
Federal Committee for Meteorological Services and Supporting Research (FCMSSR)

Neil Jacobs, Ph.D.

Assistant Secretary of Commerce
for Environmental Observation and Prediction

Deputy NOAA Administrator
and
FCMSSR Chair

October 22, 2018

Agenda

2:30 – Opening Remarks (Dr. Neil Jacobs, NOAA)

2:40 – Action Item Review (Mr. Michael Bonadonna, OFCM)

2:45 – Federal Coordinator's Update (OFCM)

3:00 – Accelerating NOAA's Next Generation Global Prediction System Research To Operations. (Dr. Jacobs)

3:30 – National Earth System Prediction Capability (ESPC) High Performance Computing Discussion. (Mr. David McCarren, ESPC Staff)

4:00 – Open Discussion (All)

4:20 – Wrap-Up (Dr. Neil Jacobs, NOAA)

FCMSSR Action Items

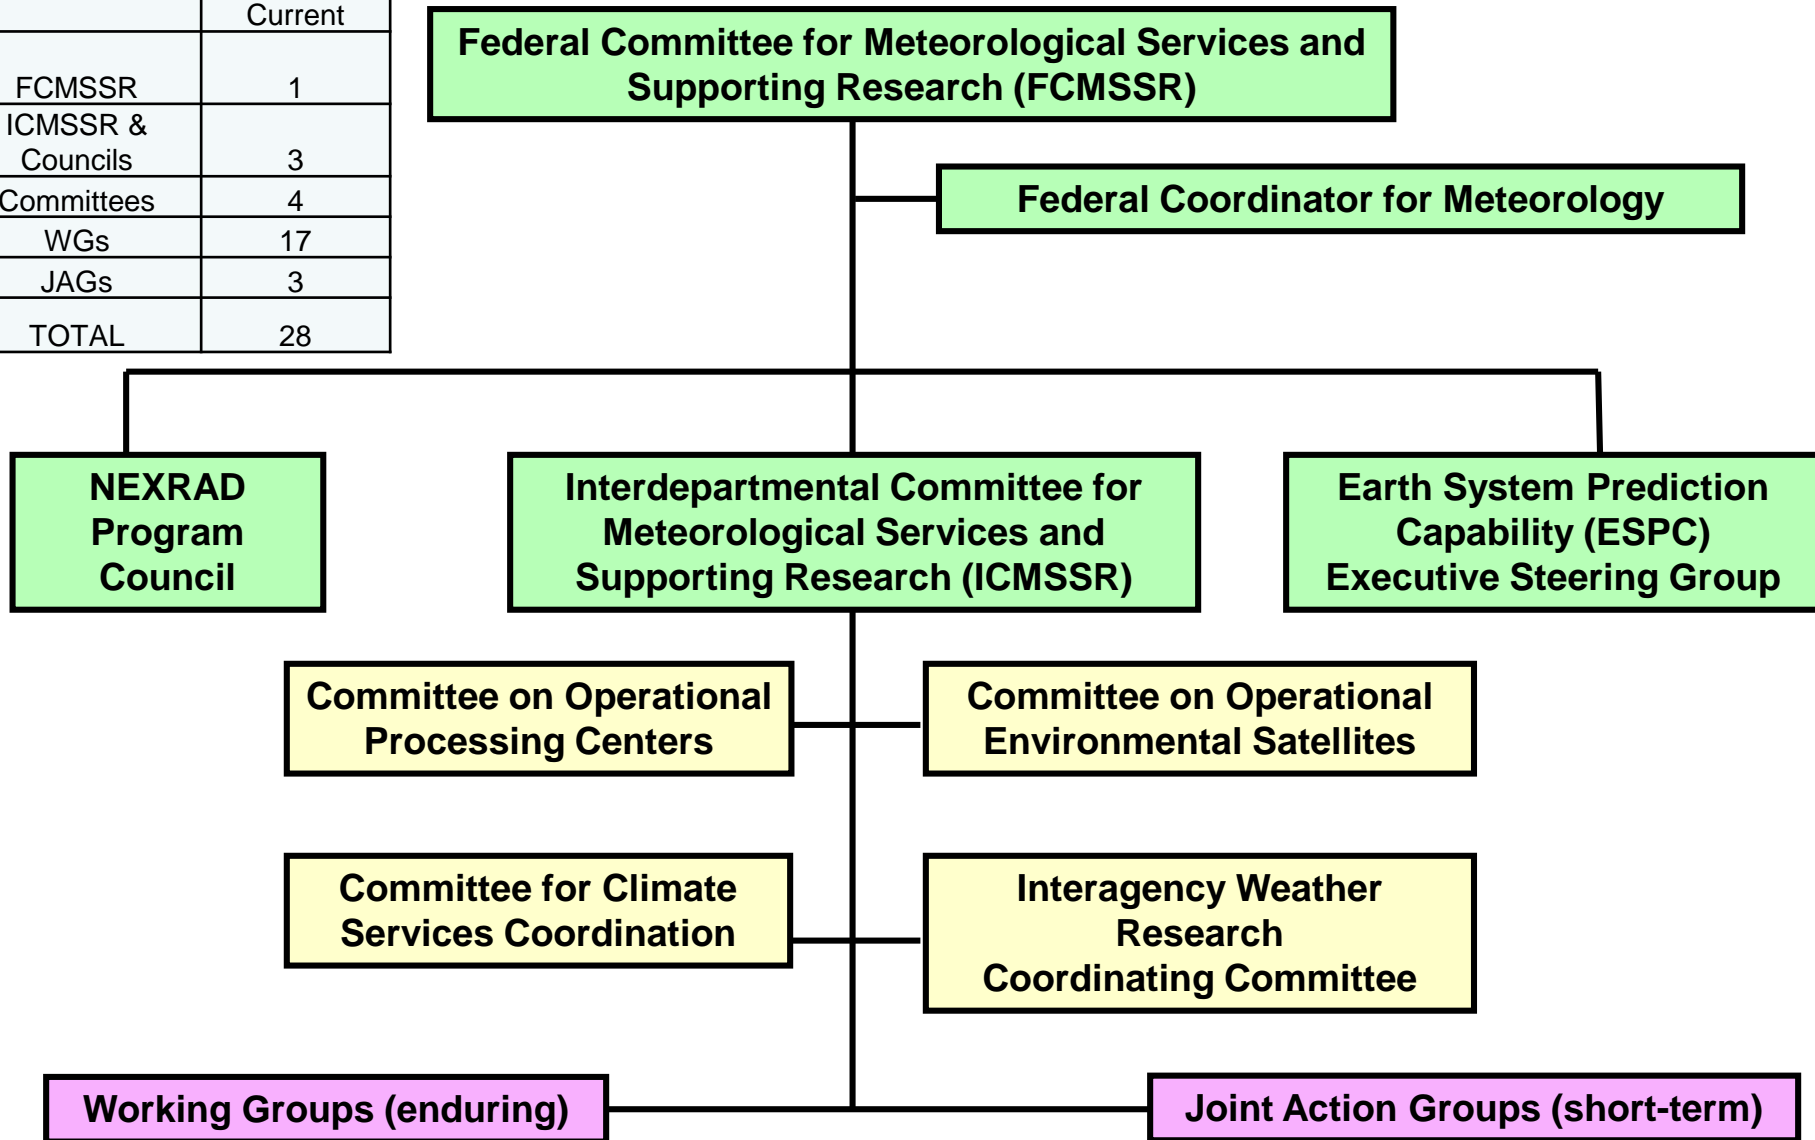
AI #	Text	Office Responsible	Comment	Status	Due Date
2018-1.1	Send email to OSTP recommending Option A: Rename the FCMSSR as ICAWS and make other changes in order to comply with the 2017 Weather Act. Request a written response from OSTP.	OFCM	5/2/18: Email has been sent. Awaiting Reply	Closed	05/04/18
2018-1.2	USAF A3W adjusts their 1340-series qualifications proposal letter as advised by FCMSSR and sends it to OFCM. OFCM drafts a cover letter for FCMSSR Chair endorsement and forwards the proposal to OPM	USAF A3W, OFCM, FCMSSR Chair	Seeking FCMSSR concurrence on revised proposed standard before sending to OPM	Open	05/11/18
2018-1.3	Review and brief FCMSSR on the impact of 1340-series qualification changes approximately one year after OPM implements the change.	USAF A3W, NWS		Open	10/31/19
2018-1.4	Brief FCMSSR on the NOAA Next Generation Global prediction Strategy as a possible framework for broader enterprise implementation	NOAA	Scheduled for this meeting	Open	10/31/18

