
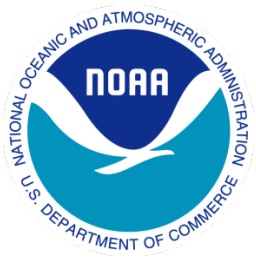


# A survey of performance characteristics of NOAA's weather and climate codes across our HPC systems

NOAA Research and Development High Performance Computing Office

Craig Tierney, U. of Colorado at Boulder

Leslie Hart, NOAA CIO Office

**Everything you read, see or hear about weather, climate and ocean forecasts in the US starts on NOAA's operational computer system**




2

# High Performance Computing and HPC Advanced Network

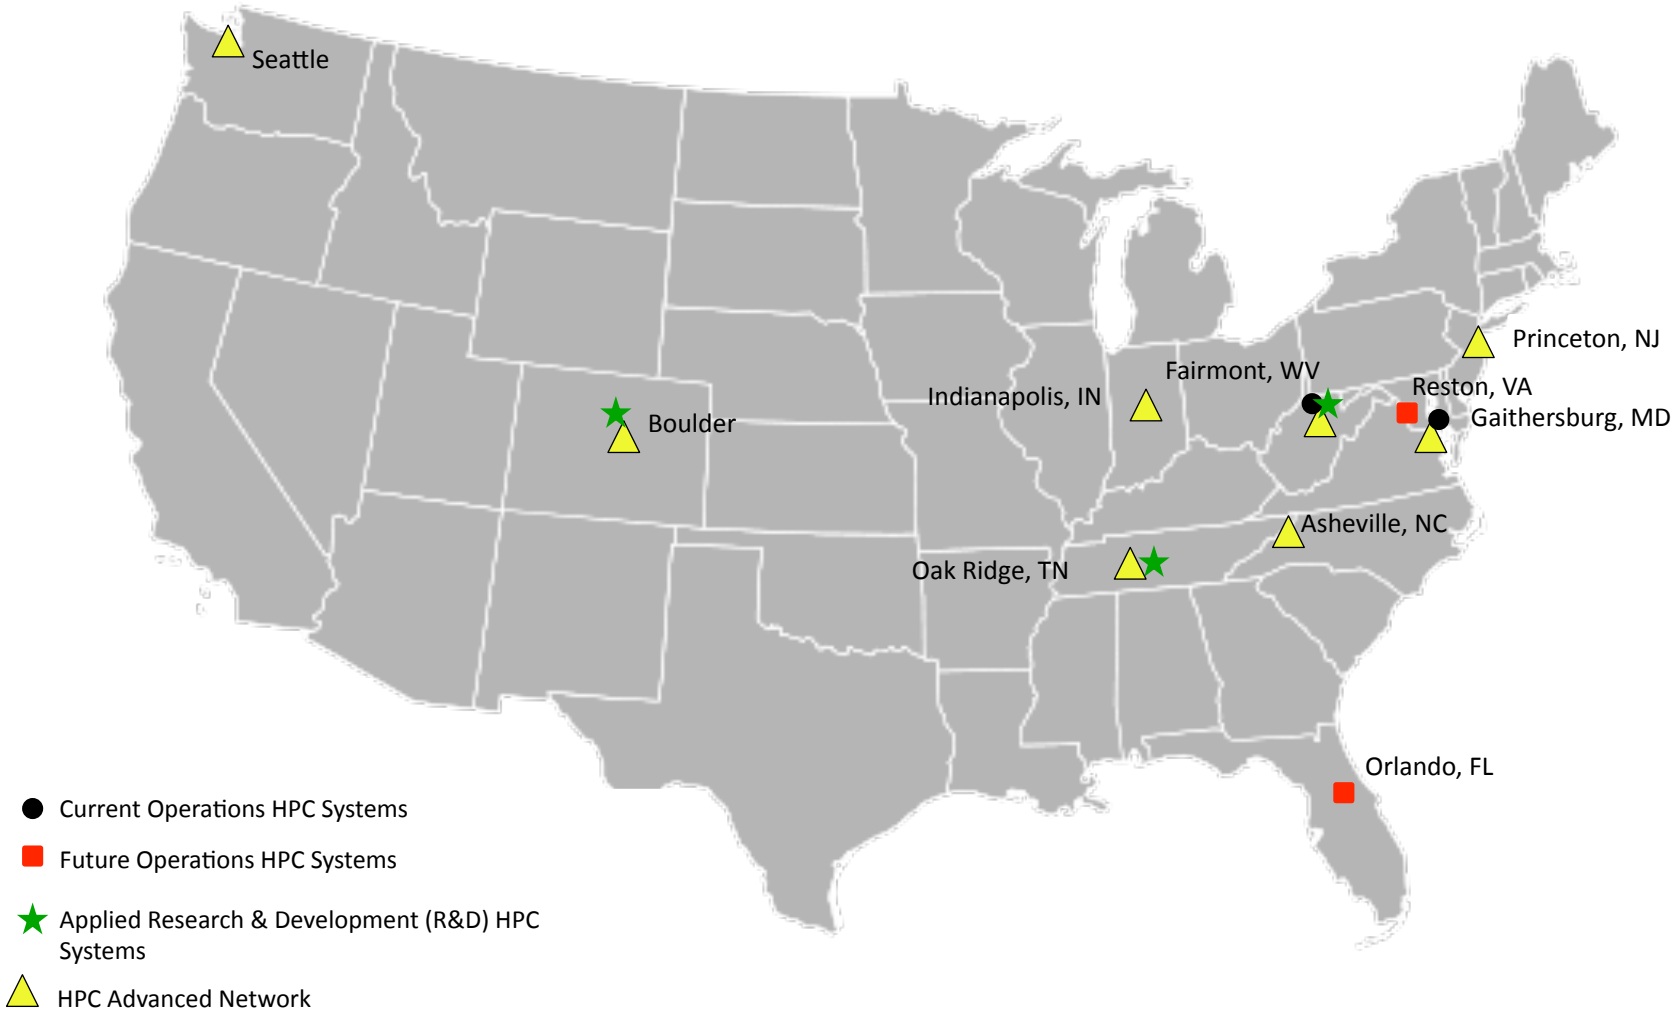

# NWAVE – NOAA Network

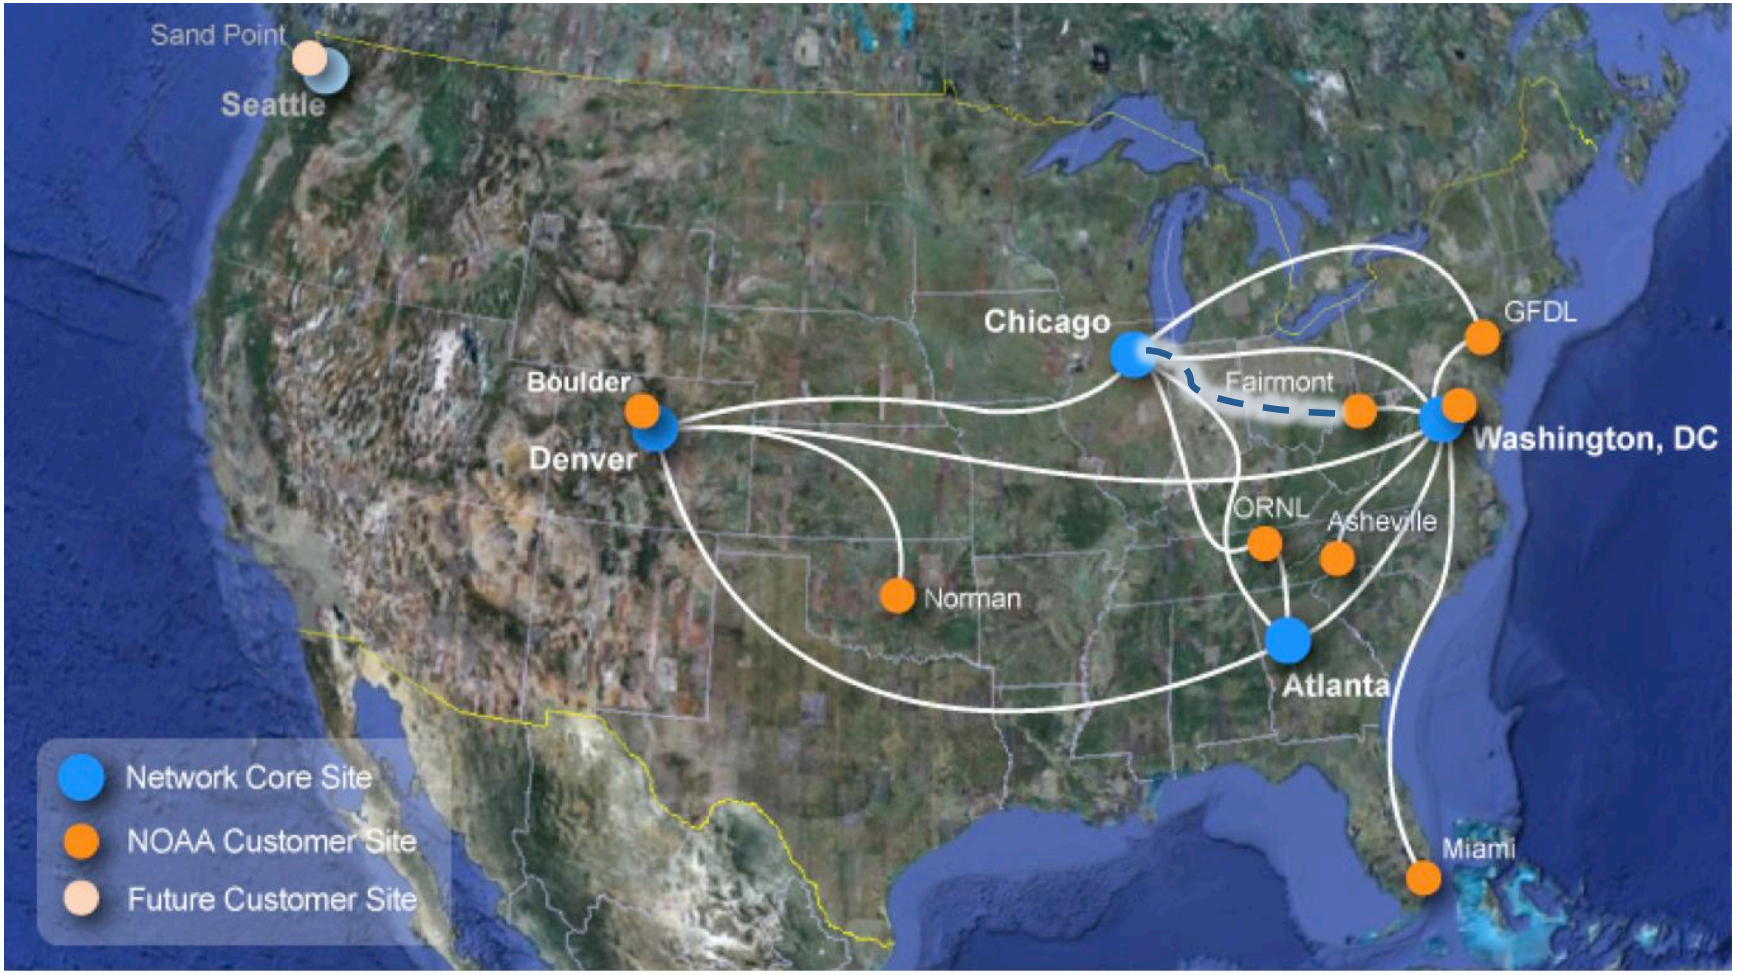

# High Performance Computing

## *Research & Development*

### Research HPC

- Location
  - Oak Ridge, TN (National Lab)
  - Interagency Agreement with DOE
- Significance

–Increase skill, resolution, and complexity of models used for climate change research and projections

### Development HPC

- Location
  - Fairmont, WV (The Vertex Center)
  - Boulder, CO
- Significance

–Support development of weather and seasonal to inter-annual climate model predictions bound for operational implementation

–Significantly improve hurricane track and intensity forecasts

–Systems include experimental architectures which allow for understanding of new technologies, drives future NOAA architectures

