NCEP HPC Transition

15th ECMWF Workshop on the Use of HPC in Meteorology

Allan Darling

Deputy Director, NCEP Central Operations

WCOSS – NOAA Weather and Climate Operational Supercomputing System

CURRENT OPERATIONAL CHALLENGE

NCEP Operational Forecast Skill

36 and 72 Hour Forecasts @ 500 MB over North America
[100 * (1-S1/70) Method]

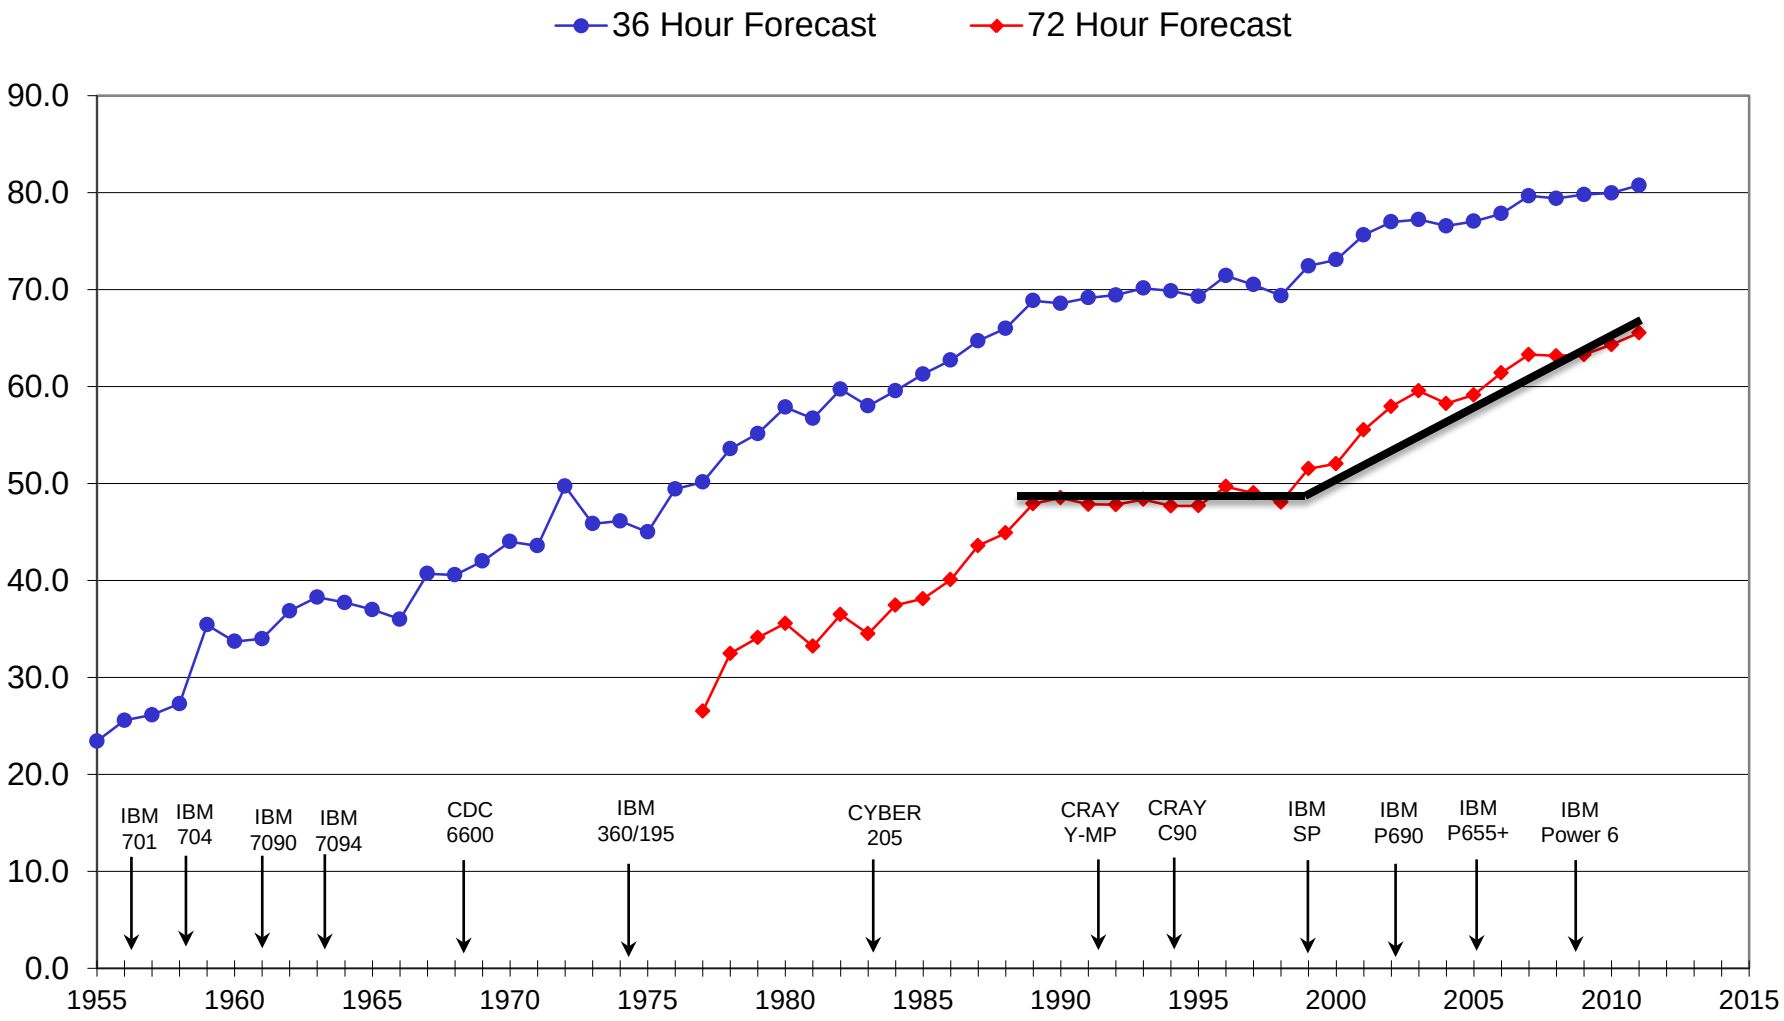

NCEP HPC Growth

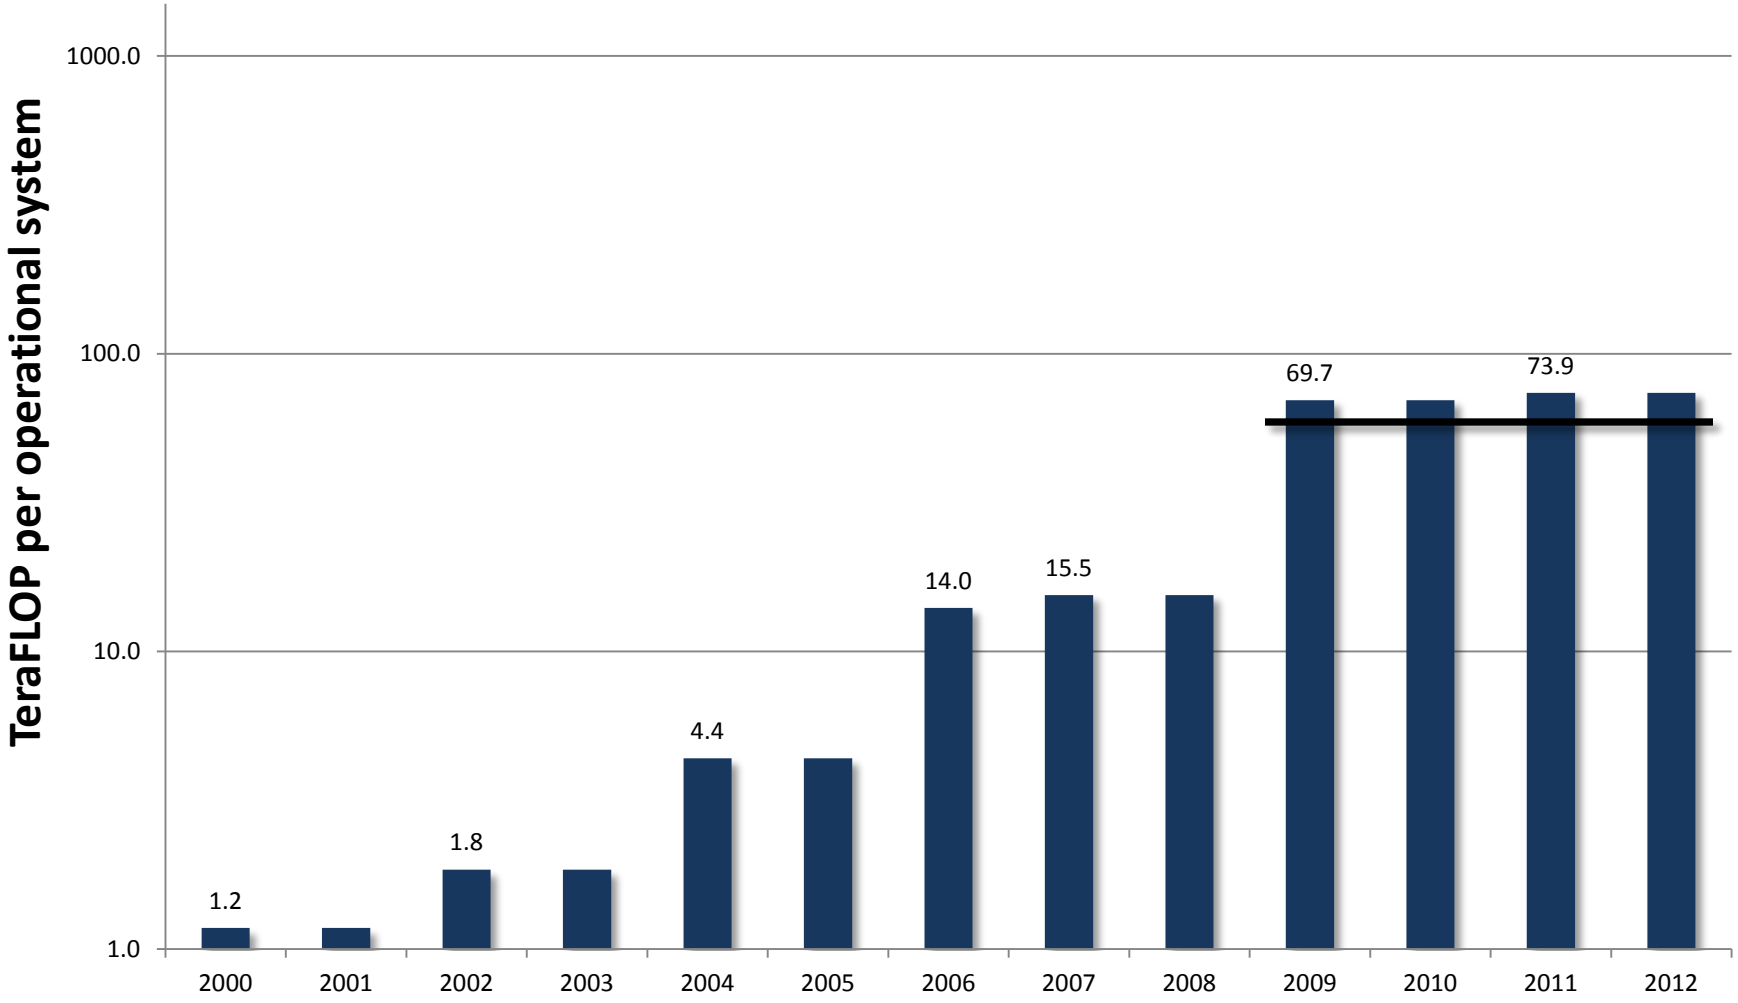

HPC Production Utilization

February 2010

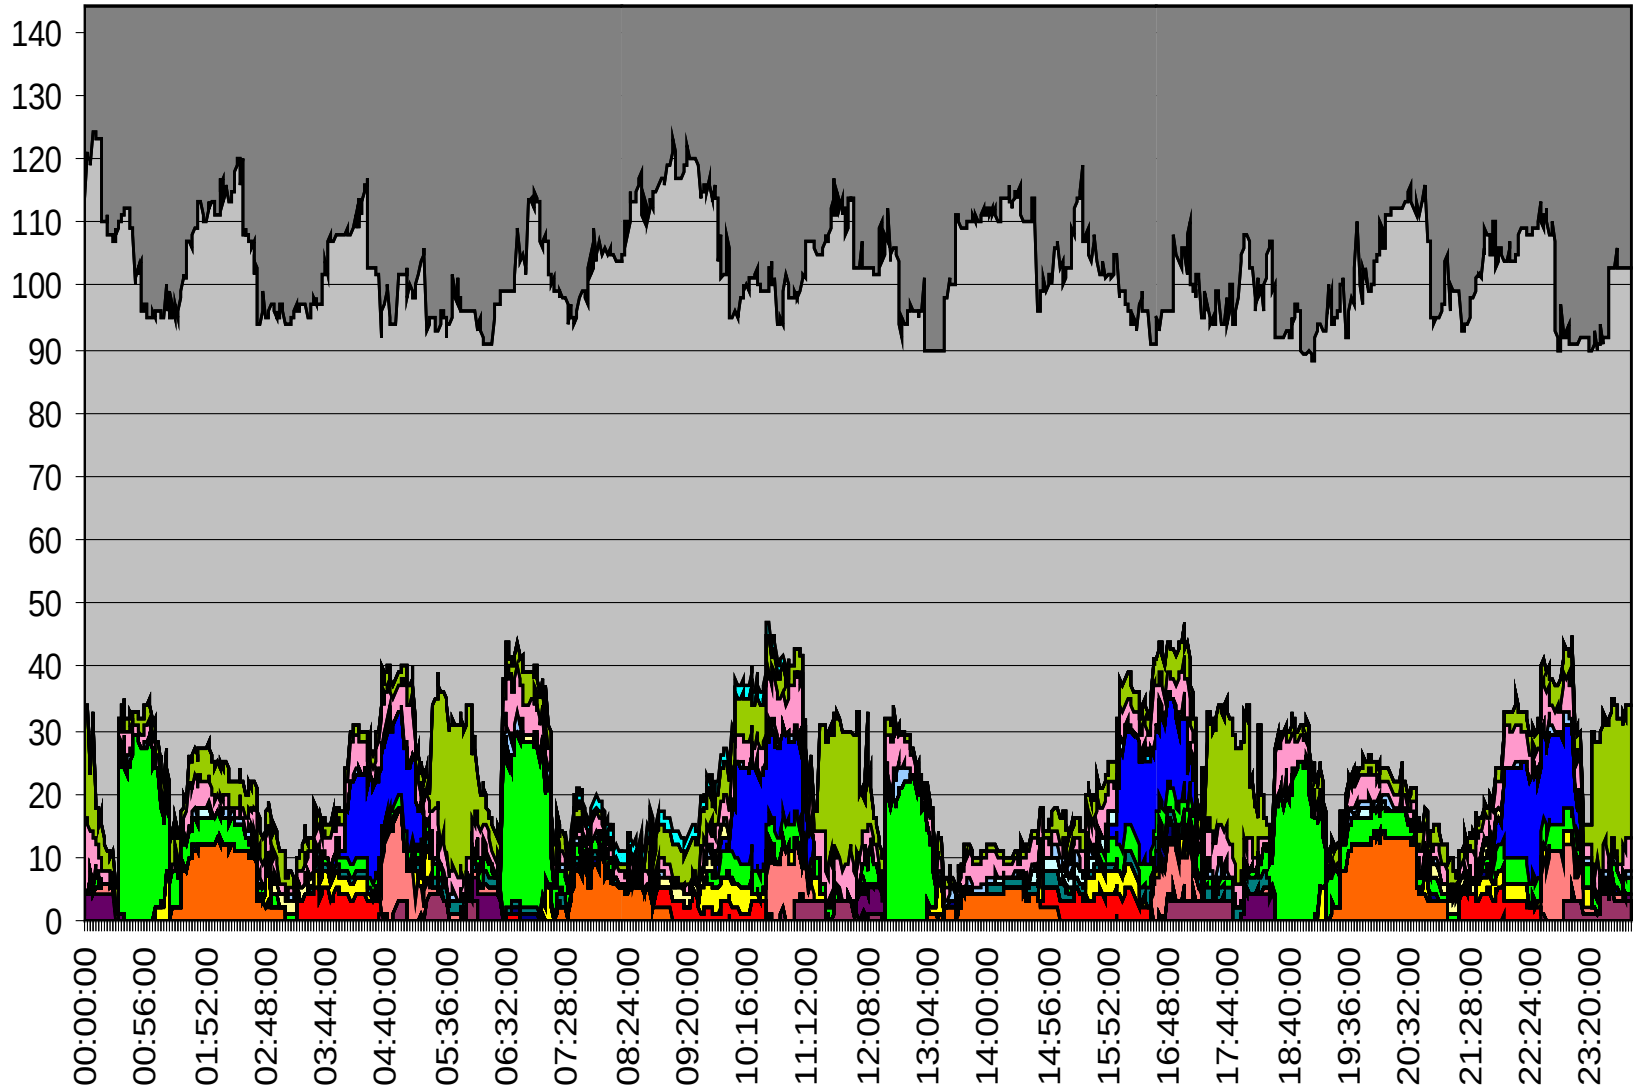

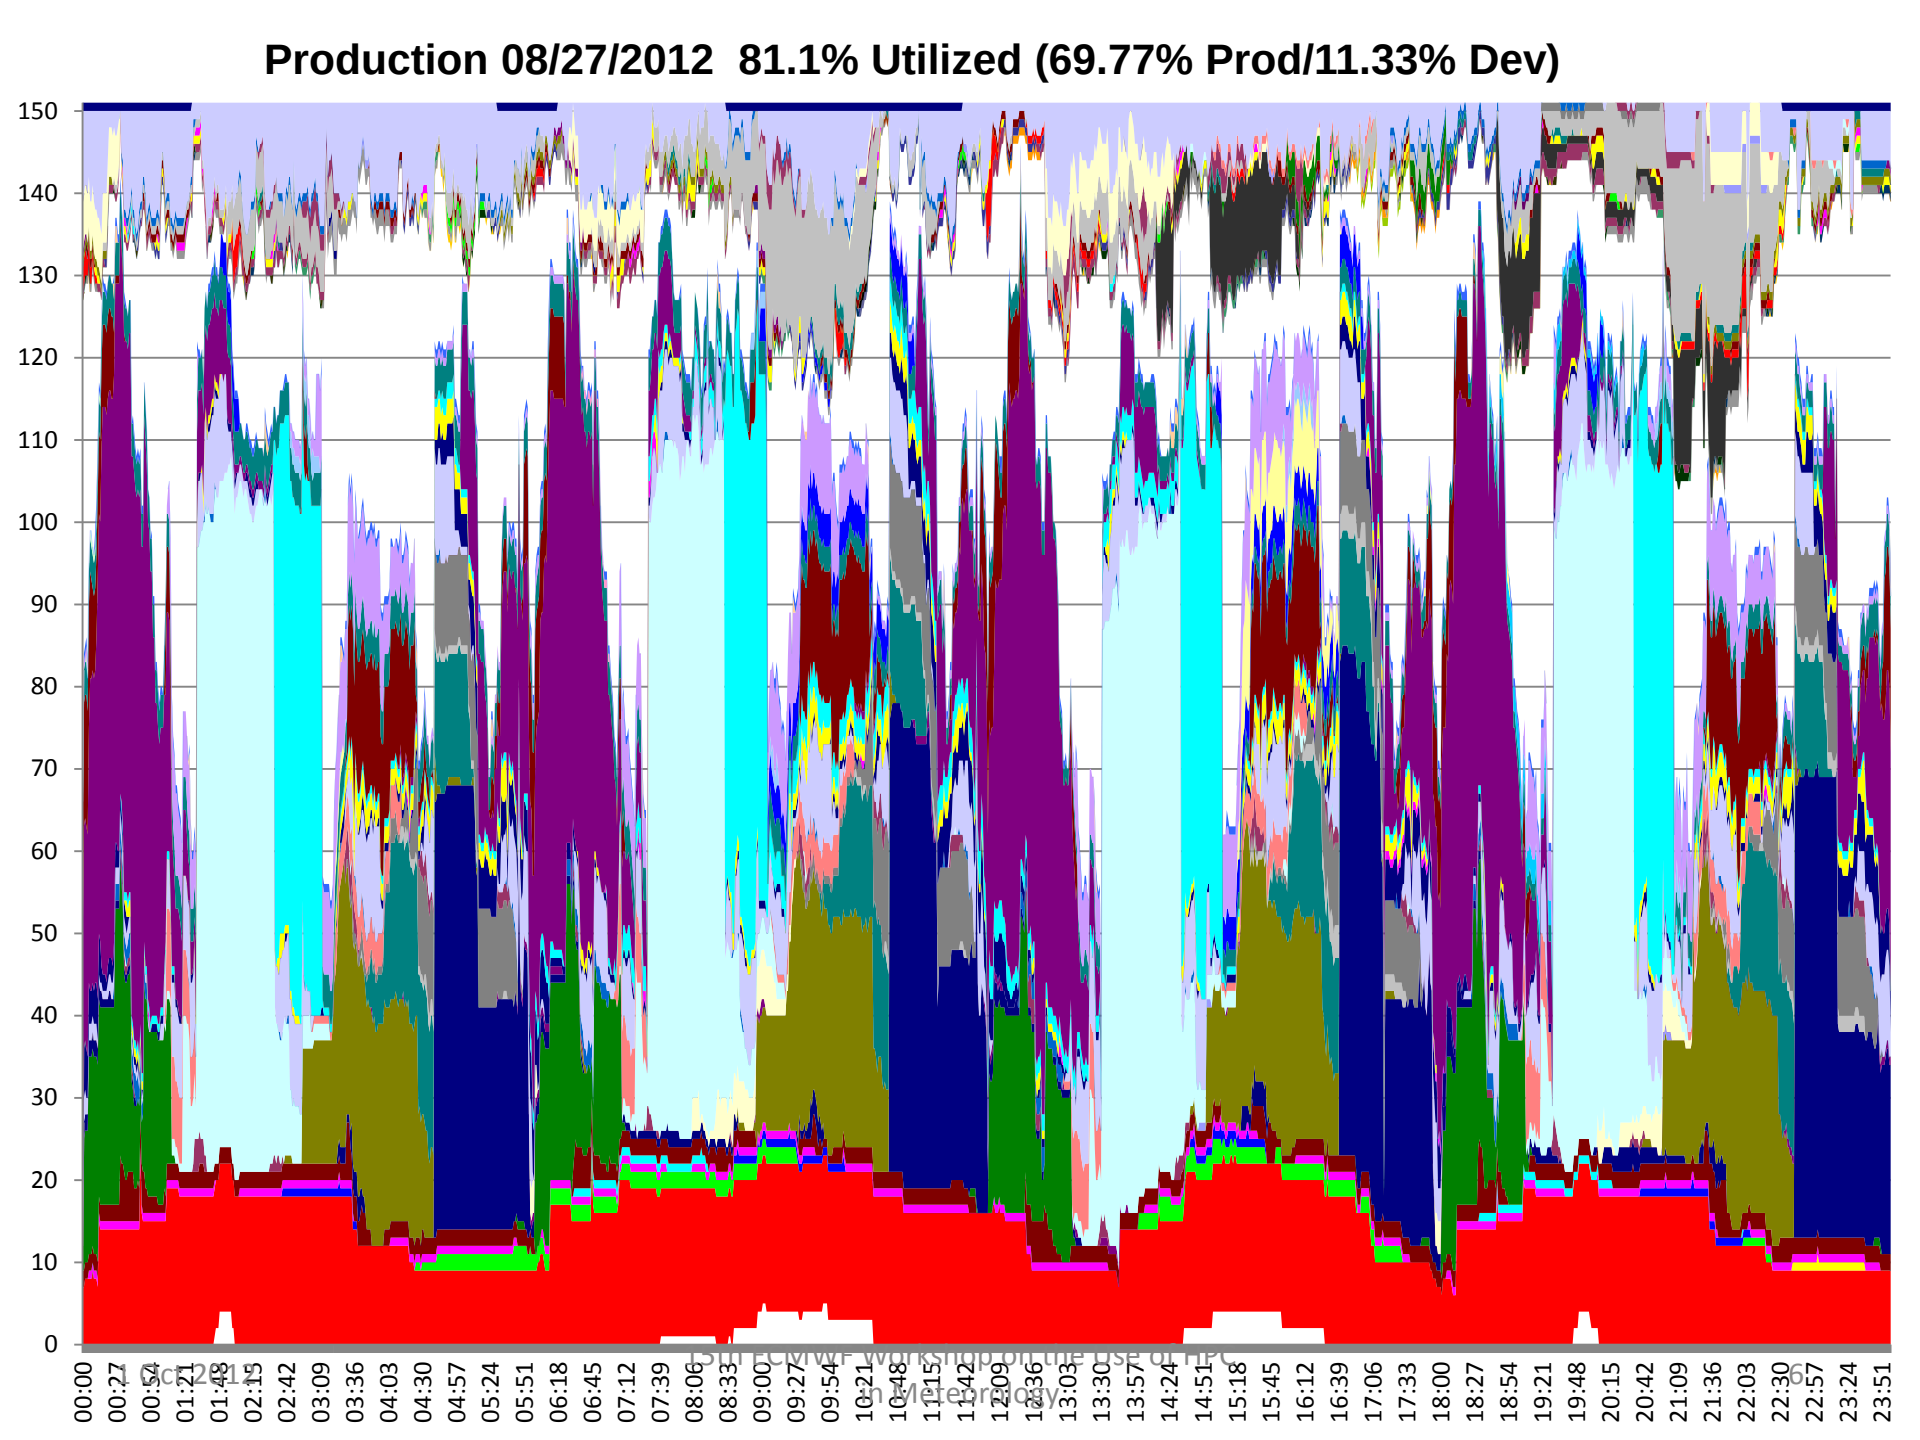
Production 08/27/2012 81.1% Utilized (69.77% Prod/11.33% Dev)
150
140
130
120
110
100
90
80
70
60
50
40
30
20
10
0
00:00
00:27
00:54
01:21
01:48
02:15
02:42
03:09
03:36
04:03
04:30
04:57
05:24
05:51
06:18
06:45
07:12
07:39
08:06
08:33
09:00
09:27
09:54
10:21
10:48
11:15
11:42
12:09
12:36
13:03
13:30
13:57
14:24
14:51
15:18
15:45
16:12
16:39
17:06
17:33
18:00
18:27
18:54
19:21
19:48
20:15
20:42
21:09
21:36
22:03
22:30
22:57
23:24
23:51

