



# The Path to Operations

Rebecca Cosgrove  
Implementation and Data Services Branch  
NCEP Central Operations

January 19, 2016



# Outline



- Implementation Process
- Computing Platforms
- Data Availability
- Dissemination Options





---

# Implementation Process



# Implementation Process

## Stage 1: Development

- Get necessary approvals for your project first
  - requirements-based
- Come talk to my team early in the lifecycle of your project!
  - May be able to help run some of your testing or retrospective runs in the “white space”
  - Review software with SPA team
  - Discuss data needs and dissemination strategy
  - Review NCO Implementation Standards
    - <http://www.nco.ncep.noaa.gov/pmb/docs/Implementation%20Standards%20v10.0.pdf>
- Scientific review and approval prior to code handoff





# Implementation Process

## Stage 2: Transition to Operations

- Submit project quad chart and resource usage to IDSB Chief one year in advance to be included in planning
  - For new applications, or those greatly expanded in scope, will need approval of the High Performance Computing Resource Allocation Council (HPCRAC)
- “EE (Environmental Equivalence) Coordination Meeting”
  - Meet with IDSB teams to discuss your application and all data-related items
  - Scheduled roughly 3 months before you submit your final code
- Once final code has been delivered:
  - IDSB conducts IT testing
    - Confirm it conforms to NCO standards
    - Run through test cases
    - Try to stress or break the application
    - Try to tune the application to maximize resources
  - IDSB conducts 30-day stability test
    - No changes to application during this period
    - Output is provided to user base for evaluation
- Conduct NCEP Director briefing to obtain approval to implement
- Development Organization must commit to providing tier 3 support



---

# Computing Platforms



# NOAA R&D HPCS Overview



## Development HPC

### Systems Integration Contract (CSC)

- May 2010-May 2019 / \$317M / IDIQ
- 9 yrs with 4-yr base, 4-yr option, 1-yr transition

### Systems Configuration

#### Theia - Fairmont, WV (Zeus Replacement)

- Short-term/seasonal/inter-annual predictions
- 1,024 teraflops – Cray
- 1 petaflop fine-grained system to be installed Q3

#### Jet - Boulder, CO (NOAA Skaggs Facility)

- Hurricane forecast improvement
- 722 teraflops - Aspen & Cray

#### Princeton, NJ (NOAA/GFDL)

- Climate post-processing and analysis
- 106 nodes (8 core Intel Xeon) – Dell

### Performance Measures

- Minimum 96.0% System Availability
- Minimum 99.0% Data Availability

## Research HPC

### Interagency Agreement (DOE/ORNL)

- Aug 2009-Aug 2016 / \$108M / Cost Reimbursable
- 5 year agreement extended 2 years

### Systems Configuration

#### Gaea - Oak Ridge, TN (Oak Ridge National Lab)

- Climate change research and projections
- 1,100 teraflops Cray (C1 + C2)
- C1 decommissioned April, C2 decommissioned February
- 1,400 teraflops Cray (C3)

#### Titan - Oak Ridge, TN (Oak Ridge National Lab)

- Applications for next generation architectures
- 500 teraflops allocation of 27,000 teraflops Cray using Nvidia Graphics Processing Units

### Performance Measures

- Minimum 96.0% System Availability
- Minimum 99.0% Data Availability



# R&D HPCS Allocation for Theia



Theia allocations assigned by Allocation Committee

## Sandy Supplemental Theia Allocation


Portfolio	Percent	Avg. Monthly Allocation	Portfolio Managers	Backup
AOML	25	4,339,690	John Cortinas Andrew Ostapenko	Thiago Quirino and Molly Baringner
EMC	29	5,034,040	Kevin Kelleher	Mary Hart
ESRL	29	5,034,040	Brian Gross	Forrest Hobbs
GFDL	9	1,562,288	Sid Boukabara	Frank Indiviglio
NESDIS	5	867,938	Jesse Feyen	Krishna Kumar
NOS	3	520,764		Mary Erickson
<b>Sum</b>	<b>100</b>	<b>17,358,760</b>		