# High Performance Computing

## *Operational System*

### Location

- Current (Gaithersburg, VA and Fairmont WV)
- New Systems (Reston, VA and Orlando, VA)

### Configuration

- Redundant Systems
- Highly Reliable / Highly Available
  - 99.5% On-time Product Generation
  - Failover tested regularly
  - Each system highly available
- New Systems
  - IBM, Intel Sandy Bridge, ~200TF Each

### Inputs and Outputs

- Processes 3.5 billion observations/ day
- Produces over 15 million products/ day

### Significance

•Where our Nation's weather forecast process starts for the protection of lives and livelihood

•Produces model guidance at global, national, and regional scales

*Examples:*

- Hurricane Forecasts
- Aviation / Transportation
- Winter Weather
- Fire Weather

# Gaea

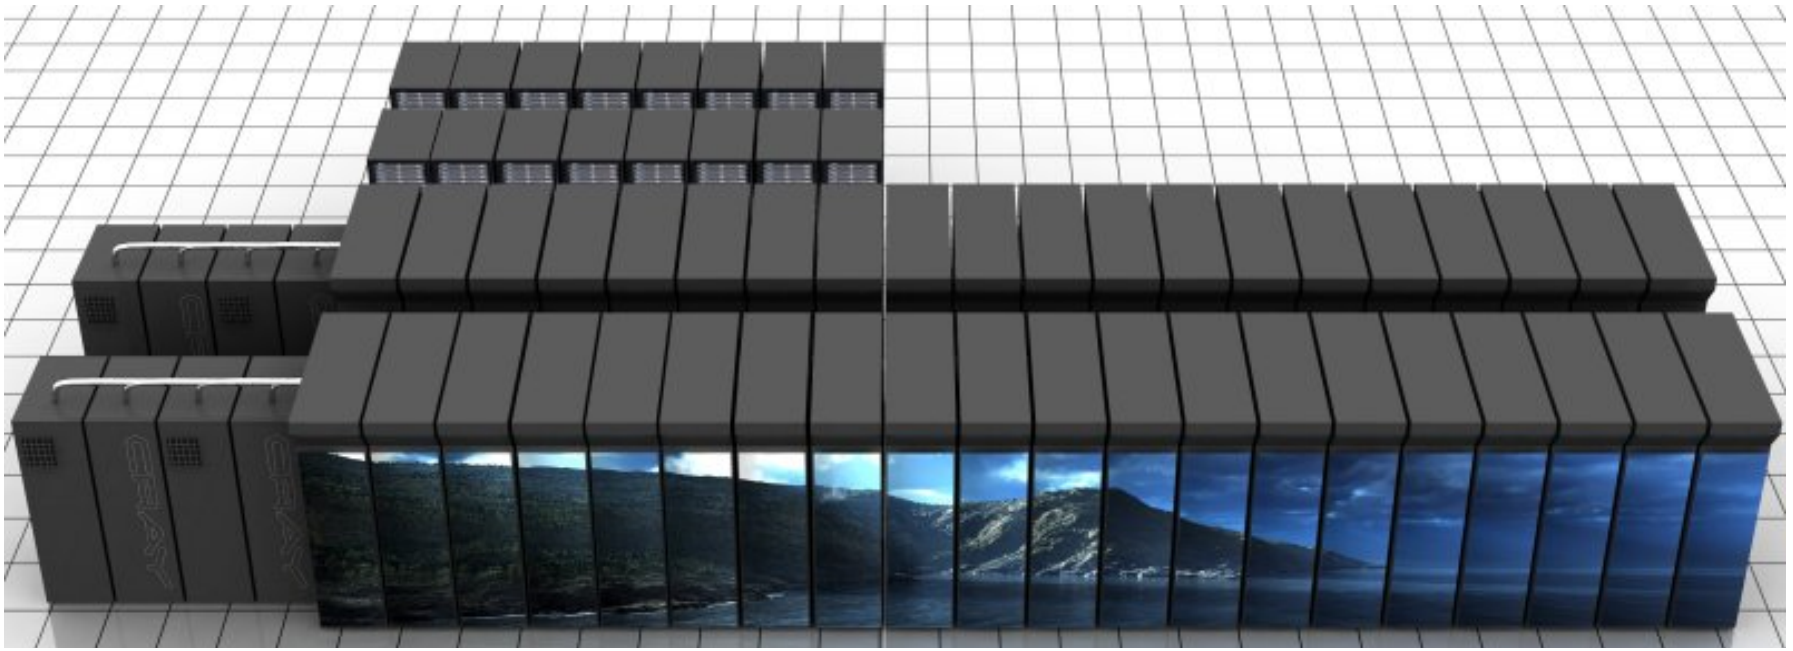

- Installed 2011/2012
- Oak Ridge, TN
  - Managed by DOE
- Cray XT6

- 120,320 Cores
  - AMD Bulldozer, 2.7GHz
- Cray Gemini
  - 3D Torus
- Workload – Mostly Climate

# Zeus

- Installed 2012
- Fairmont, WV
- SGI EX8400
- 27648 Cores
  - Intel Westmere, 3.46 GHz
- QDR Infiniband
  - Dual Rail Hypercube
- Workload
  - Mostly Weather/Climate
  - Supports real-time, deadline driven, experiments