FEDERAL COORDINATOR'S UPDATE

***Mike Bonadonna
Acting Federal Coordinator***

Federal Weather Enterprise Infrastructure

	Current
FCMSSR	1
ICMSSR & Councils	3
Committees	4
WGs	17
JAGs	3
TOTAL	28



Federal Coordinator's Update

1. Seeking FCMSSR concurrence on adding an objective to the Strategic Plan for Federal Weather Coordination.
2. Status of the FY2020 Federal Weather Enterprise Budget and Coordination Report
3. Seeking FCMSSR concurrence on changes to the U.S. Office of Personnel Management 1340-series Meteorologist qualification standard.
4. Implementing the Weather Act of 2017: Status of the proposal to rename FCMSSR as the “Interagency Committee for Advancing Weather Services (ICAWS)”

New Objective for the Strategic Plan for Federal Weather Enterprise Coordination

Goal 4. Conduct productive, synergistic interagency research efforts.

Objective 4.1: Exercise leadership in coordinating U.S. efforts in international weather research priorities including the current WMO Grand Challenges.

Objective 4.2: Foster interagency collaboration of research initiatives starting at the planning stage.

Objective 4.3: Support efforts among FWE participants to coordination task definition and sponsorship of National Academies research initiatives.

Objective 4.4: Expand interagency use of data and information for research.

New Objective 4.5: Develop coordination processes that facilitate operational feedback to the research community, and that accelerate the integration of promising research from federal, commercial and academic partners into operational improvements in observing, forecasting, warning and threat communication.

- ICMSSR approved 4.5, after noting that the Weather Act of 2017 placed an emphasis on improving research-to-operations processes, but these were not mentioned in the Strategic Plan.
- Likely vehicles for execution of Objective 4.5 are the IWRCC, WG/TCOR, and COES.
- Recommend FCMSSR approval.

Budget and Coordination Report FY20



- **Working Group for BCR formed.**
- **Terms of Reference in review.**
- **Initial guidance to agencies for FY20 in draft; has been discussed at agency working levels.**
- **Formal request for FY20 budget information will be issued in December.**
- **Due date for getting budget information to OFCM will be ~2 weeks after release of FY20 President's Budget Request.**

Update 1340 Meteorologist Series Qualification

Recap/path forward:

1. USAF proposed changing the OPM qualification standard for the 1340 series in order to widen the applicant pool and prevent inadvertent automated de-screening of prospective successful meteorologists.
 2. ICMSSR and FCMSSR reviewed the USAF proposal and approved it for forwarding to OPM.
 3. OPM requires a department-level Chief Human Capital Officer (CHCO) endorsement before considering a proposal.
 4. DOC CHCO office recommended the proposal be revised to ensure compliance with Merit System Principles and to conform to OPM's qualification framework.
 5. NOAA Deputy Administrator and USAF A3W agreed to new proposal
-
6. NOAA leadership will secure concurrence for the new proposal and resubmit the package to DOC CHCO.
 7. Any additional changes will be cleared through the FCMSSR.

Update 1340 Meteorologist Series Qualification:

Existing requirements

- A. Degree in applicable science that includes 24 semester hours in meteorology/atmospheric science, 6 semester hours in physics, 3 semester hours of ordinary differential equations, 9 semester hours of course work appropriate to physical science major, or....
- B. Combination of education and experience: Course work shown above **plus** appropriate experience or additional education.