Product Generation Summary

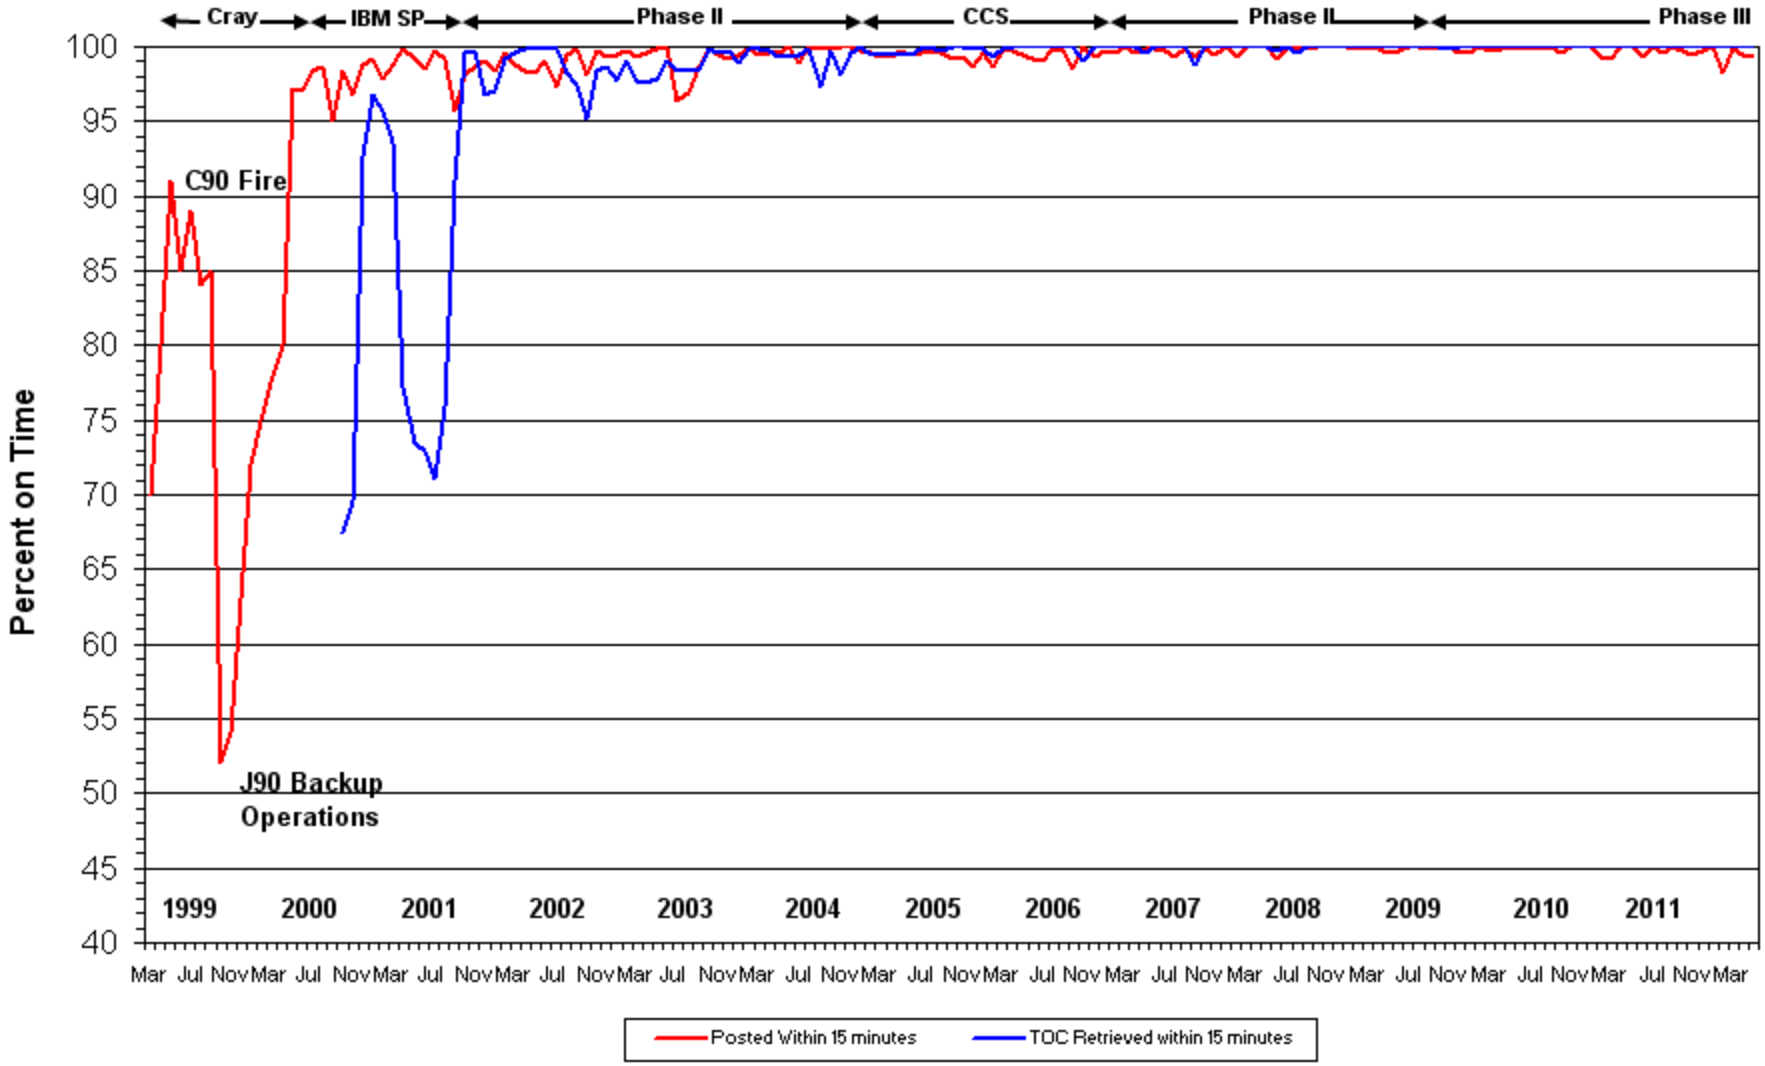

May 2012

FY11 On-Time Product Generation

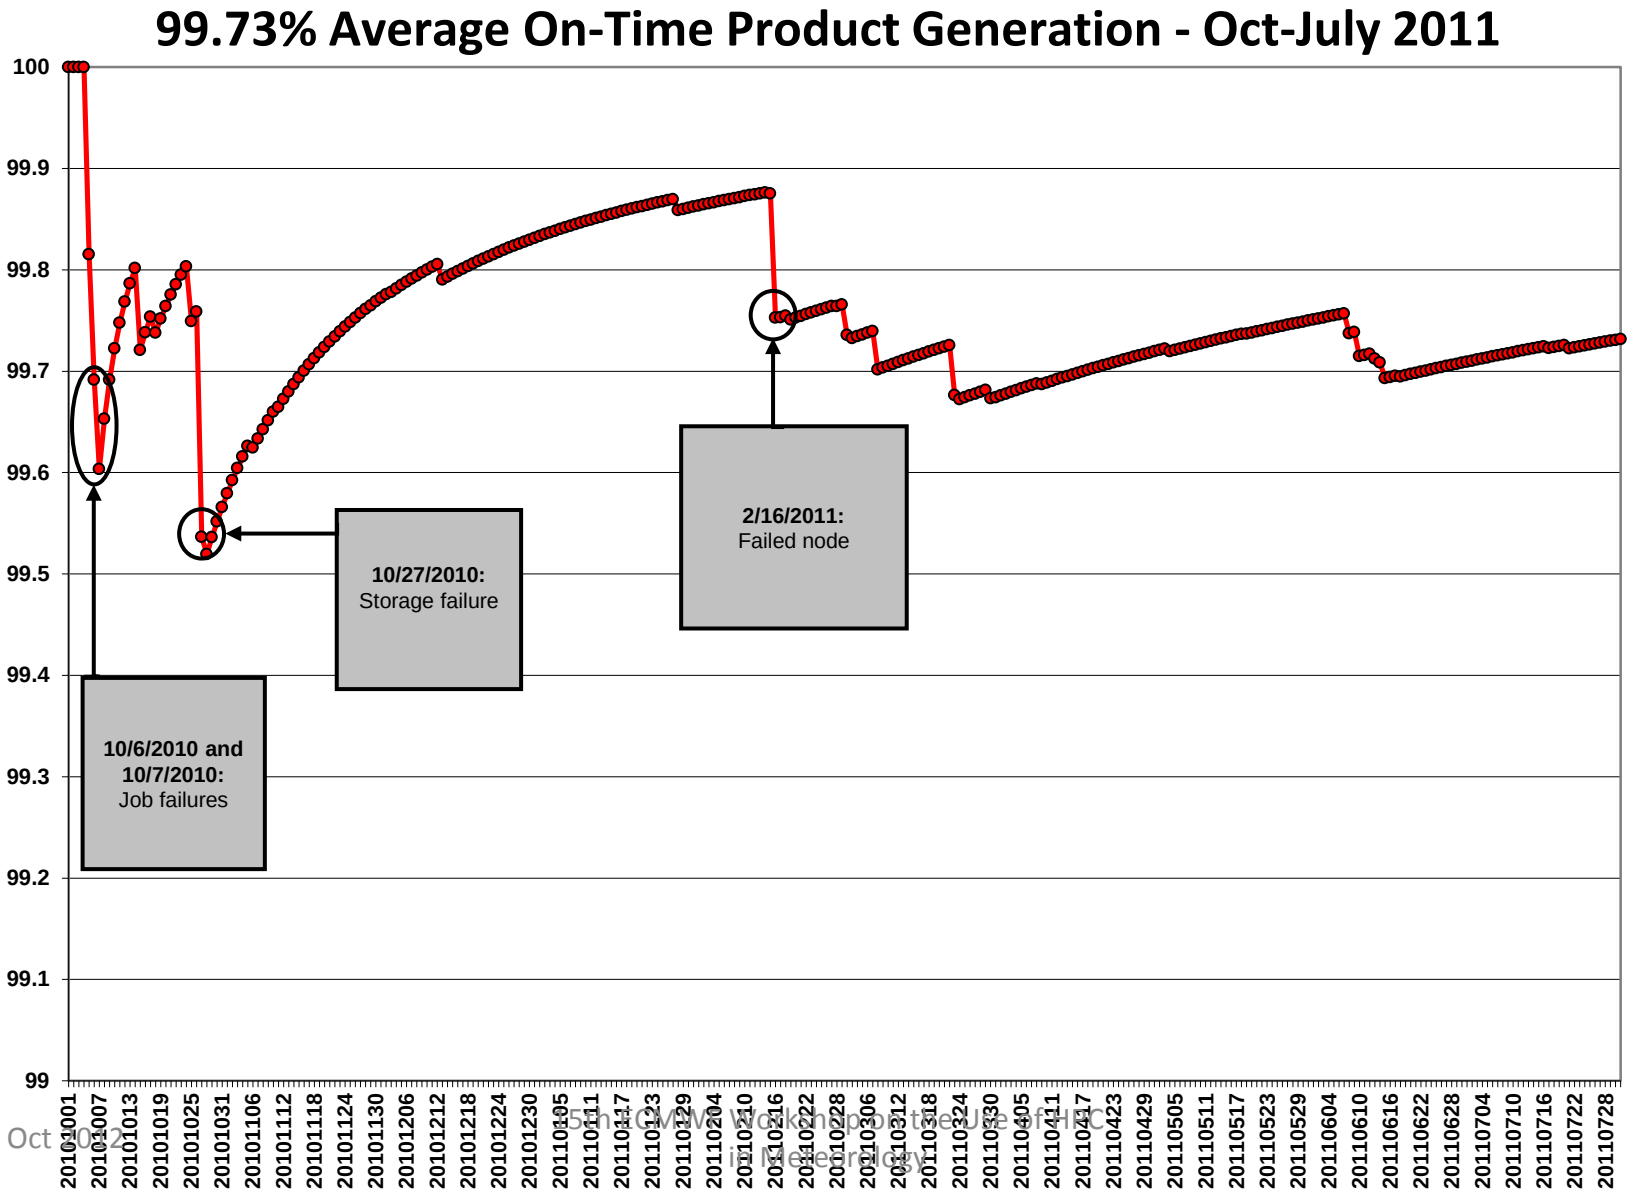

FY12 On-Time Product Generation

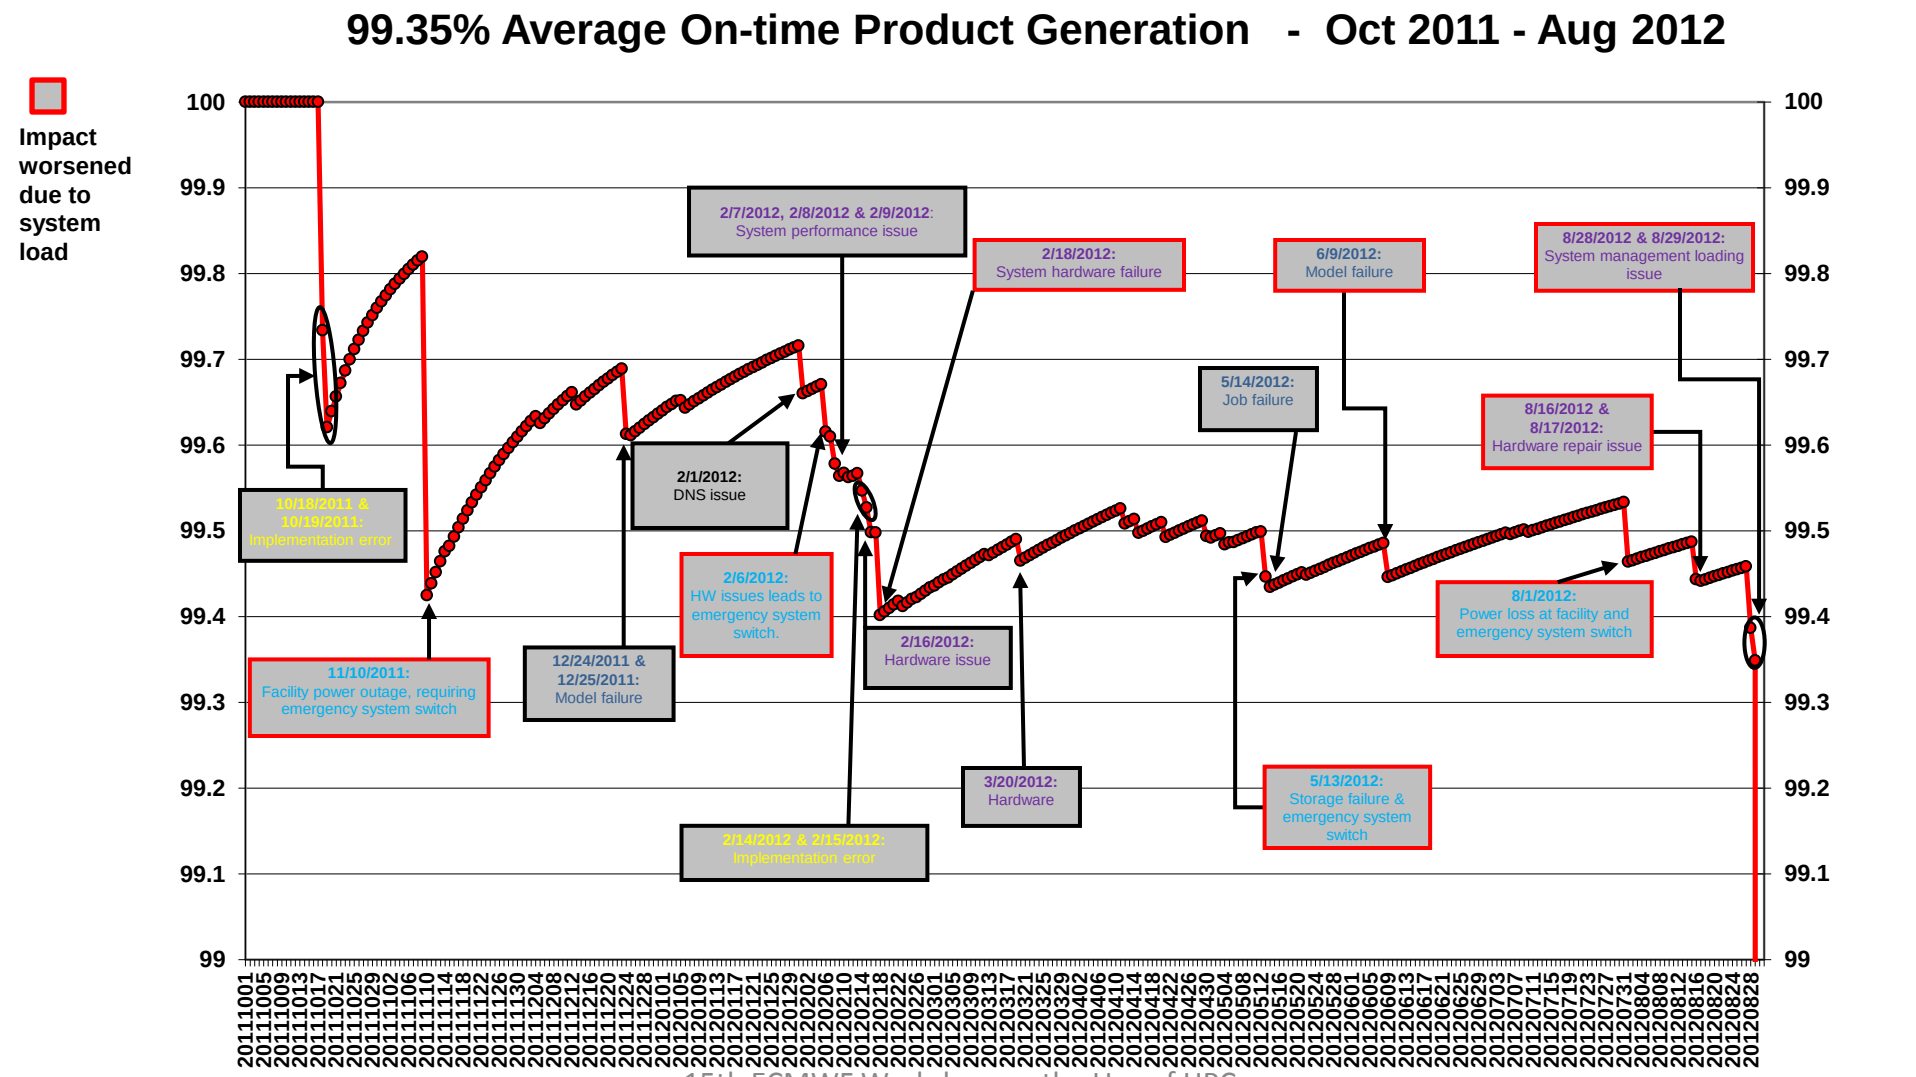

WCOSS – NOAA Weather and Climate Operational Supercomputing System

CONTRACT & SYSTEM OVERVIEW

WCOSS Contract Overview

- Sustains current operational systems during build, acceptance and transition to new systems
  - Bridge contract expires September 2013
- One 5-year base period, one 3-year option period, one 2-year transition option period
- Provides phased delivery of new Primary and Backup systems in the 5-year base period
  - Will transition from IBM/Power6/AIX to IBM/Intel/Linux by end of FY13
- Provides new sites
  - Primary Reston, VA
  - Backup Orlando, FL