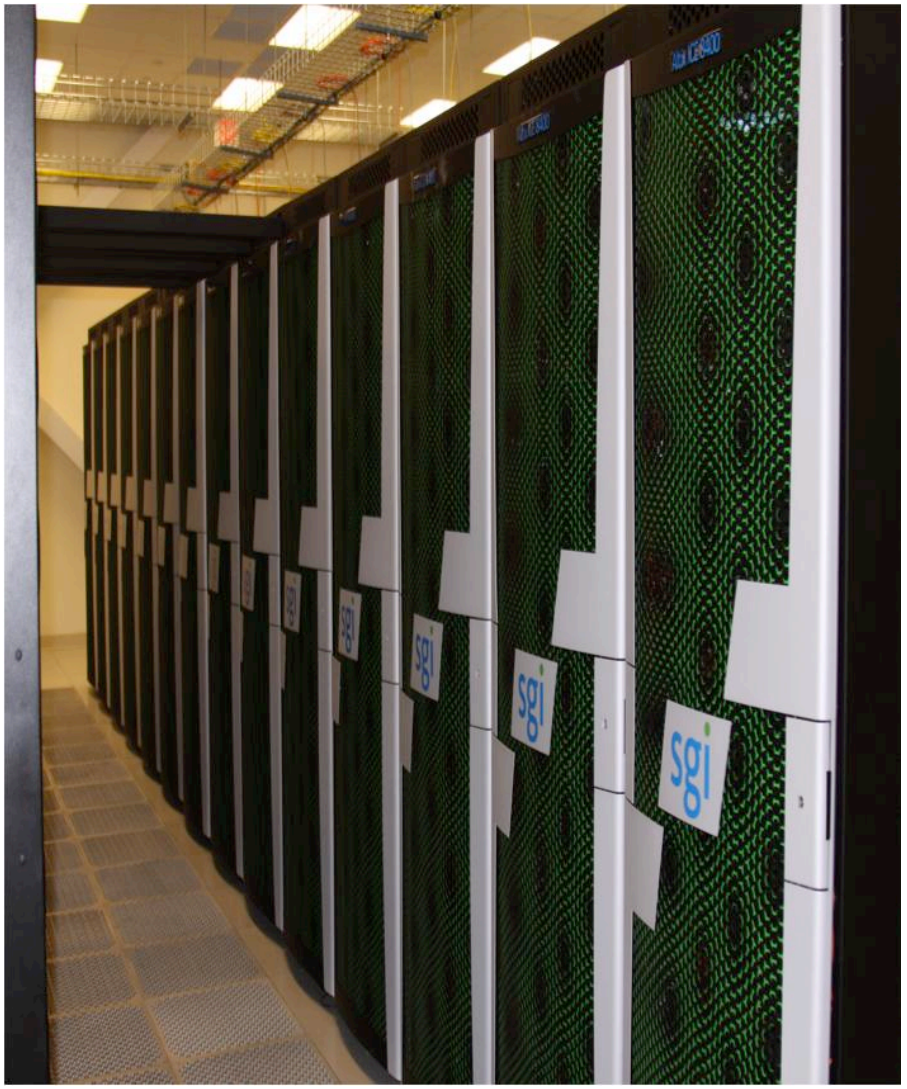

# nJet

- Installed 2009
- Boulder, CO
- Aspen Systems/Raytheon
- 3584 Cores
  - Intel Nehalem, 2.8 GHz
- QDR Infiniband
  - 70% Fat Tree
- Workload
  - Mostly Weather/Hurricane
  - Rapid Refresh Forecasting
  - Supports real-time, deadline driven, experiments

# tJet

- Installed 2010/2011
- Boulder, CO
- Aspen Systems/CSC
- 16176 Cores
  - Intel Westmere, 2.66 GHz
- QDR Infiniband
  - 70% Fat Tree
- Workload
  - Mostly Hurricane
  - Supports real-time, deadline driven, experiments

# sJet

- Installed 2012
- Boulder, CO
- Appro Supercomputer/CSC
- 5440 Cores
  - Intel Sandy Bridge, 2.6 GHz
- QDR Infiniband
  - 70% Fat Tree
- Workload
  - Mostly Hurricane
  - Supports real-time, deadline driven, experiments

Supercomputer Solutions

# Weather Research Framework (WRF)

- Non-hydrostatic, mesoscale, atmospheric simulation system
- Used for both research and operational forecasting
  - Regional Weather and Hurricane Forecasting
- Curve-linear, rotated grid
- Standard WRF benchmark used
  - Benchmark, 3km, (1501x1201x35)
  - 3 hour forecast
  - I/O timing not included

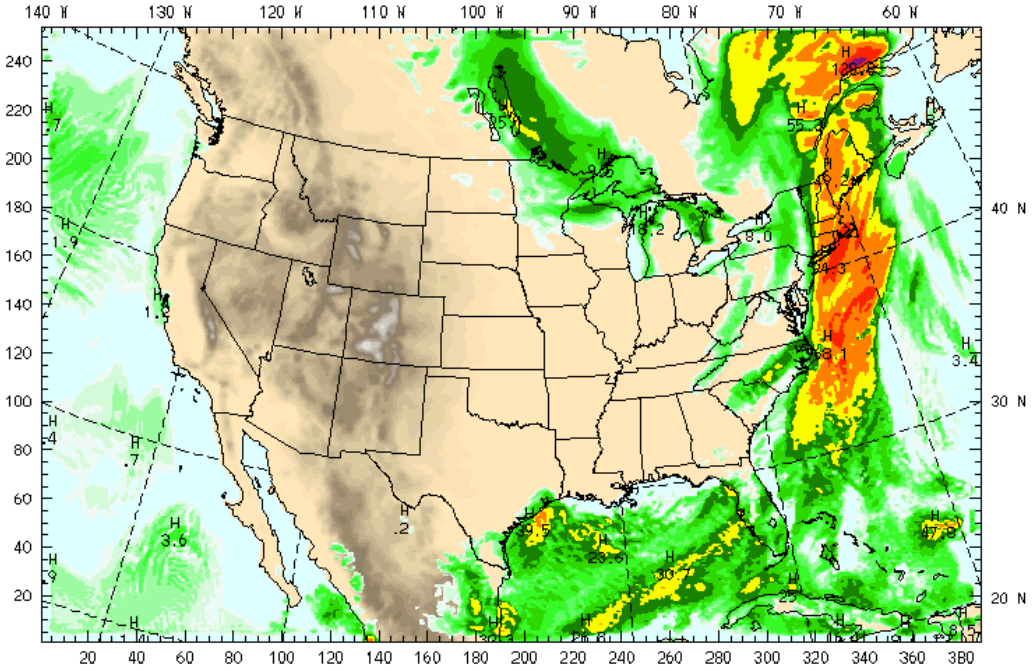


Website: http://www.wrf-model.org
WRF Benchmark Website: www.mmm.ucar.edu/wrf/WG2/bench
Image source: http://box.mmm.ucar.edu/