# WCOSS

## Weather and Climate Operational Supercomputing System



- ★ Tide/Luna – Reston, VA
- ★ Gyre/Surge – Orlando, FL
- combined 2.8 petaflop system with connected IBM and Cray components

One system designated as the “production” machine, the other the “development” machine

- Open access to development machine
- Production machine access restricted



# WCOSS

## Weather and Climate Supercomputing System

Tide/Gyre

Luna/Surge

WCOSS Components	Task Order 002 IBM Phase I July 2013	Task Order 002 IBM Phase 2 Jan 2015	Task Order 003 IBM July 2013	Task Order 004 (Cray) Nov 2015	Totals
Compute NODES	440	1,080	180	2,048	3,748
Compute Racks	9	15		12	
Spare Nodes	20	18		30	68
Peak TFs	167	599	64	2,060	2,890
Cores (Compute and Service)	7,040	25,920	2,880	50,176	86,016
Spare Cores	320	432		768	1,520
Processor Type	Intel Sandy Bridge	Intel Ivy Bridge	Intel Sandy Bridge	Haswell & Sandy Bridge	
Processor Clock Speed	2.6 Ghz	2.7 Ghz	2.6 Ghz	2.6 Ghz	
Cores/node	16	24	16	24	
Service Nodes	54	58		100	212
Memory/core	2 GiB	2.66 GiB	2 GiB	2 GiB	
Disk Storage (useable)	1.152 PB	2.034PB	1.438 PB	3.5 PB	8.124 PB
Shared Storage	259 TB	266TB		300 TB	825 TB
Backup Tape Capacity	600 TB				600 TB
Interconnect Fabric	Mellanox FDR	Mellanox FDR	Mellanox FDR	Mellanox FDR	
Operating System	Red Hat Linux	Red Hat Linux	Red Hat Linux	CRAY and SUSE Linux	
Filesystem	GPFS	GPFS	GPFS	GPFS	
Scheduler	LSF	LSF	LSF	LSF	
Workflow Scheduler	ecFLOW	ecFLOW	ecFLOW	ecFLOW	



---

# Data Availability



# HPSS

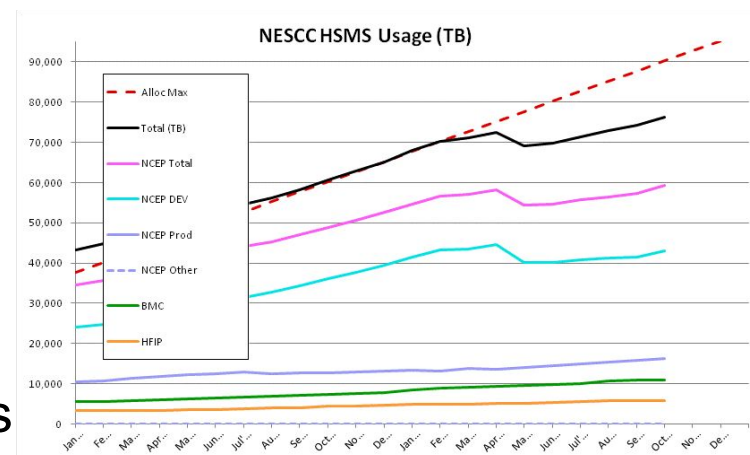
## High Performance Storage System

- Tape storage system located in Fairmont, WV with connectivity to all R&D machines and WCOSS
- Space is allocated by project, but we are nearing capacity
- Developers can save any data they choose within their quota
- NCO sends input data and operational output to HPSS as the “runhistories”

- Input data needed to rerun models kept forever, and select output for 1 to 2 years
- Catalog of datasets:

<http://www2.nco.ncep.noaa.gov/pmb/docs/runhist.html>

- Note that what is saved is being revisited due to HPSS being at capacity





# Data on WCOSS

- Input Data (resides in /dcom for up to 10 days):
  - Real-time satellite data
  - Real-time observational data in “tanks” in BUFR format
    - NCEP/EMC submits code to NCO that controls the content of the tanks
  - Other US and international modeling centers’ model output (some restrictions apply)
    - FNMOC, Air Force, ECMWF, CMC, UKMET
  - Data from the IDP system – MRMS, MADIS
  - Radar and precipitation amount data
- NCEP Production Suite model output data (resides in “/com” directories for 2 to 60 days, depending on model)
- Operationally-critical data is mirrored to development machine, and a subset sent to Theia



---

# Dissemination Options



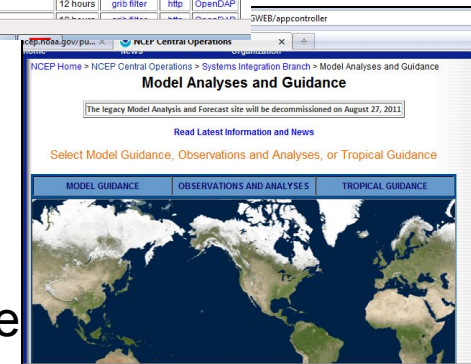
# Outgoing Data

- Major data destinations:
  - **NWSTG**
    - NWS FTP server, aka “tgftp”, which includes National Digital Guidance Database (NDGD)
    - AWIPS2 and NOAAPORT
    - International Partners via GTS
    - Direct Dissemination to Partners
    - Pathways to NCEI
  - **NCO Dissemination Systems**
    - NCEP FTP server, aka “ftpprd”
    - NOMADS, providing “slice and dice” capability
    - CONDUIT – dissemination to university community
    - MAG – imagery from WCOSS applications
  - **NCEP Centers/NAWIPS/AWIPS2**
  - **Integrated Dissemination Program (IDP) System**
    - nowCOAST and other GIS systems
- The majority of our data is made available to the public for free
  - Exception is restricted data that can not be widely distributed

NCEP Central Operations

NOAA Operational Model Archive and Distribution System

Data Set	freq	grib filter	http	gds
<b>Global Models</b>				
FNL	6 hours	grib filter	http	OpenDAP
GFS 1.0x1.0 Degree	6 hours	grib filter	http	OpenDAP
GFS 0.5x0.5 Degree	6 hours	grib filter	http	OpenDAP
GFS 2.5x2.5 Degree	12 hours	grib filter	http	OpenDAP
GFS Ensemble high resolution	6 hours	grib filter	http	OpenDAP
GFS Ensemble high-resolution Bias-Corrected	6 hours	grib filter	http	OpenDAP
GFS Ensemble NDGD resolution Bias-Corrected	6 hours	grib filter	http	OpenDAP
GFS Ensemble Precip Bias-Corrected	daily	grib filter	http	OpenDAP
NAEFS high resolution Bias-Corrected	6 hours	grib filter	http	OpenDAP
NAEFS NDGD resolution Bias-Corrected	6 hours	grib filter	http	OpenDAP
<b>Regional Models</b>				
AGM	06Z_12Z	grib filter	http	OpenDAP
HIRES Alaska	daily	grib filter	http	OpenDAP
HIRES East CONUS	12 hours	grib filter	http	OpenDAP
HIRES Guam	12 hours	grib filter	http	OpenDAP





# Issues to Consider

- HPSS tape system is nearing capacity
- R&D network is full
- Disk space on WCROSS has been allocated
- Making data available to users may require large amounts of disk space on ftp servers
- All facets of the IT infrastructure must be considered when planning/funding your projects!