Bridge Contract

Current NOAA Weather and Climate Operational Supercomputer System

Location

- Primary
  - Gaithersburg, MD (IBM provided facility)
- Backup
  - Fairmont, WV (GFE NASA IV&V facility)

Configuration

- Identical Systems (per site)
  - IBM Power 6/P575/AIX
  - 73.9 trillion calculations/sec
  - 5,314 processing cores
  - 0.8 petabytes of storage
- Performance Requirements
  - Minimum 99.0% Operational Use Time
  - Minimum 99.0% On-time Product Generation
  - Minimum 99.0% Development Use Time
  - Minimum 99.0% System Availability
  - Failover tested regularly

Inputs and Outputs

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

Significance

- Where United States 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
  - Air Quality
  - Fire Weather
- Contract ensures no gap in operations

WCOSS Contract

Future NOAA Weather and Climate Operational Supercomputer System

Location

- Primary
  - Reston, VA (IBM provided facility)
- Backup
  - Orlando, FL (IBM provided facility)

Configuration

- Identical Systems (per site)
  - IBM iDataPlex/Intel Sandy Bridge/Linux
  - 208 trillion calculations/sec
  - 10,048 processing cores
  - 2.59 petabytes of storage
- Performance Requirements
  - Minimum 99.9% Operational Use Time
  - Minimum 99.0% On-time Product Generation
  - Minimum 99.0% Development Use Time
  - Minimum 99.0% System Availability
  - Failover tested regularly

Inputs and Outputs

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

Significance

- Where United States 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
  - Air Quality
  - Fire Weather

Key Performance Metrics

- Bridge
  - At least 99.0% operational use time
- WCOSS
  - At least 99.9% operational use time
- WCOSS and Bridge
  - At least 99.0% system availability
  - Sustain on-time generation of model guidance products at a rate of 99% or better; within 15 minutes of target completion times
  - No degradation in product delivery to occur during planned, routine failover testing between the Primary and Backup WCOSS
  - Delivery of committed computing capability upgrades objectively measured and accepted through the use of benchmark suite

WCOSS Locations

NCEP will insert the new Reston and Orlando facilities into NCEP's high speed network

Reston WCOSS Primary Facility

- SAS 70 Type II certified
- 2N and N+1 AC UPS power, DC power available
- N+1 cooling systems
- 24x7 staffing and site access, CCTV surveillance, hardened data center perimeter, and biometric access control

Orlando WCOSS Backup Facility

- SAS 70 Type II Certified
- N+1 or 2N UPS power
- 20-ton CRAH units delivering 2,750 tons of chiller capacity
- 24x7 on-site security staffing, CCTV surveillance, 8" thick, steel-reinforced concrete walls, and biometric access

Alternate Process Site (APS) space and power included at the Reston facility

System Characteristics

(per site)

	Life Cycle	Architecture	OS	Average Capability	Average Capacity	Cores	TeraFLOP
Bridge System	Oct 2012 - Sep 2013	Power6	AIX 5.3	1.0X	1.0X	5,314	73.9 TF
WCOSS Phase 1	Planned accept Dec 2012 – FOC Aug 2013	iDataPlex	Linux (RHEL)	*2.0X over Bridge P6	*2.3X over Bridge P6	10,048	208 TF
WCOSS Phase 2	Planned accept Dec 2014 – FOC Jul 2015	iDataPlex	Linux (RHEL)	*1.9X over Phase 1	*4.4X over Phase 1	*44,400	*920 TF

*Estimated values; assumes FY13 NWS Reallocation Budget

System Characteristics – cont.

(per site)

	Life Cycle	Operational Use Time Requirement	Storage Useable	Power KW	Floor Space
Bridge System	Oct 2012 - Sep 2013	99%	0.80 PB	689 KW	3,100 SF
WCOSS Phase 1	Plan accept Dec 2012 – FOC Aug 2013	99.9%	2.59 PB	469 KW	4,060 SF
WCOSS Phase 2	Plan accept Dec 2014 – FOC Jul 2015	99.9%	*7.2PB	*1050 KW	4,060 SF