# Flow-following, Finite-Volume, Icosahederal Model (FIM)

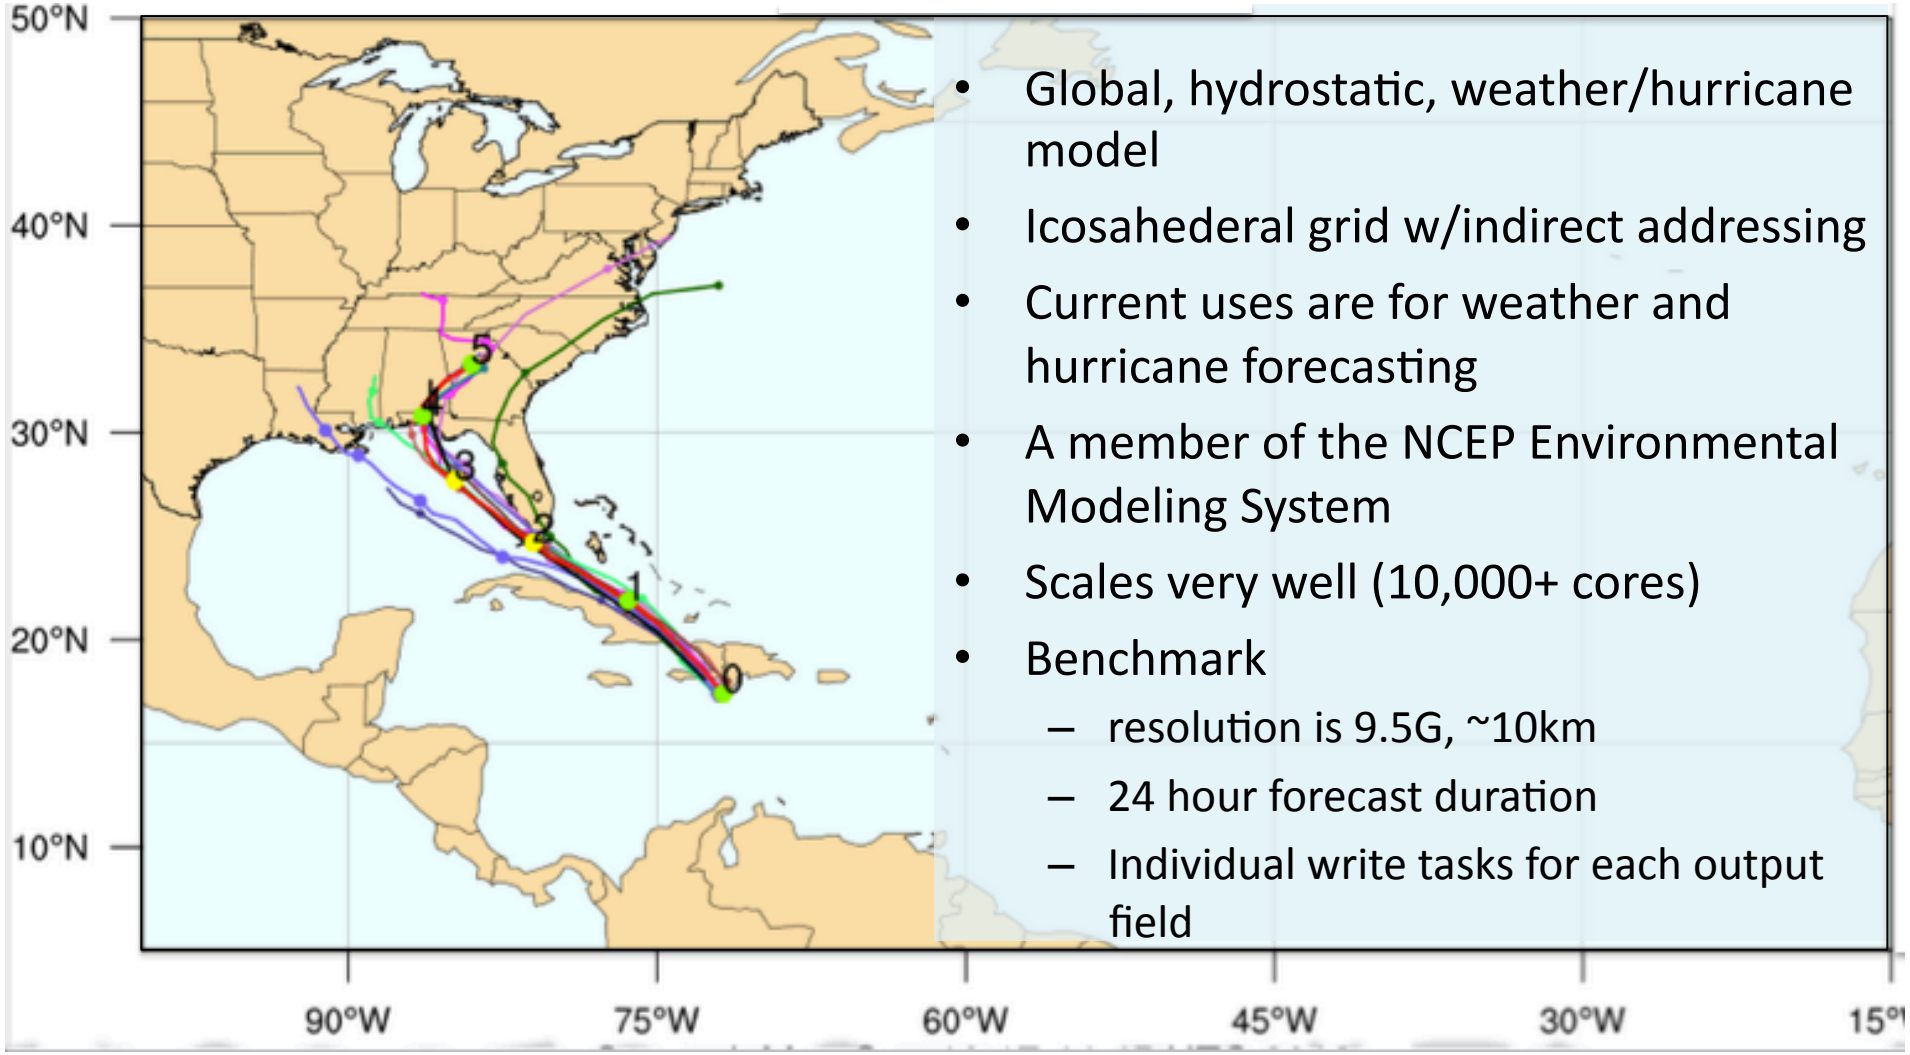


- Global, hydrostatic, weather/hurricane model
- Icosahederal grid w/indirect addressing
- Current uses are for weather and hurricane forecasting
- A member of the NCEP Environmental Modeling System
- Scales very well (10,000+ cores)
- Benchmark
  - resolution is 9.5G, ~10km
  - 24 hour forecast duration
  - Individual write tasks for each output field

# Global Forecast System (GFS)

- Global Numerical Weather Prediction System
- Member of the NCEP Environmental Modeling System

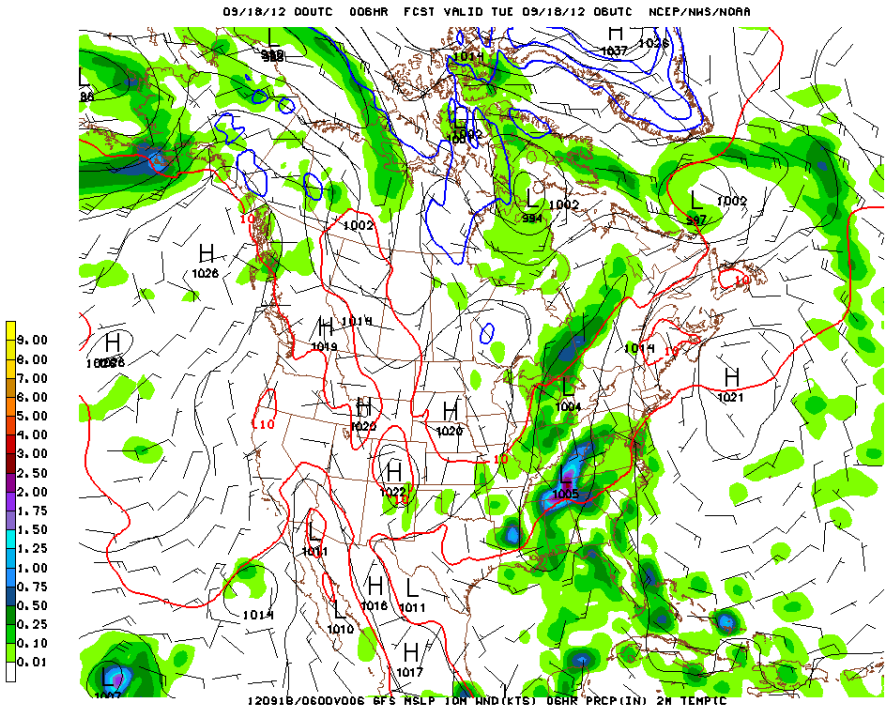


- Spectral horizontal representation
- Benchmark configuration
  - resolution is t574, ~27km
    - Matches current operational configuration
  - 24 hour forecast
  - I/O cache enabled

Website: http://www.emc.ncep.noaa.gov/index.php?branch=GFS

# Performance Comparisons

- Codes compiled as average user would (-O3 –x $ARCH etc).
  - Additional compiler options may provide a few percent
- Comparisons are per core and per node
  - Solely per core comparisons not fair
- NOAA cares more about high-throughput systems, not grand challenge problems
  - Forecasts, Ensembles, Retrospective studies