Implementing Weather Act of 2017 § 402

Mandate: The 2017 Weather Act directs OSTP to establish “an Interagency Committee for Advancing Weather Services (ICAWS),” with duties including identifying, prioritizing, and coordinating top forecast needs, and sharing needs and improvements across agencies.

Proposal:

- Rename FCMSSR as ICAWS, combine ICAWS duties and existing FMCSSR duties into one ICAWS charter.
- OSTP and NOAA co-chair ICAWS
- OSTP submits legislative change request to make OFCM Director the ICAWS Exec. Sec. (vice co-chair.)

Status:

- FCMSSR approved recommending this plan to OSTP at the April 2018 FCMSSR.
- OSTP received recommendation from FCMSSR via OFCM in May 2018.
- OSTP Director (President’s Science Advisor) nominee awaiting confirmation. OSTP staff will hold approval of this plan until a Director is confirmed and can see the plan.

Agenda

2:30 – Opening Remarks (Dr. Neil Jacobs, NOAA)

2:40 – Action Item Review (Mr. Michael Bonadonna, OFCM)

2:45 – Federal Coordinator's Update (OFCM)

3:00 – Accelerating NOAA's Next Generation Global Prediction System Research To Operations. (Dr. Jacobs)

3:30 – National Earth System Prediction Capability (ESPC) High Performance Computing Discussion. (Mr. David McCarren, ESPC Staff)

4:00 – Open Discussion (All)

4:20 – Wrap-Up (Dr. Neil Jacobs, NOAA)



NOAA

Improving Operational Numerical Guidance and Forecast Skill

Neil Jacobs, Ph.D.

Assistant Secretary of Commerce for Environmental Observation and Prediction
Deputy NOAA Administrator






Improving forecast skill and performance



- Strategic Implementation Plan (SIP) for Unified Forecasting System (UFS)
- Quality control of observations
- Data assimilation
- Dynamic core and model physics
- Code efficiency
- Optimized hardware



Anomaly correlation: higher is better





Inherent barriers with the status quo



- Fractured internal strategy and mission creep
- Fractured external strategy across various agencies with different priorities
- Obtuse HPC procurement process (both hard iron and cloud)
- Security clearance procedures for visiting scientists
- Cultural (internal and external)
- Funding allocation process disincentivizes collaboration
- Risk aversion (incentive not to fail >> incentive to improve)
- Too many committees with overlapping and conflicting input
- Lack of documented, supported, and portable community code





Accelerate R2O and O2R overview



- End-to-end community model (harness collective advancements)
- Focus on UFS infrastructure (NOAA-NCAR MOA)
- VM (cloud HPC) for on-demand parallel “surge” development
- Visiting scientists (PrepIFS)
- Formalizing the R2O funnel (requirements, gates and transitions)
- Fast-tracking satellite DA / Drive up benefit in cost-benefit ratio
- Agile/nimble “skunk works” sandbox
- Governance, funding, and streamlining committees, etc.
- Accelerating SIP: building off momentum and addressing process deficiencies



End-to-end community model



- Weather Research and Forecasting Innovation Act of 2017
- Code development force multiplier once community has access
- UMAC / grad student test
- WRF-ARW
- Initial heavy lift on us / help from vendors
- System architecture agnostic





VMs for on-demand parallel “surge” development



- VM (cloud HPC) for on-demand parallel “surge” development
- Problem : NOAA scientists have long queue to run jobs
- NOAA research : operations compute = 1:1
- ECMWF roughly 5:1
- VMs: X:1 (X=scalable on demand)
- Parallel experiments and testing already going on at GFDL
- Not meant to replace NCO “hard iron”





Virtual machines and cloud HPC



- Virtual machines (cloud HPC) for parallel community development
- Not all HPC “clouds” are the same
- NWP needs 2 things:
 1. Remote Direct Memory Access (RDMA)
 - Remote memory location read/write
 - Direct processor interface bypasses kernel in I/O path
 2. Fast interconnect speeds
 - AWS* currently has 25 Gbps (moving to 100 Gbps)
 - Azure currently has 100 Gbps (moving to 400 Gbps)
 - WCOSS has 100 Gbps