* Estimated values; assumes FY13 NWS Reallocation Budget

WCOSS HPC Growth

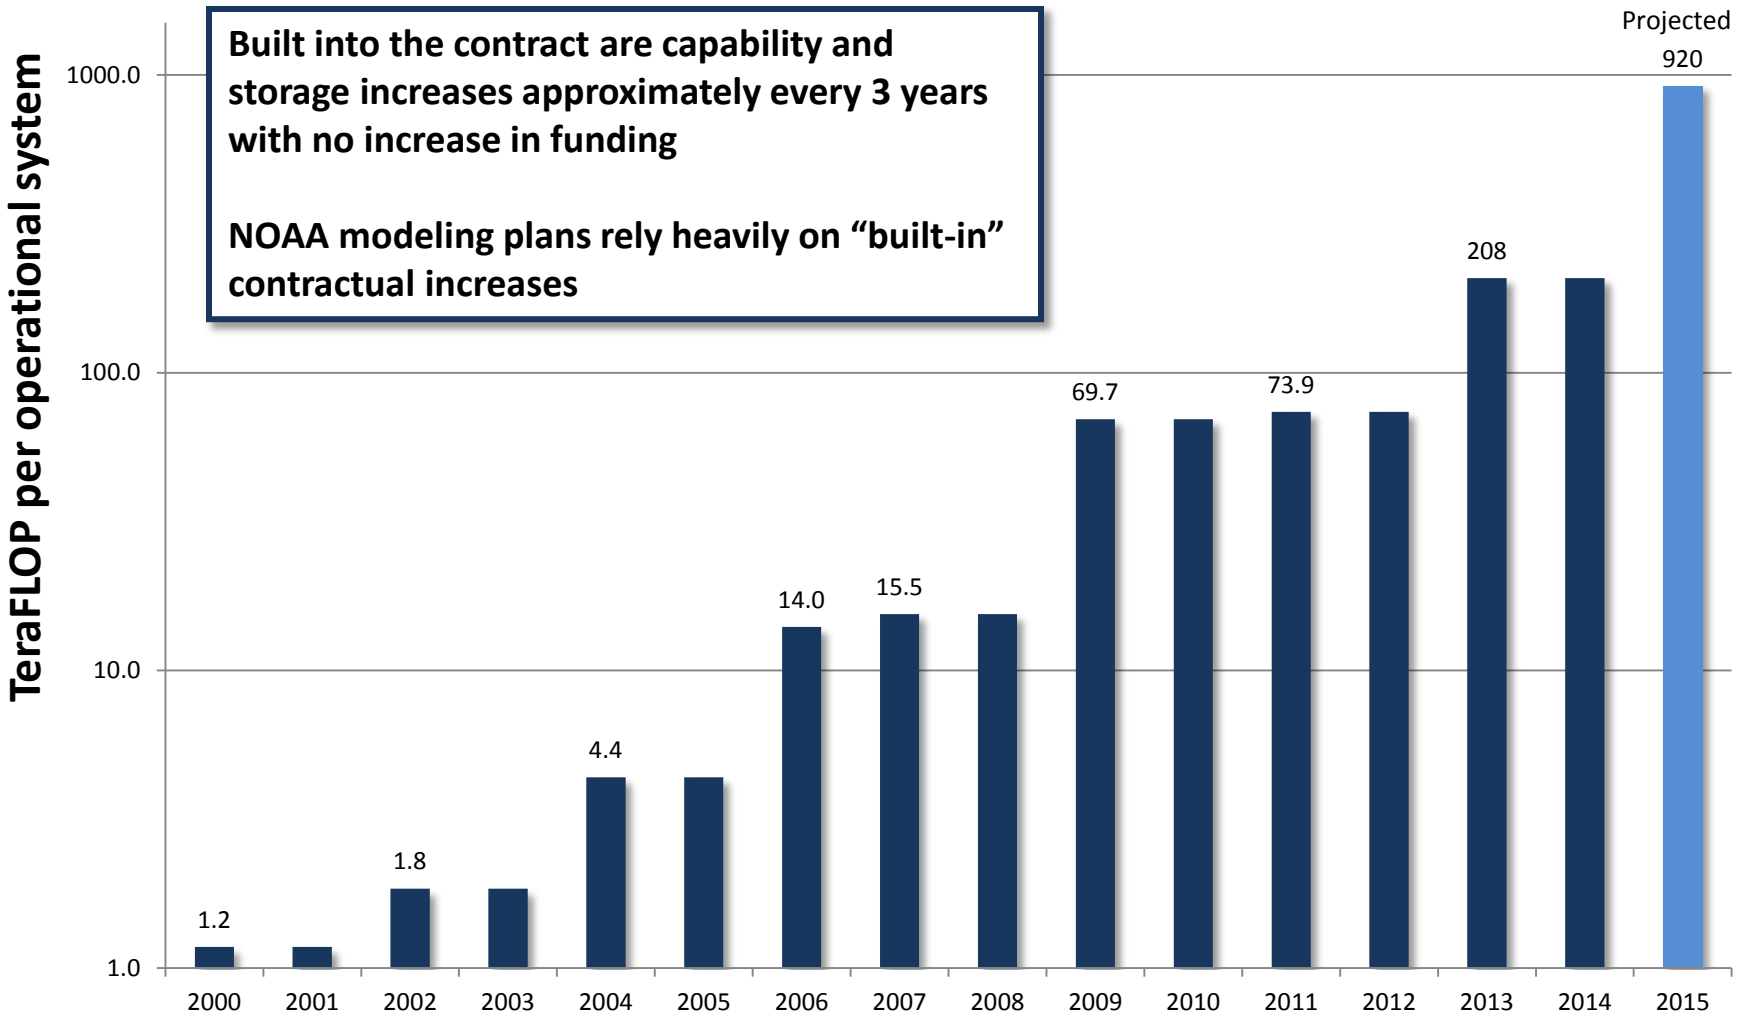

WCOSS – NOAA Weather and Climate Operational Supercomputing System

SCHEDULE & TRANSITION

Bridge Contract, WCOSS Contract and Transition from Bridge to new WCOSS Services

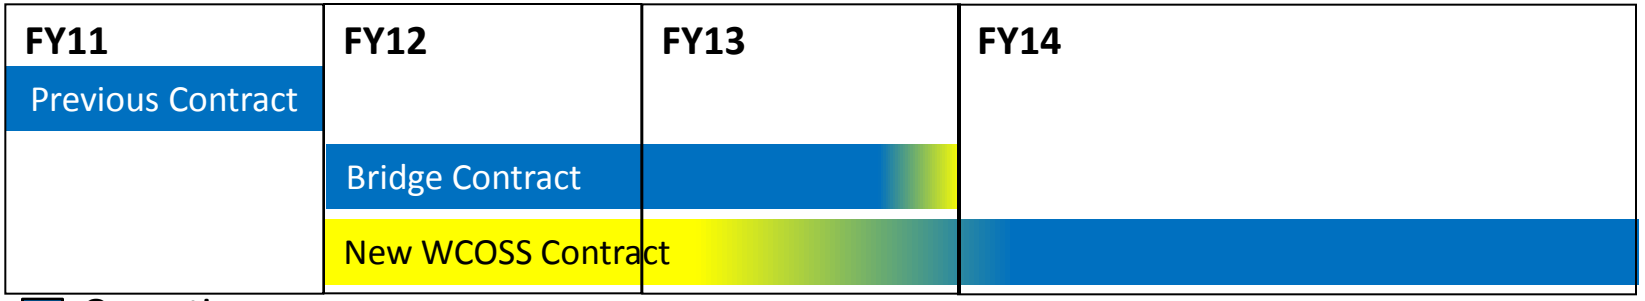

Major Milestones

Date	Milestone
✓Aug 2011	Awarded 2-year Bridge Contract
✓Sept 2011	Previous WCOSS contract expired
✓Nov 2011	Awarded new WCOSS ID/IQ Contract
✓Apr 2012	Installed early access system & started early testing with new architecture
✓Sep 2012	Complete installation of Primary system (Reston, VA)
Nov 2012	Complete installation of Backup system (Orlando, FL)
Dec 2012	Acceptance of all systems completed
May2013	Achieve Authority to Operate (ATO) for new WCOSS (NOAA8866)
July 2013	Complete transition porting, testing, validation and tuning
Aug2013	Go operational on new WCOSS (Ops on new WCOSS Contract)
Sep 2013	Close out Bridge contract

Transition End-to-End Validation

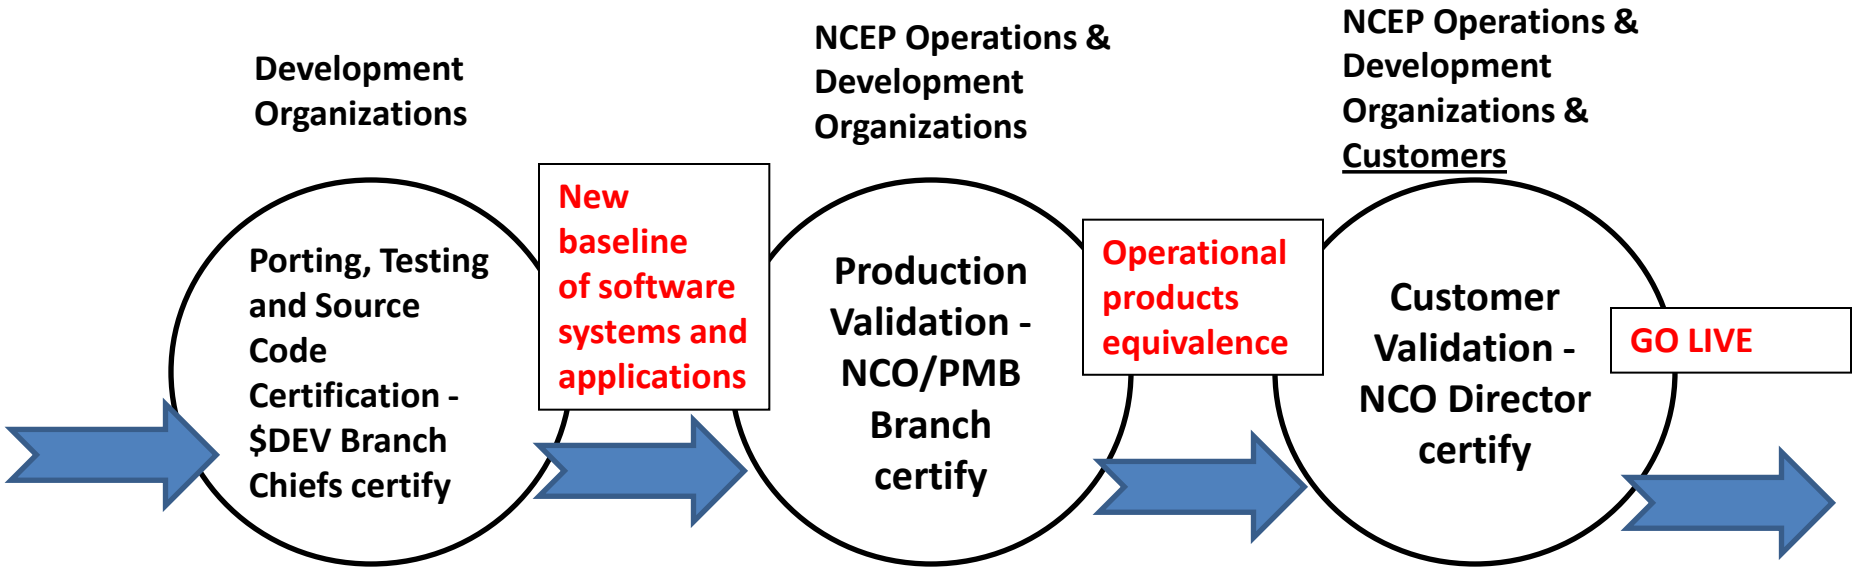

- Using the WCOSS Transition Test Checklist & Certification Document:
- Part I:Approval for planned changes
- Part II: Evaluation and certification
- Vertical structure

- Re-compile
- Integrate into production scheduler
- Tuning for resources and timing
- Post production/ product generation
- Validate output
- Downstream testing for other software systems
- Vertical structure

- Dataflow actions
- Validate network capacity
- Notify customers
- Receive customer feedback
- Make decisions and changes based on customer feedback

WCOSS – NOAA Weather and Climate Operational Supercomputing System

SYSTEM DETAILS

Phase 1 and 2 Software Overview

- Red Hat Linux OS
- IBM xCAT
- IBM Parallel Environment including runtime and developer editions
- Platform Computing Load Sharing Facility (LSF)
- ECMWF ecFlow Workflow Manager
- IBM General Parallel Filesystem (GPFS)
- Intel Cluster Studio
- Rogue Wave TotalView Debugger
- IDL
- Community Supported Software, e.g., BUFR, GRIB, netCDF, GEMPAK, GIS, GoogleMaps, GrADS, NCAR Graphics, VIS5D, Subversion

System Architecture

- IBM dx360M4 iDataPlex server based on the Intel Sandy Bridge Processor
- Clustered with a Mellanox InfiniBand FDR fabric
- Cluster components include the compute nodes, login nodes, high-performance interconnect, and the management infrastructure
- Configured with “hot spares” for high system availability
- Storage based on the IBM DCS3700 disk subsystem and GPFS
- Includes services to address the implementation, acceptance testing, and training for the transition from the current to new systems