---

# Questions?


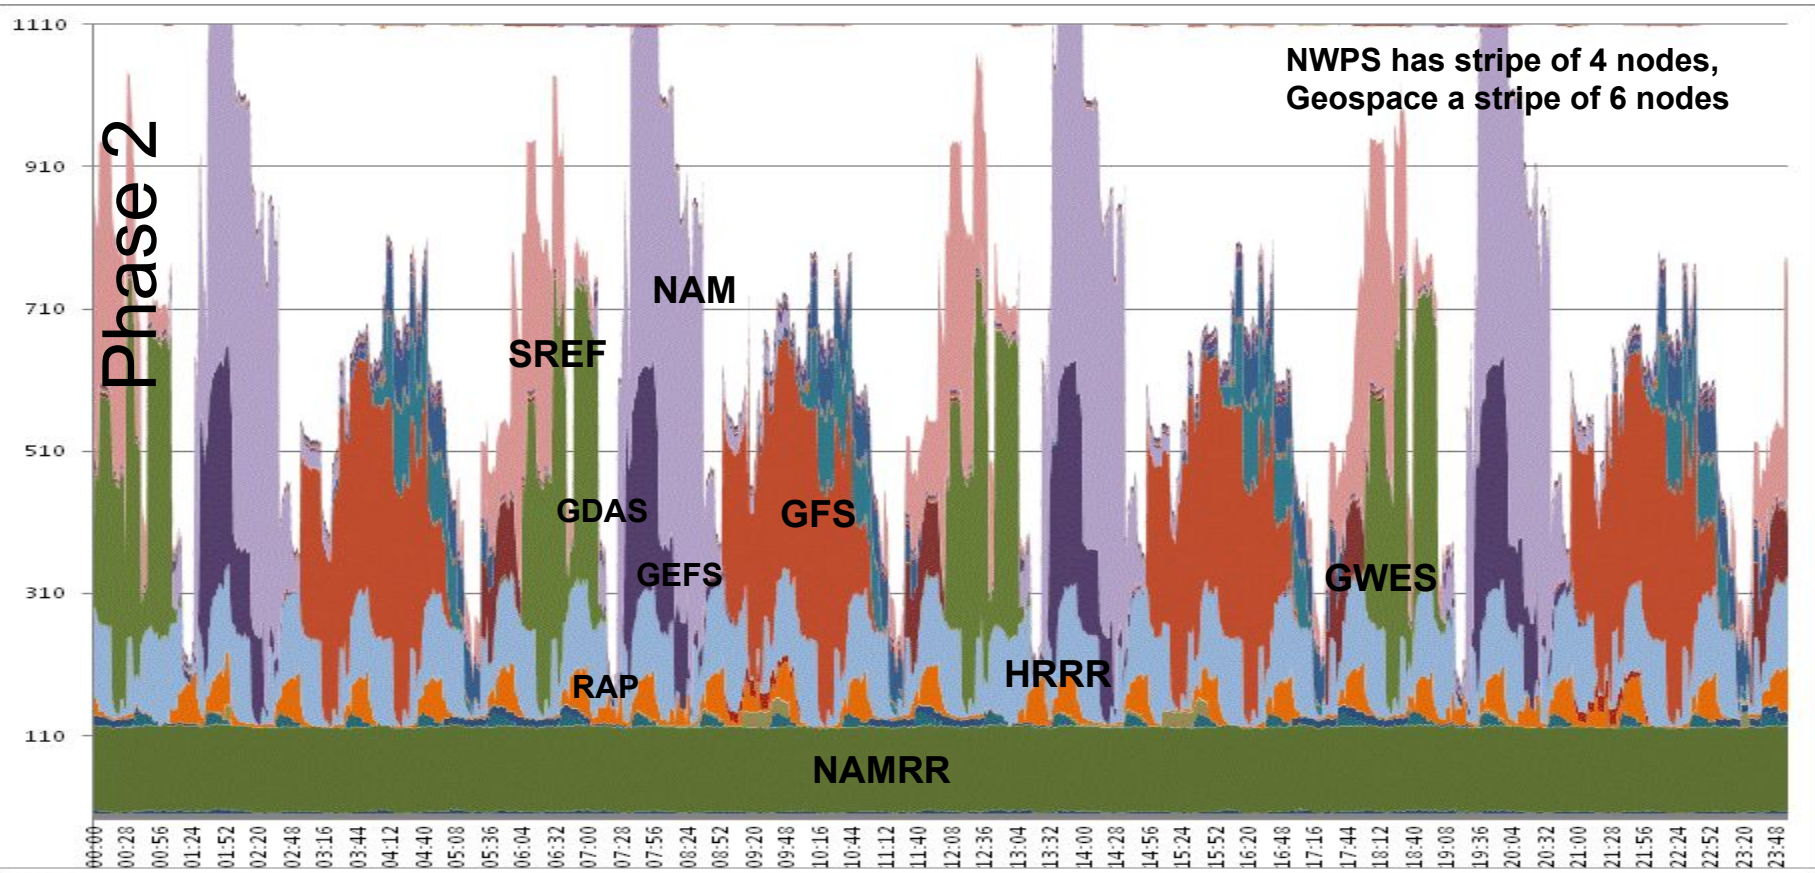
[rebecca.cosgrove@noaa.gov](mailto:rebecca.cosgrove@noaa.gov)



