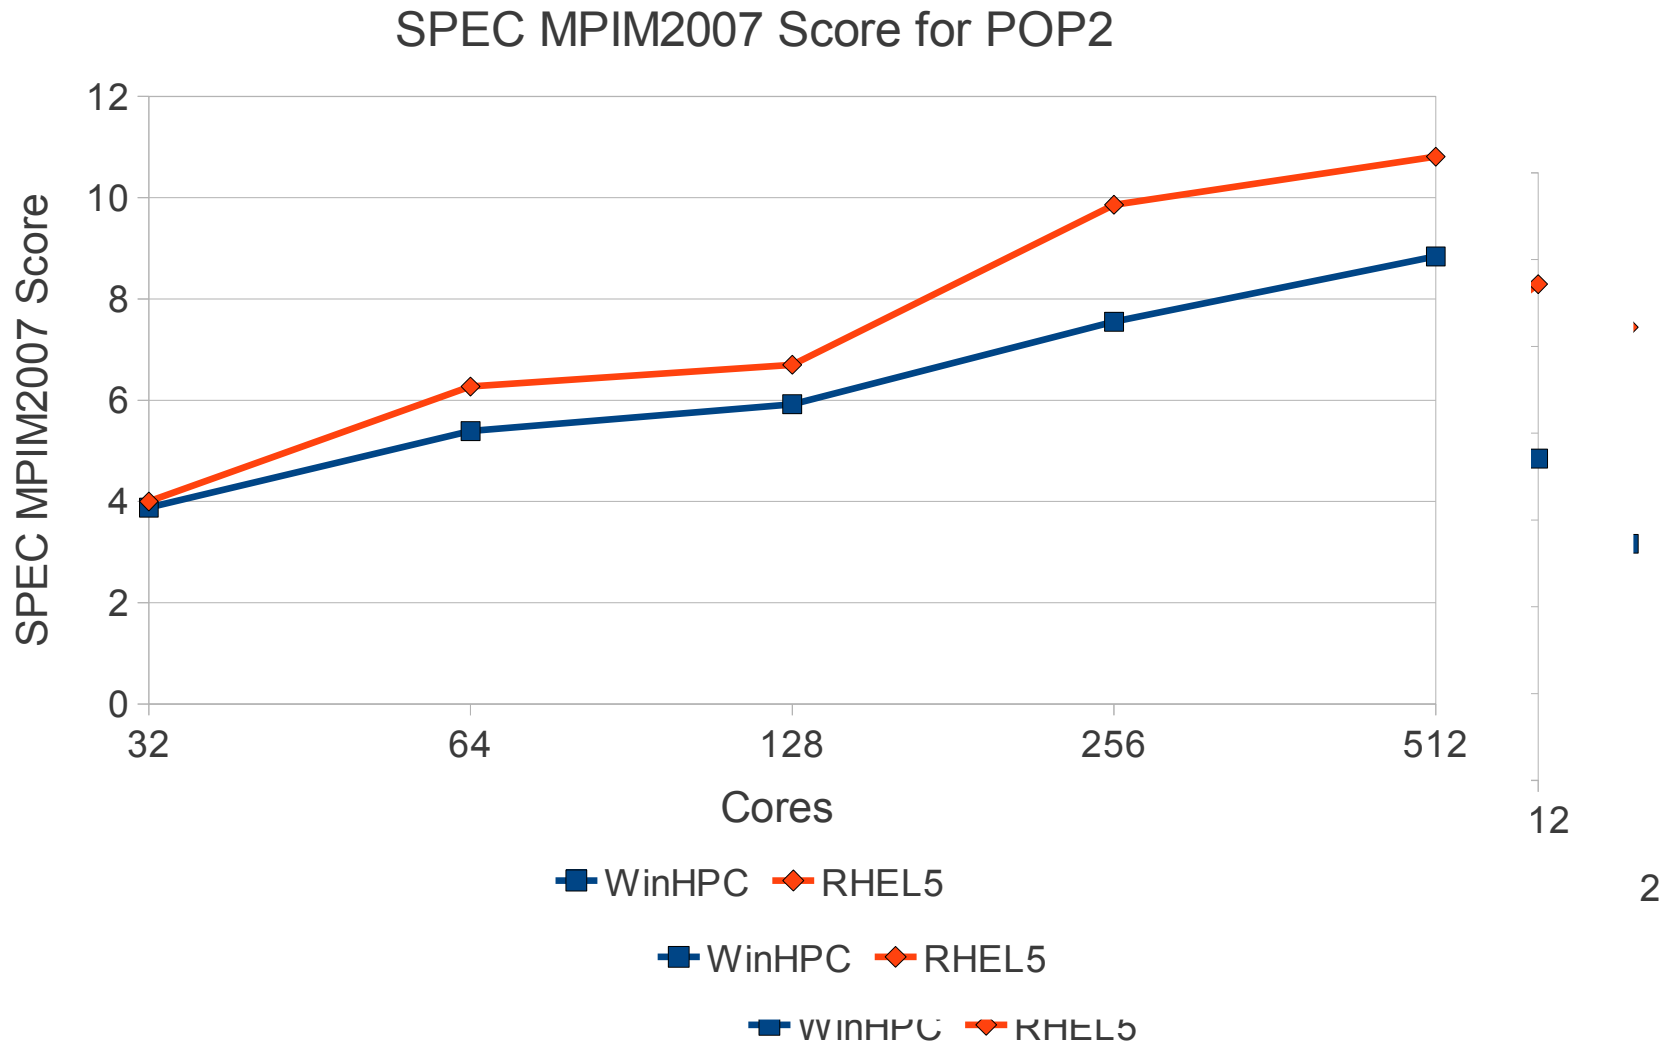
SPEC MPIM2007

- RHEL5 outperforms WinHPC in 5 