# Stream Results

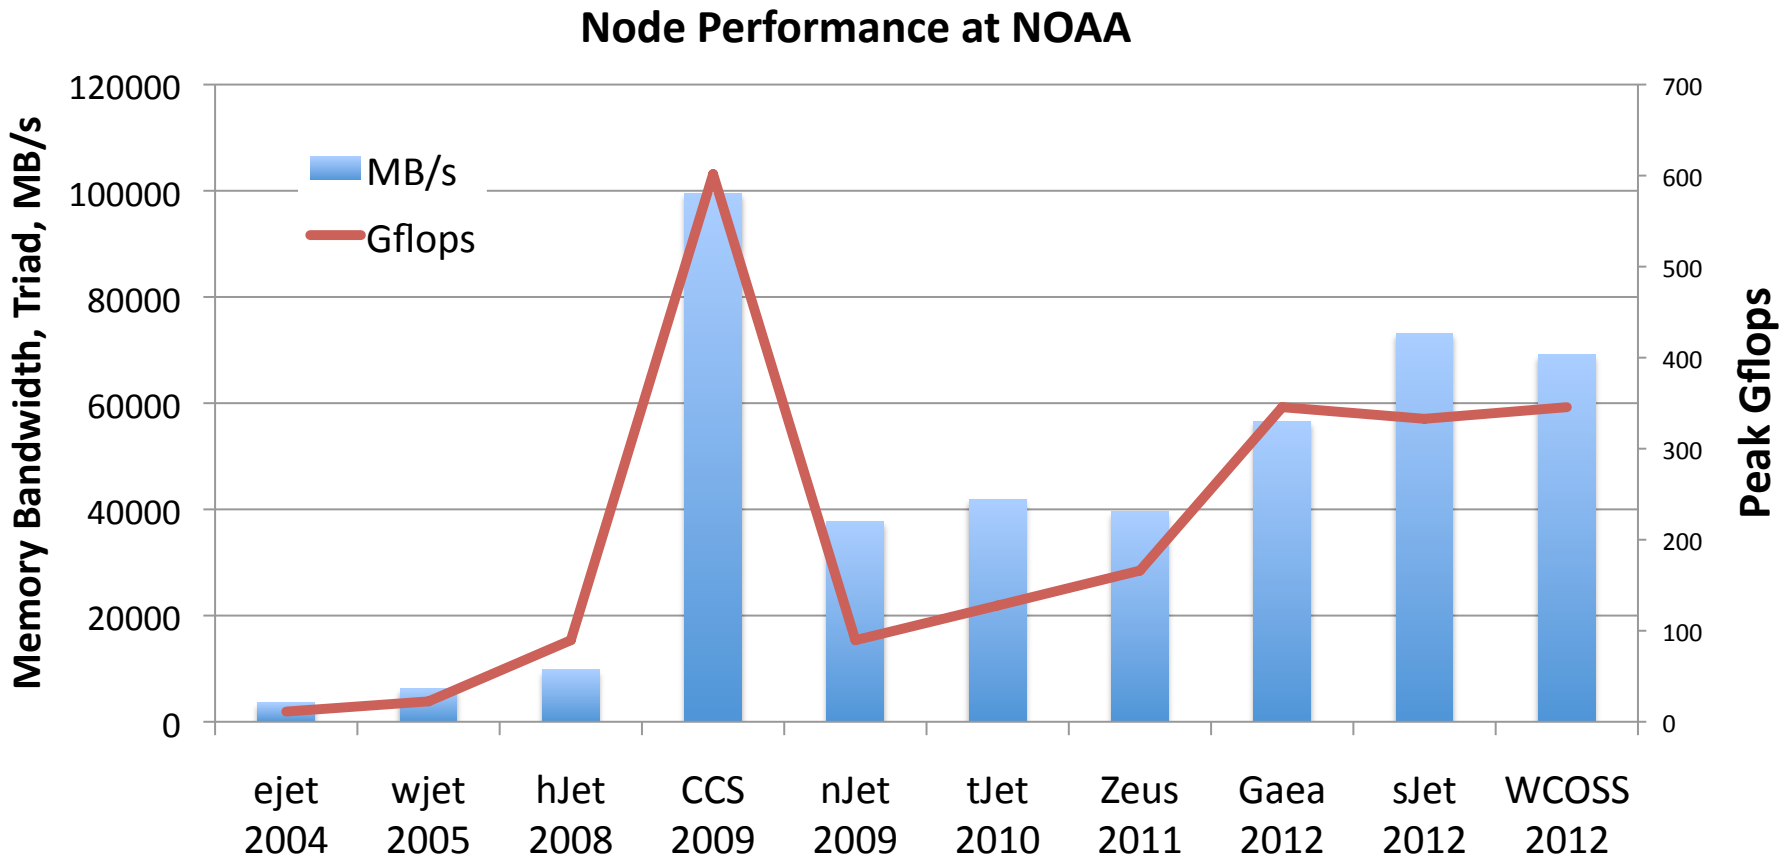


Stream Benchmark: http://www.cs.virginia.edu/stream/

# WRF Performance Results

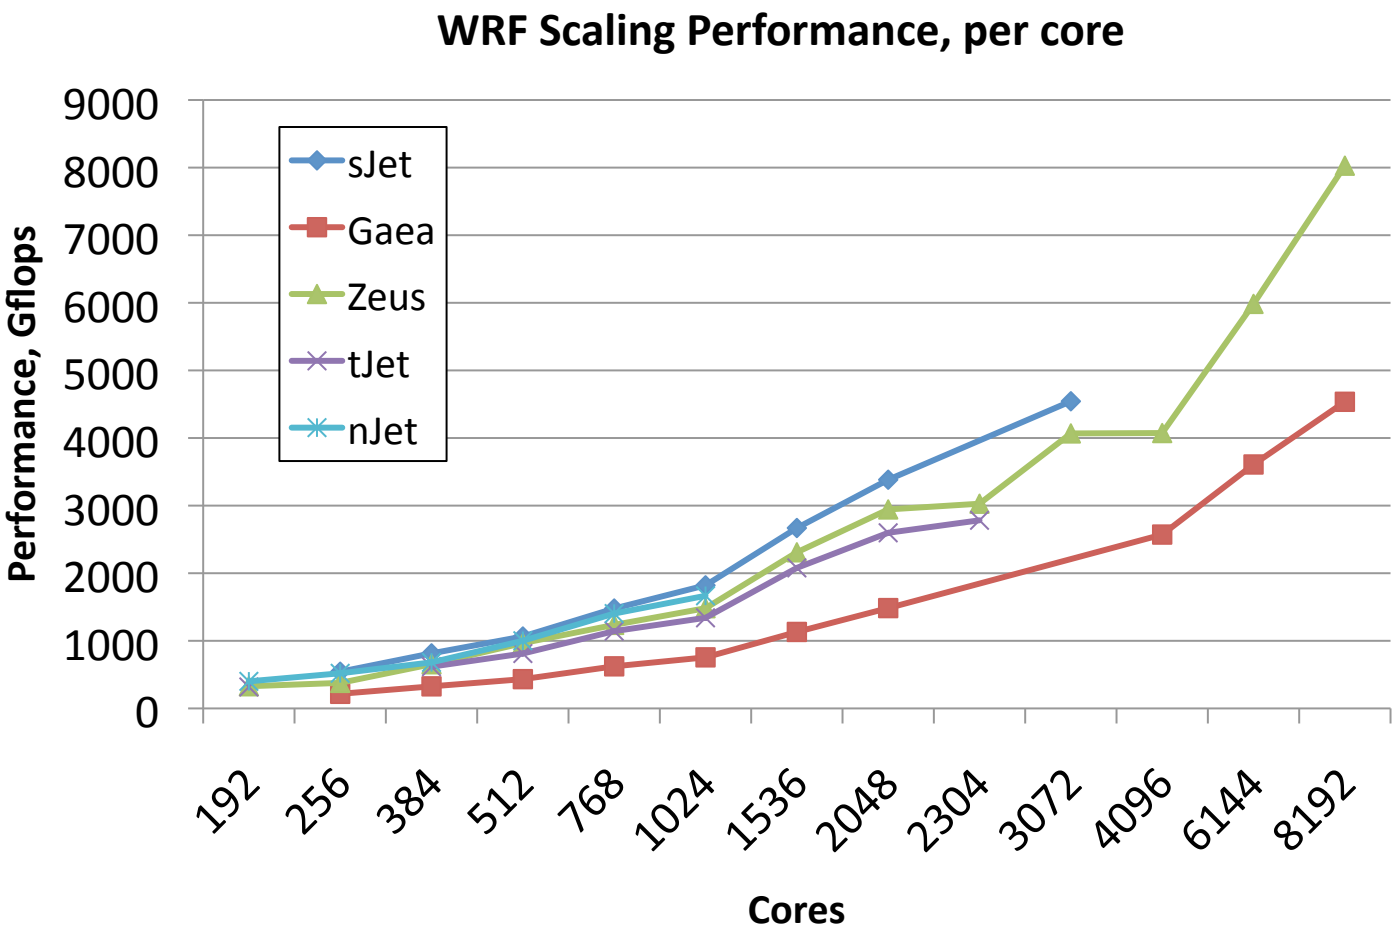