Phase 1 System Architecture

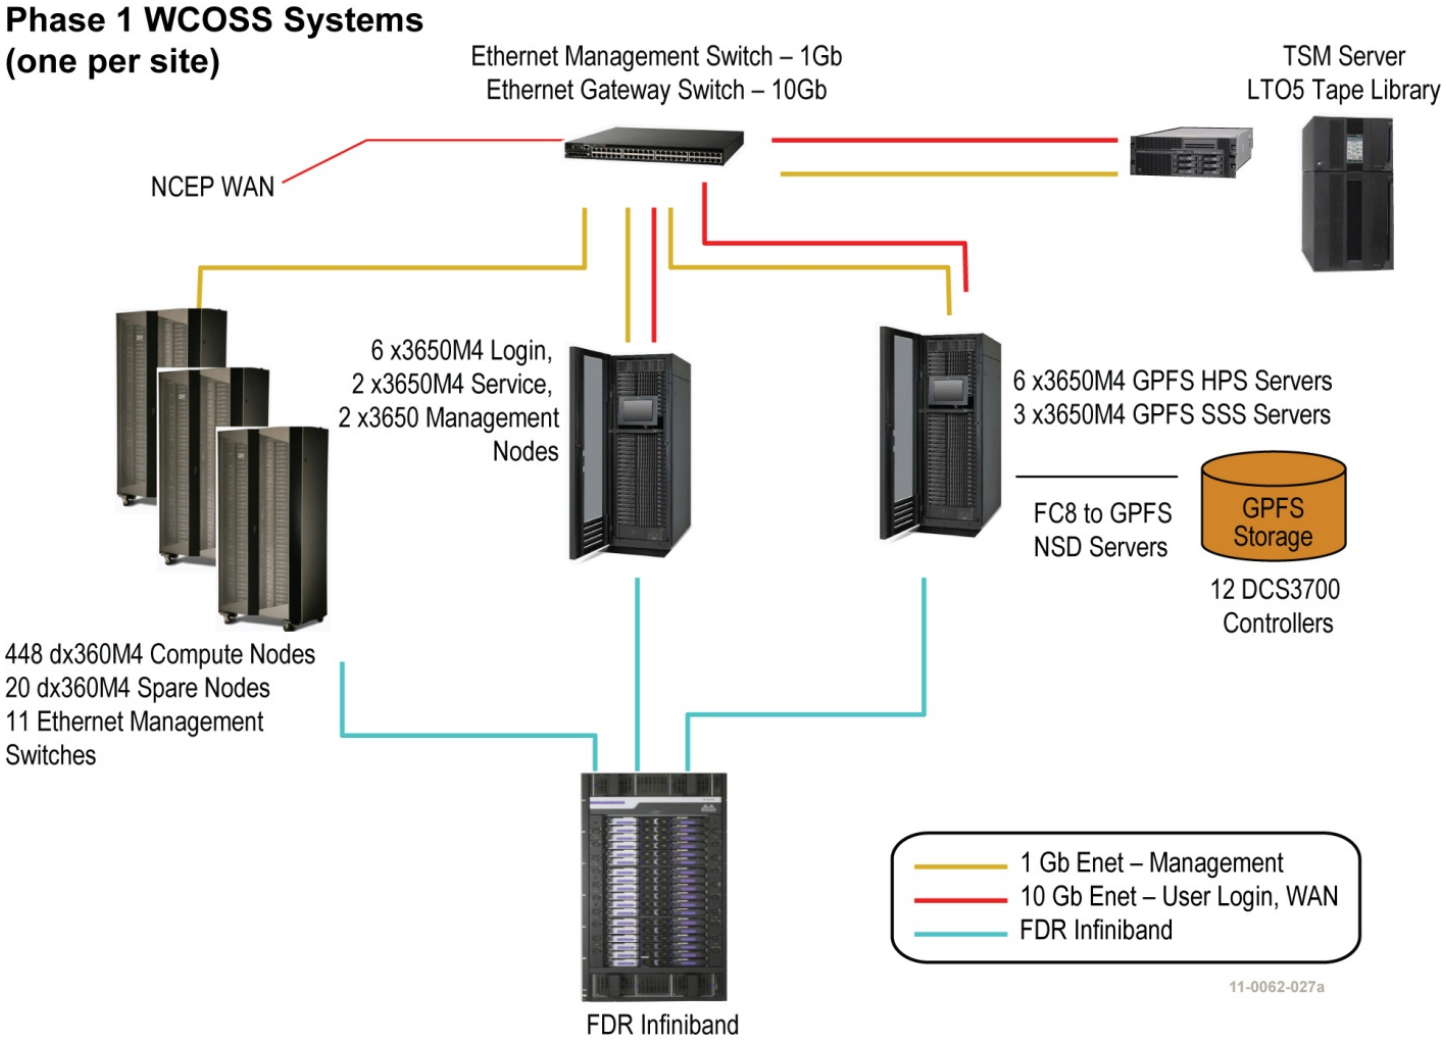

Phase 2 System Architecture

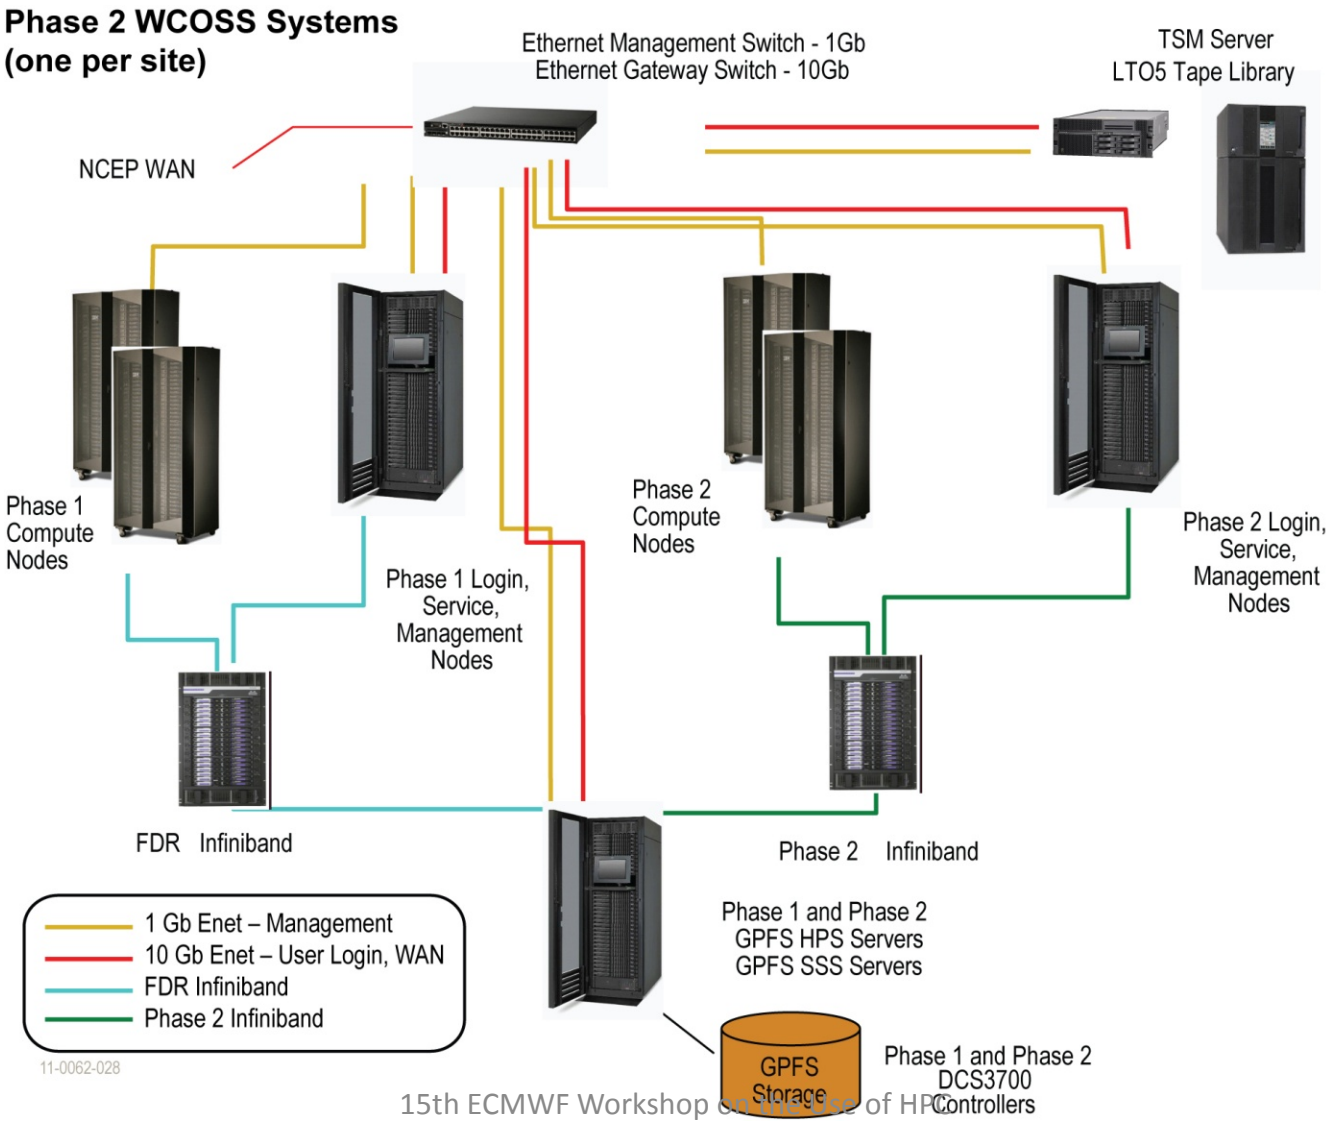

WCOSS – NOAA Weather and Climate Operational Supercomputing System

WRAP-UP

Transition

Key Considerations

- Schedule constraint with high impact consequences
  - Transition before bridge contract expires in September 2013
  - Current operational systems at peak load nearing end of life
- Resource dependencies external to investment
  - Transition highly dependent on resources from the NOAA Data Assimilation & Modeling/Science investment
- Technical
  - NOAA transitioning operational software systems/applications from legacy systems (Bridge) to a new/different architecture (WCOSS)
  - Technical challenges with integration of IBM's new hardware and system software architecture

Summary

- System growth curve delayed but resuming and accelerating (for now)
- Model implementations continued at the cost of system recoverability
- Contract architecture allows for agile response to budget changes
- System architecture designed with agility to match contract capability

WCOSS – NOAA Weather and Climate Operational Supercomputing System

QUESTIONS

ALLAN.DARLING@NOAA.GOV