***Note: NASA study often cited used 25 Gbps AWS and mentions slow interconnect speed as the primary limiting factor**





Global model cloud pilot (GMCP)



- Port and test NGGPS in various cloud HPC and VM environments
- Conduct “bake off” with 3rd party software engineering firms and cloud vendors
- Port and containerize FV3 GFS code for various cloud architectures
- NOAA provides raw data and source code
- Vendors are free to select/design cloud architecture and optimize as needed
- Evaluation based on run time, output precision, core count/scalability, portability between vendors, and cost
- Success defined as
 - Ability to match or beat production model run time
 - Numerically replicate production model output
 - Match or beat cost of running on owned on-prim HPC
- Testing initiated by late Q3 FY19
- Results evaluated by Q1 FY20



Visiting scientists and external expertise



- Move the development sandbox outside NOAA
- Avoid clearance process time lag
- Avoid limited compute resources
- Need accessible interface / PrepIFS
- Will need robust secure ingest (DMZ) for external sandbox
- Modularized code - don't have to know/run entire system
- We need better agency collaboration with NASA, DOE, USAF, NAVY, etc.
- Code sprints / JEDI academy – already have a true 4D-Var FV3 running!!!





Formalizing the R2O funnel



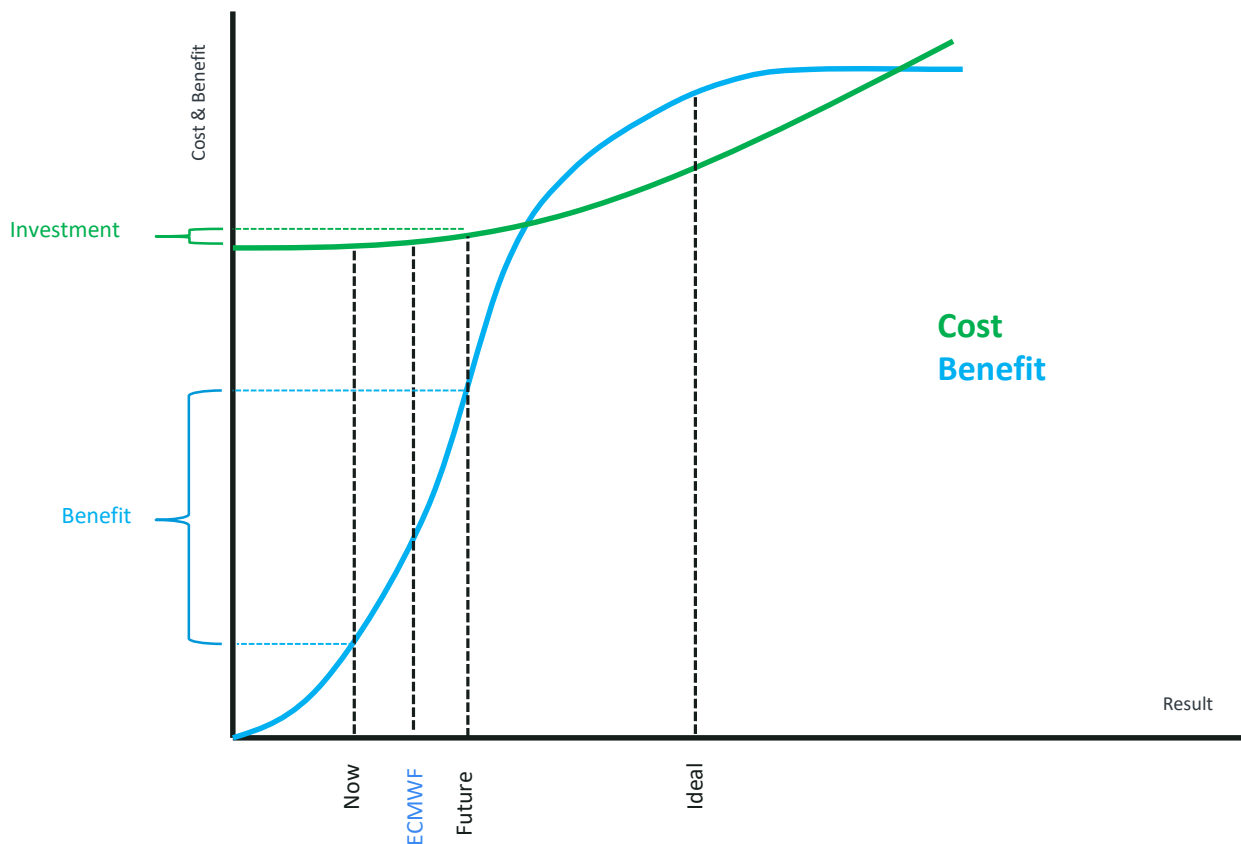
- Follow SIP strategy, but build on momentum and speed up process
- Initial baseline requirements with operations in mind (gates and transitions)
- Problems should be suggested (versus tasks)
- Objective evaluation process to transition through gates
- UMAC: evidence-based decisions
- Parallel production environment (possibly many)
- Academia and labs (community)
- Software engineers brought in at initial stages
- EMC involved throughout the process (avoid forklift approach)