# WRF Performance Results

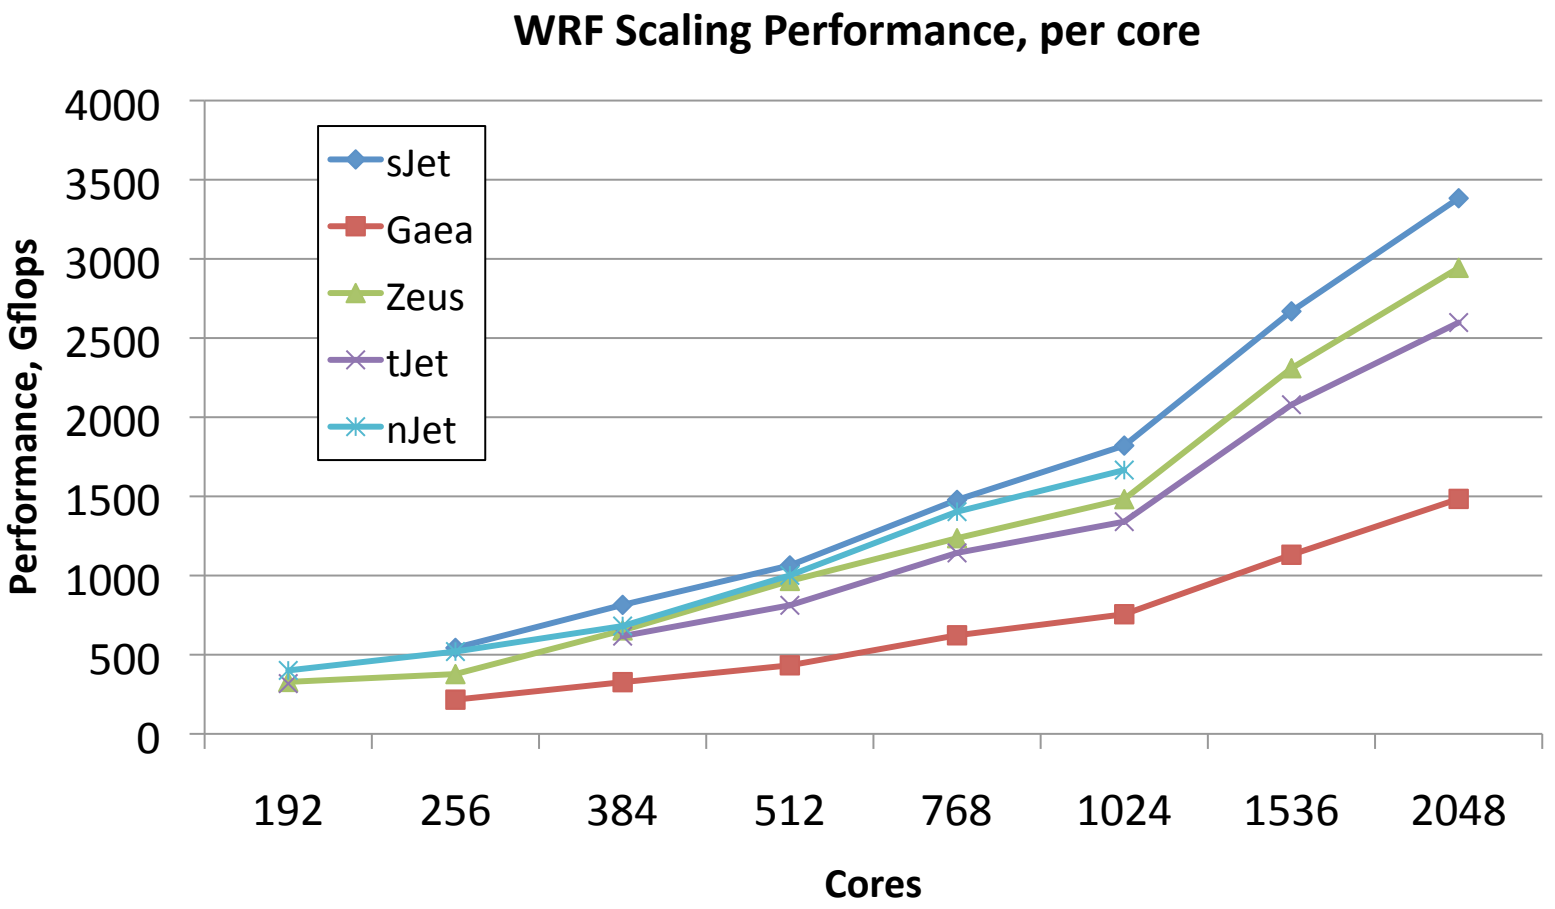

# WRF Performance Results

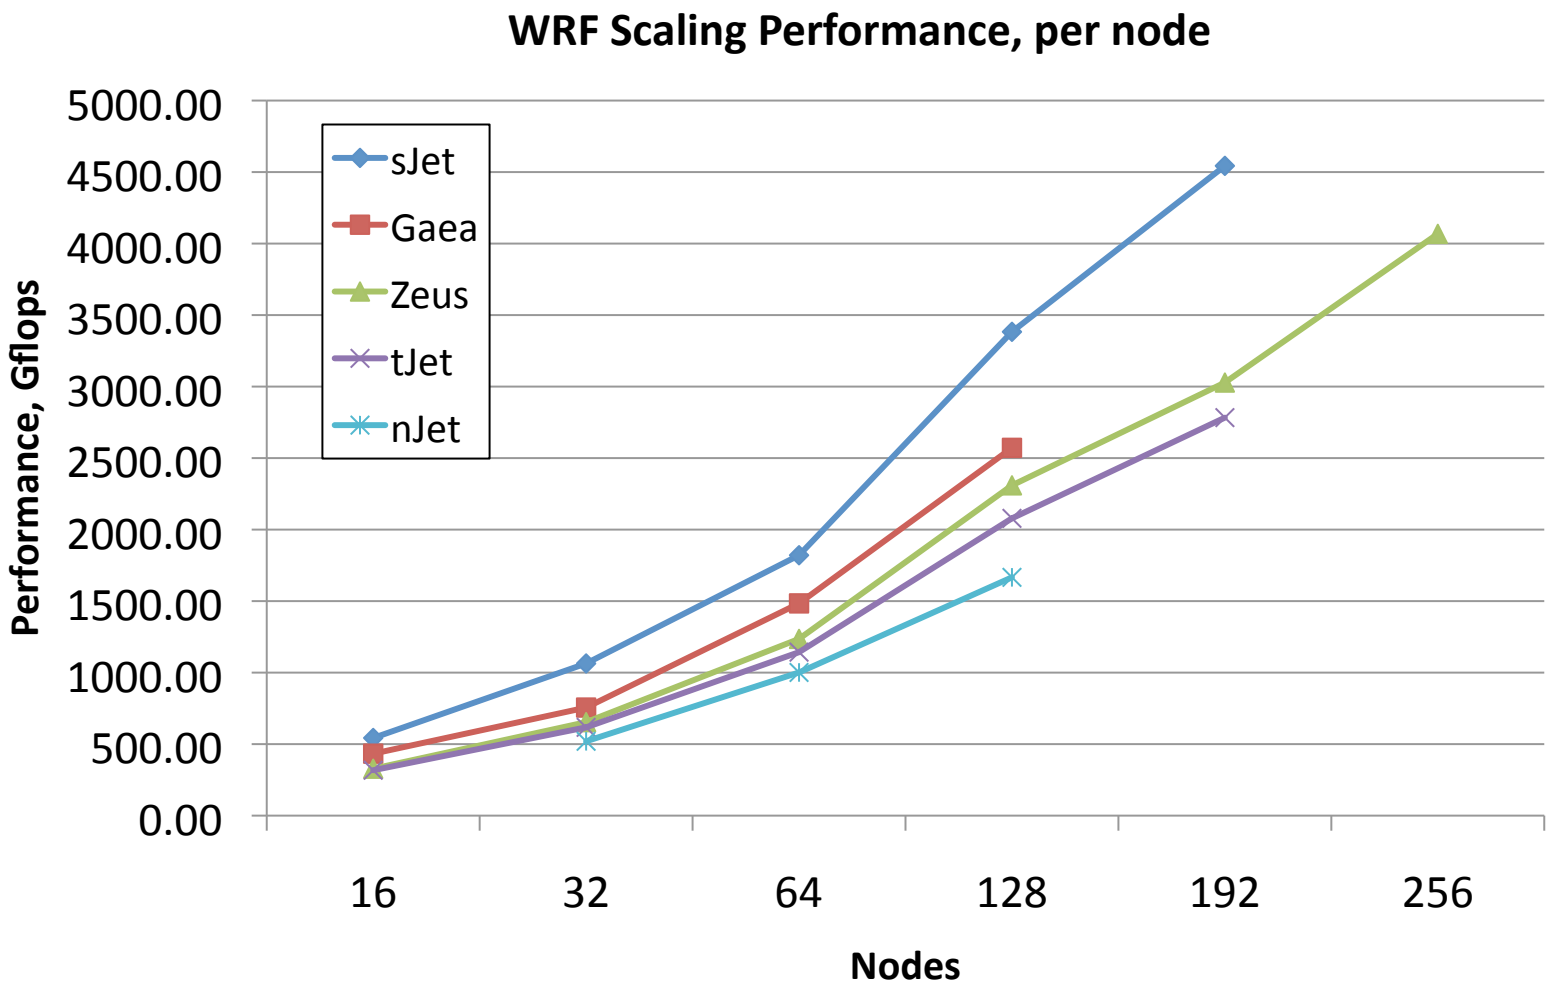

# GFS Performance Results

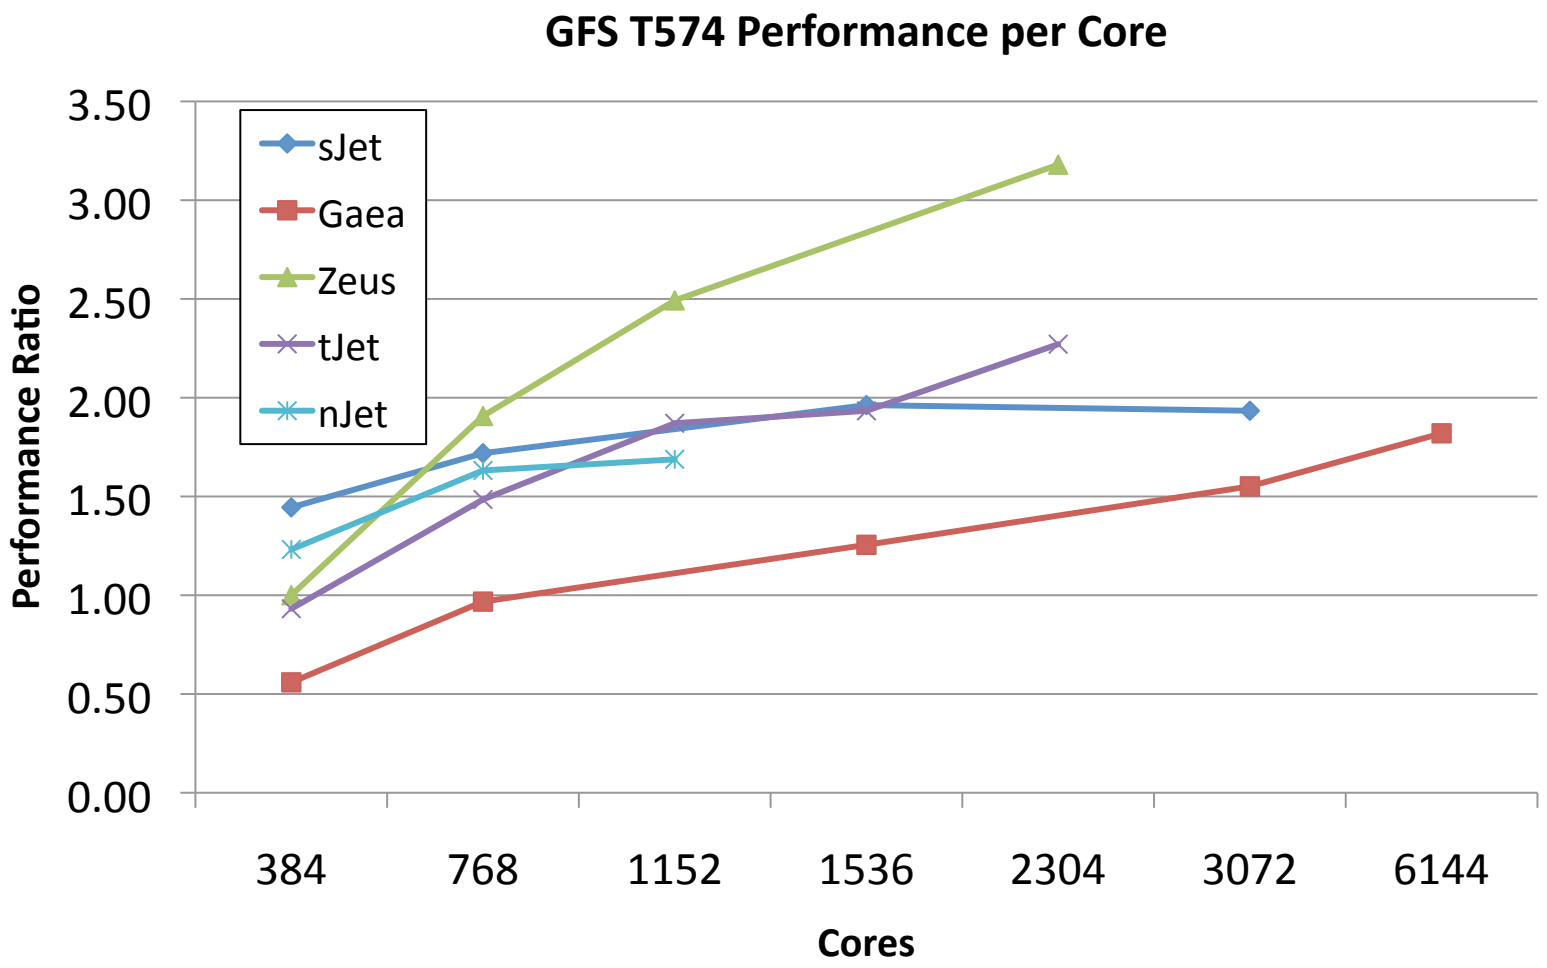

# FIM Performance Results