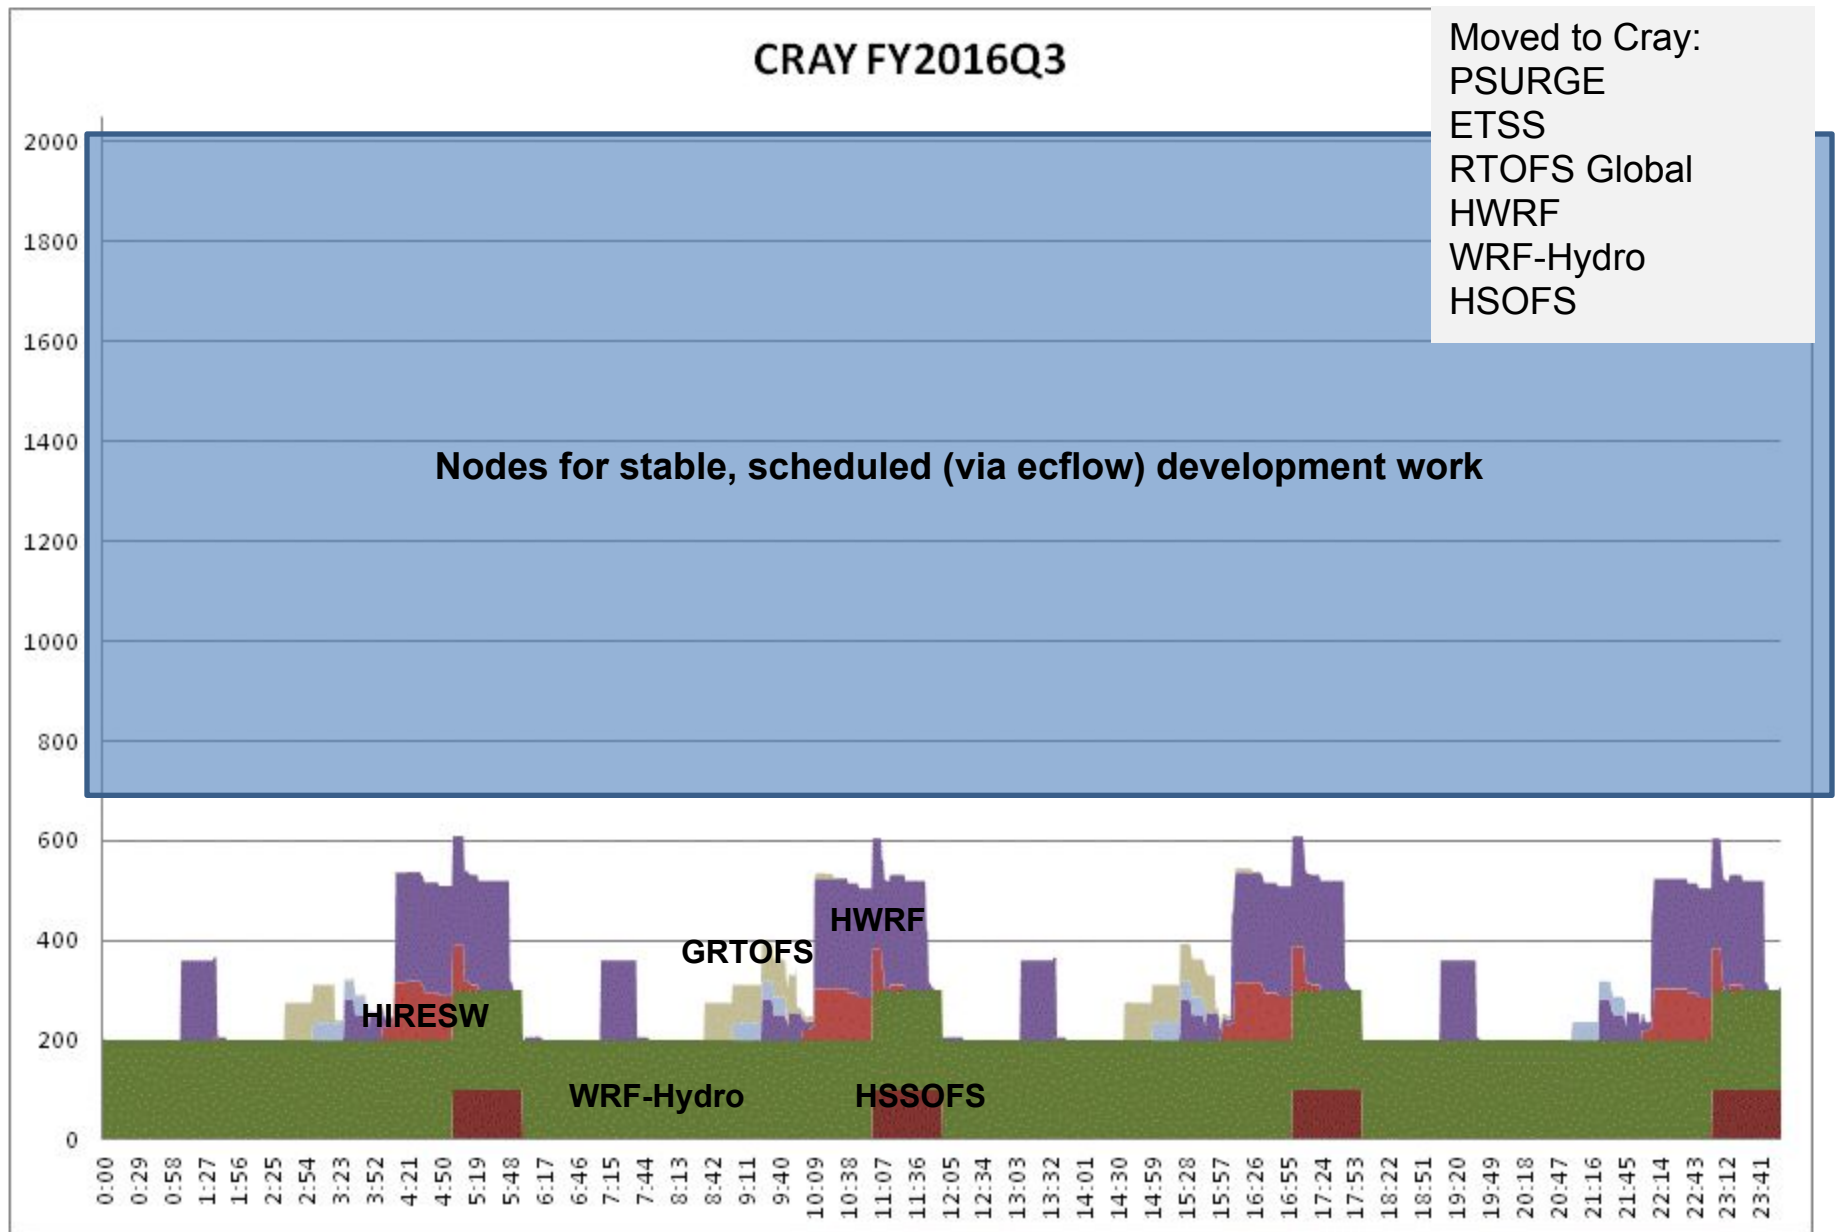
---

# Backup Slides

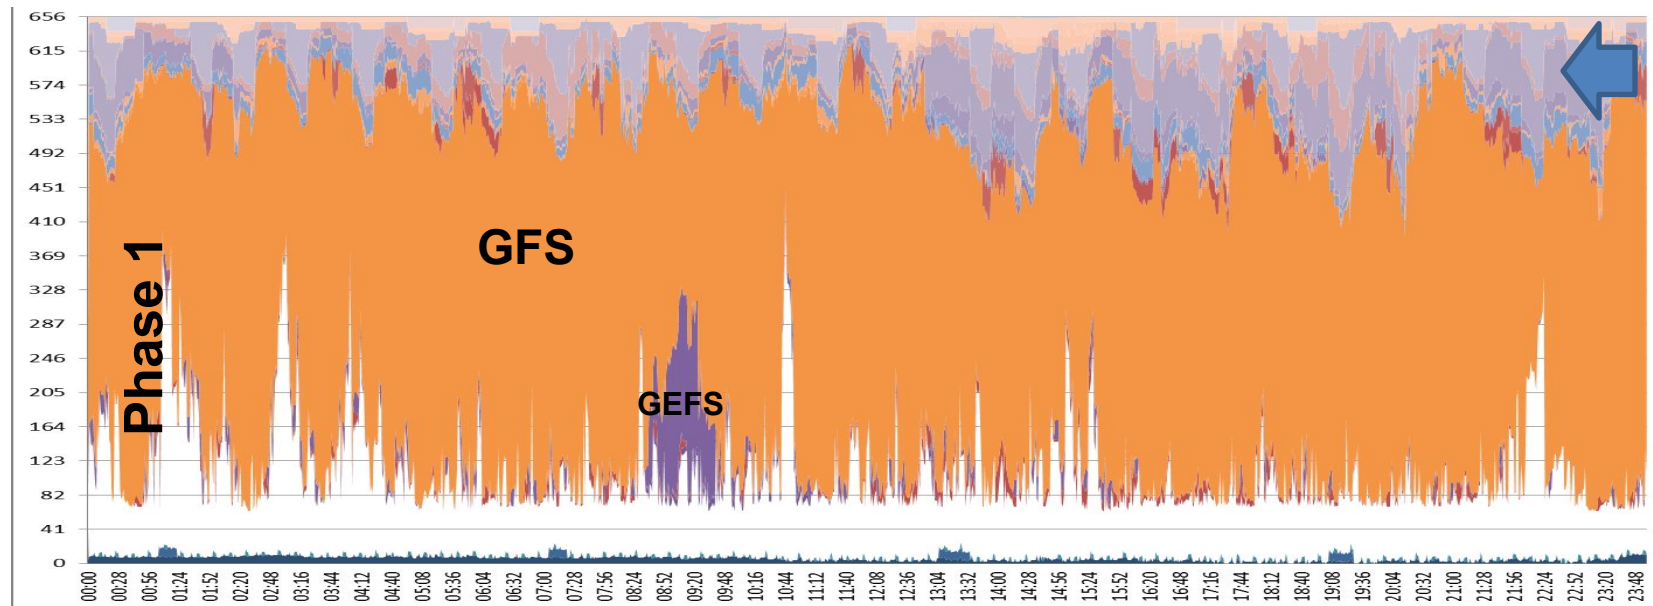

# Projected Operational IBM Node Usage by end of FY16Q3



# Projected Operational Cray Node Usage by end of FY16Q3



# Development System - 12/9/2015



- Geospace
- RAP
- NOS models
- NWPS
- etc