applications




SPEC MPIM2007

- RHEL5 outperforms WinHPC in 5 applications



SPEC MPIM2007

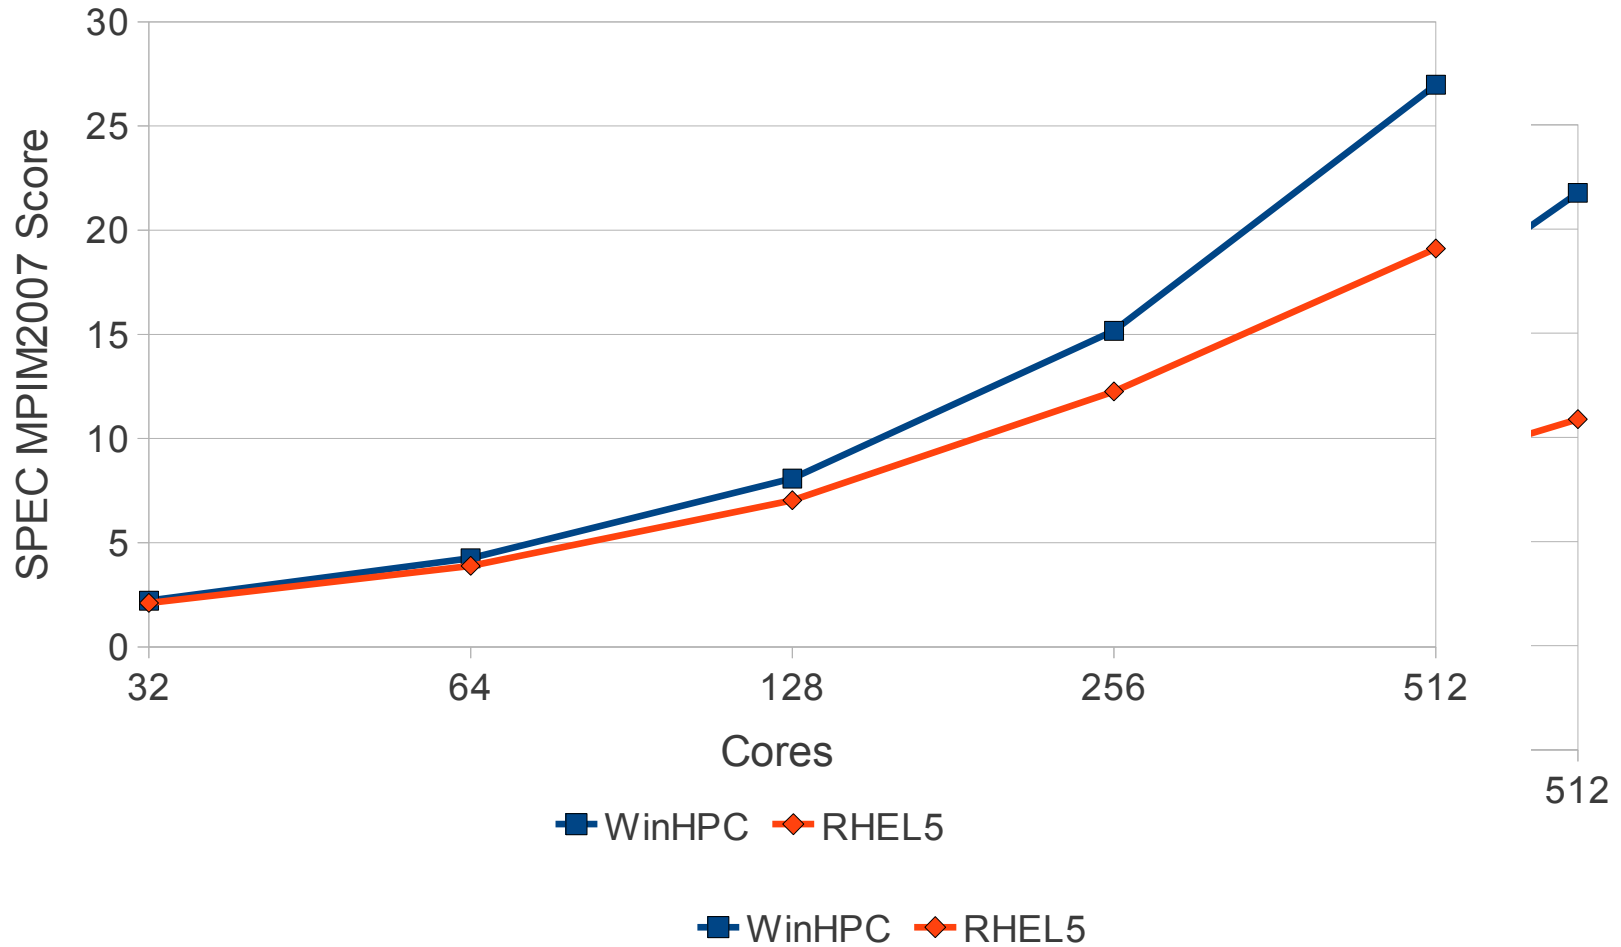
- WinHPC outperforms RHEL5 in 2 applications