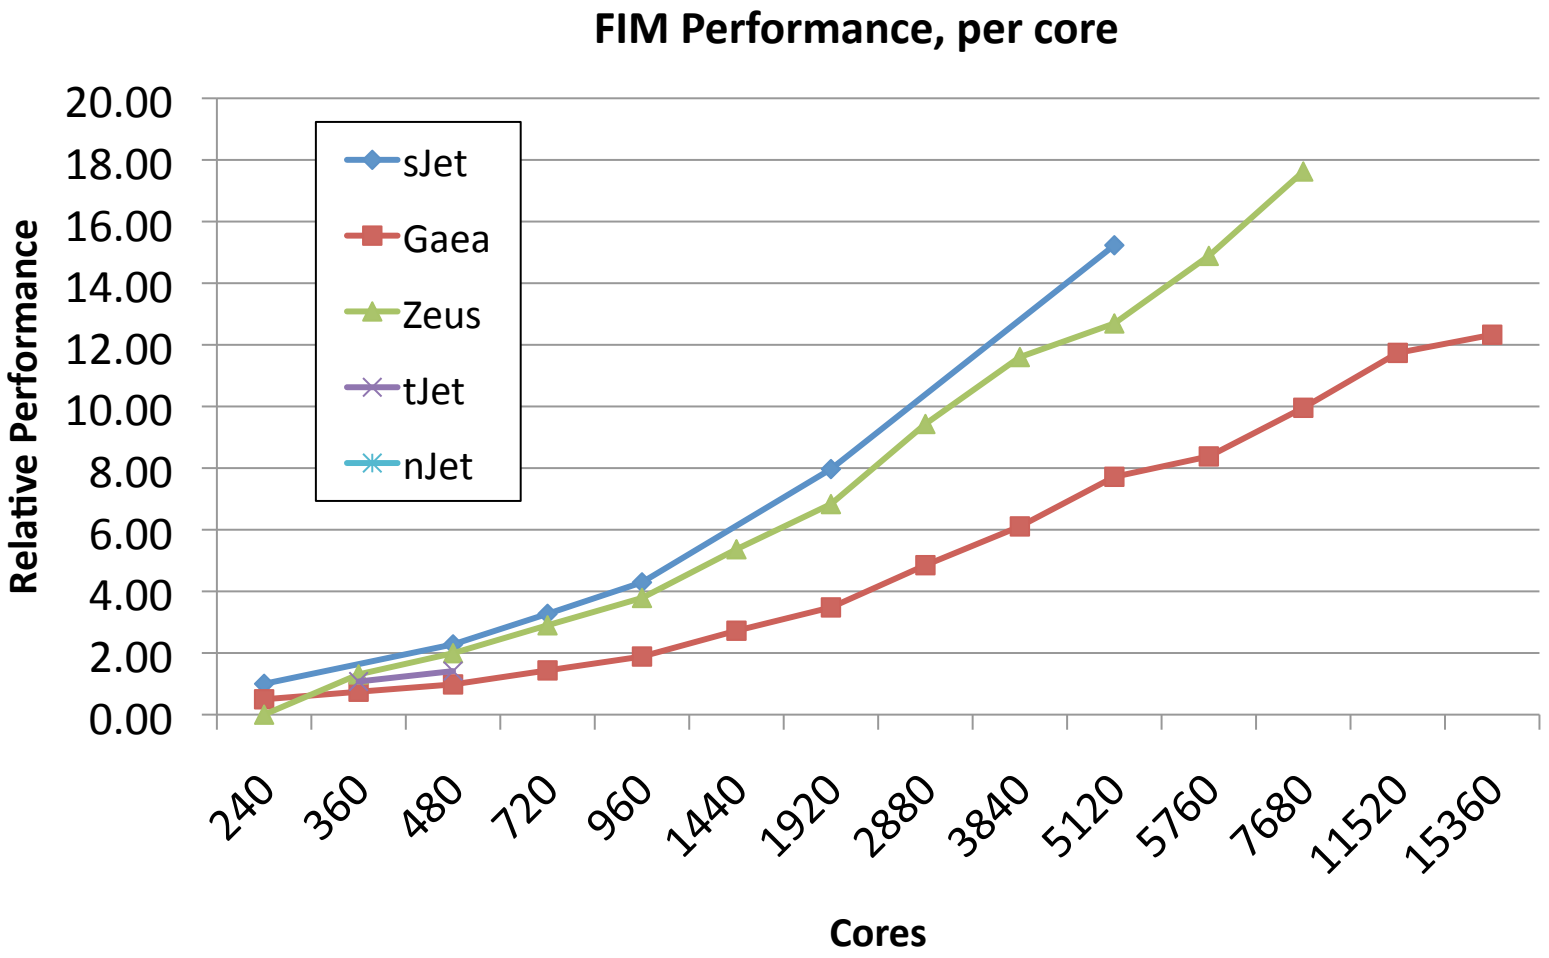

# Standing Reservations on tJet

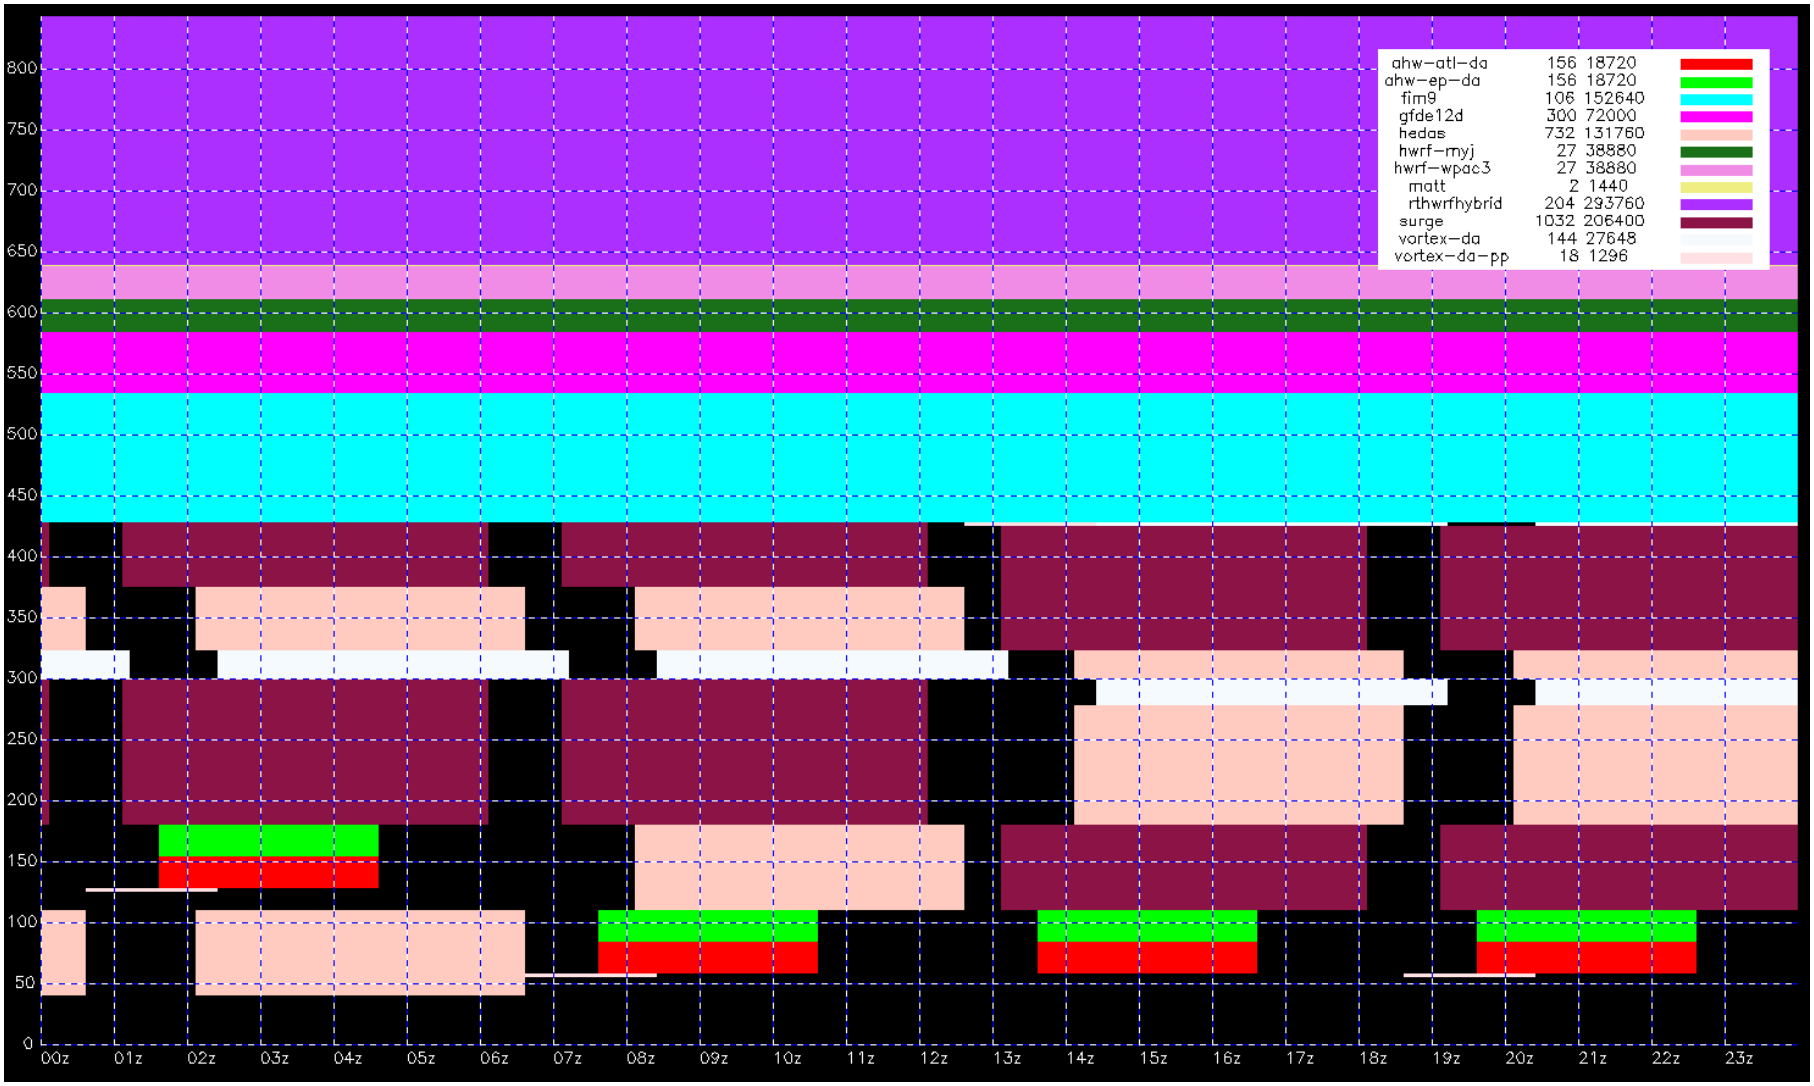

# FIM Performance Results

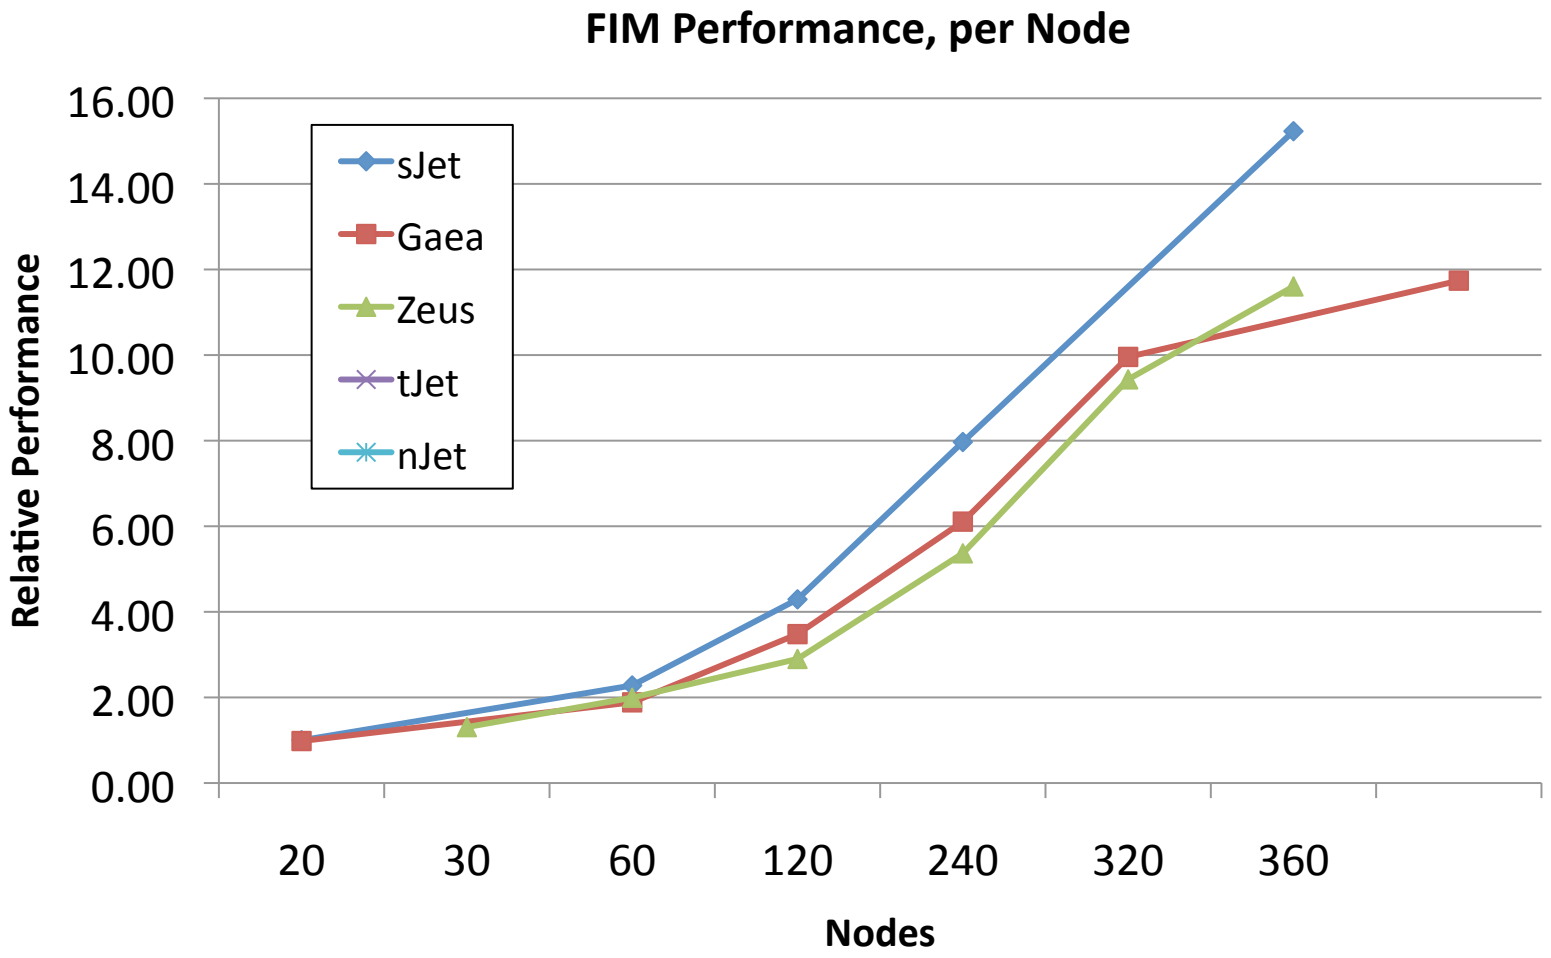

# Summary

- Each architecture has their strengths
- Scalability can vary widely by application and interconnect
  - We have some things to look into!
- Heterogeneity (technology and operations) is a good thing, as one system is not best for all