SPEC MPIM2007

- WinHPC outperforms RHEL5 in 2 applications

SPEC MPIM2007 Score for LESLIE3D




NSF Acceptance Test Suite

- NSF Track 2 program
 - Sun constellation at TACC, Cray XT5 at University of Tennessee, FutureGrid
- Benchmarking Information Referenced in NSF 05-625 “High Performance Computing System Acquisition: Towards a Petascale Computing Environment for Science and Engineering” (2005)
 - <http://www.nsf.gov/pubs/2006/nsf0605/nsf0605.jsp>
- Initially comprised of HPCC, 6 applications, SPIOBENCH
- Now, just HPCC and 4 applications
 - PARATEC, HOMME, MILC, WRF

NSF Acceptance Test Suite


- MILC, PARATEC, HOMME
- WRF not included




- Innovative Computing Laboratory at the University of Tennessee
- Version 3.1.3
- 3 categories:
 - Floating Point
 - Memory
 - Interconnect

HPCC

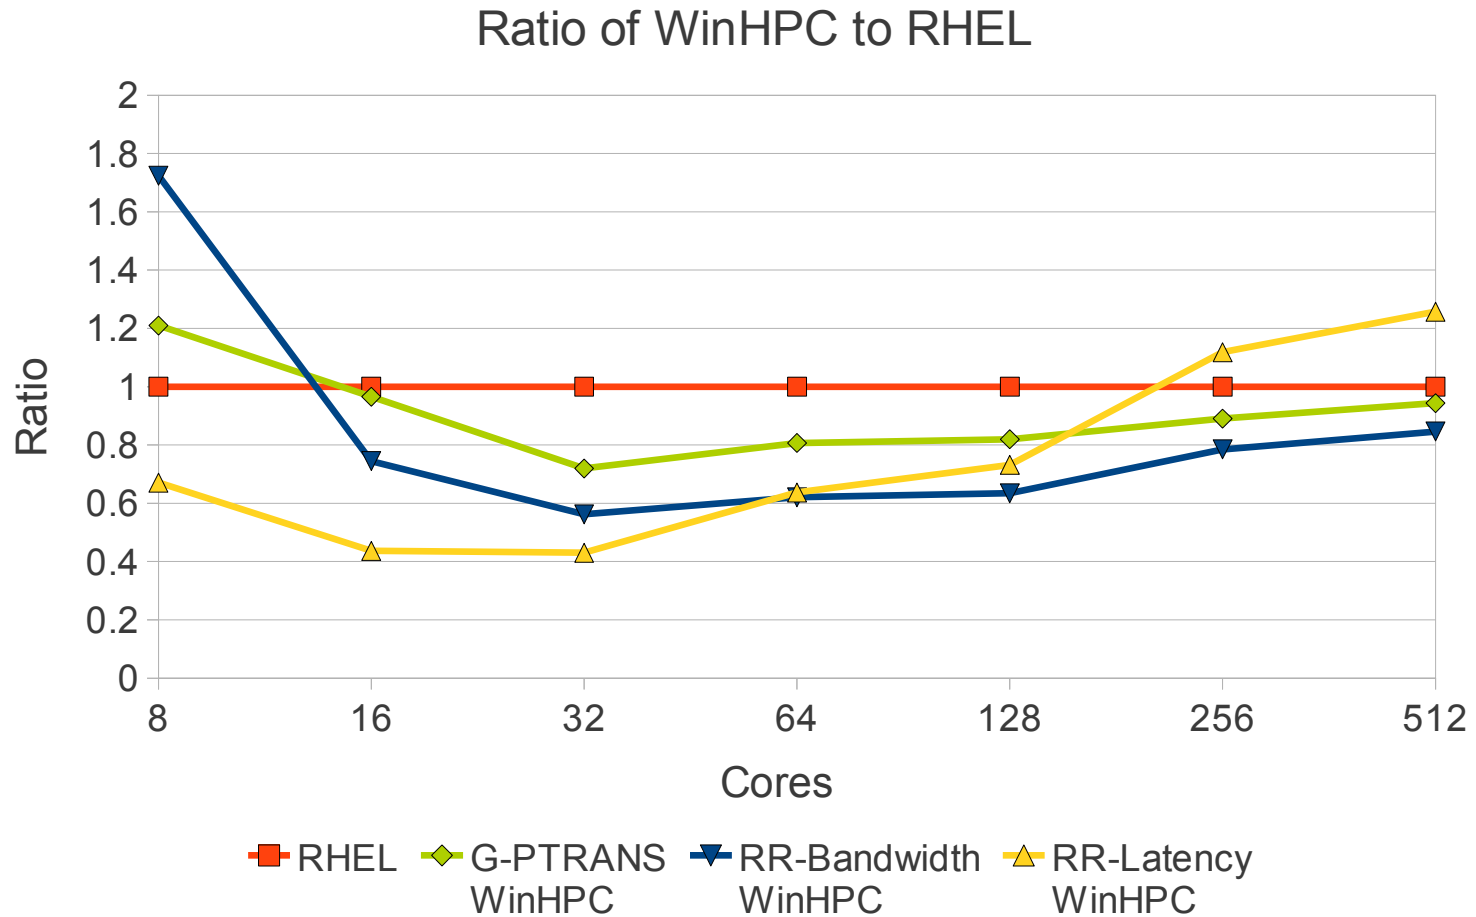
- Floating point tests (HPL, G-FFT)



- Memory performance tests (Random Access, Stream)



- Interconnect performance tests (G-PTRANS, RR-Bandwidth and Latency)



Conclusion

- Overall performance of WinHPC and RHEL5 is almost identical
- Certain applications scale better on RHEL5 than on WinHPC for larger core counts, while they perform very similar on smaller core counts
- When applications scale better on WinHPC, they do so across all core counts
- Building and running the benchmarks is more challenging on WinHPC

Acknowledgements

This document was developed with support from the National Science Foundation (NSF) under Grant No. 0910812 to Indiana University for "**FutureGrid: An Experimental, High-Performance Grid Test-bed.**" Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.

This material is based upon work supported by the National Science Foundation under Grants No. **ACI-0338618I, OCI-0451237, OCI-0535258, OCI-0504075, OCI-0910812 and CNS-0521433.**

This research was supported in part by the Indiana METACyt Initiative. The **Indiana METACyt Initiative** of Indiana University is supported in part by Lilly Endowment, Inc.

Thank You

Questions?