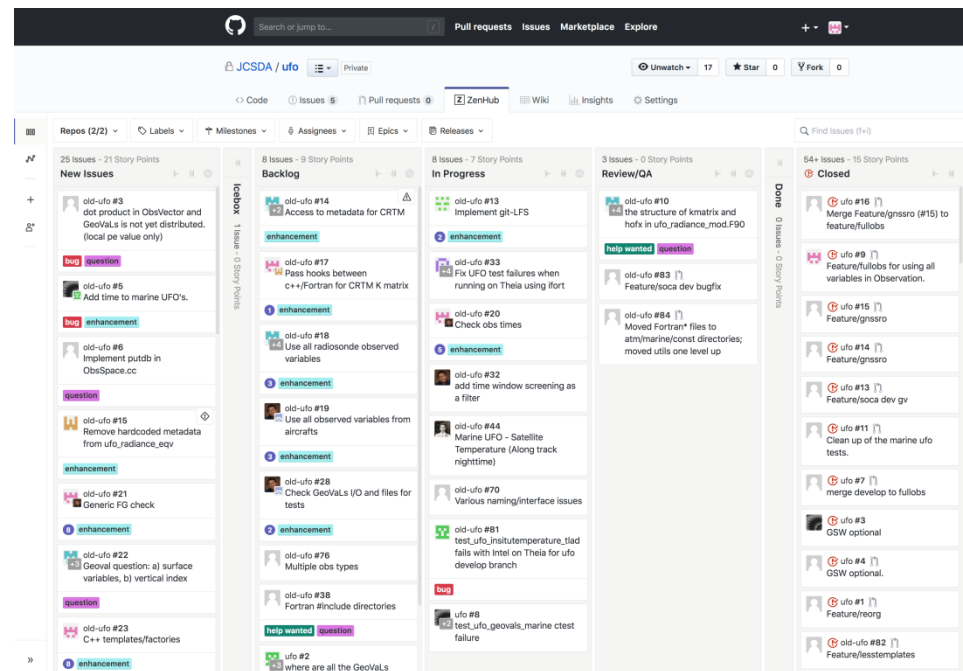
Fast-tracking satellite DA

- Optimize all data, but satellite has a lot of value left on the table
- Critical area we are lagging ECMWF
- Drive up *benefit* in cost-benefit ratio
- JCSDA JEDI



Agile/nimble sandbox

- Avoid risk aversion:
 - IBM's Watson Research Center
 - Lockheed Martin's Skunk Works
 - Google X
 - Boeing Phantom Works
 - Amazon's Lab126
- UK Met Office (offsite dev)
- Avoid outside distractions, etc.

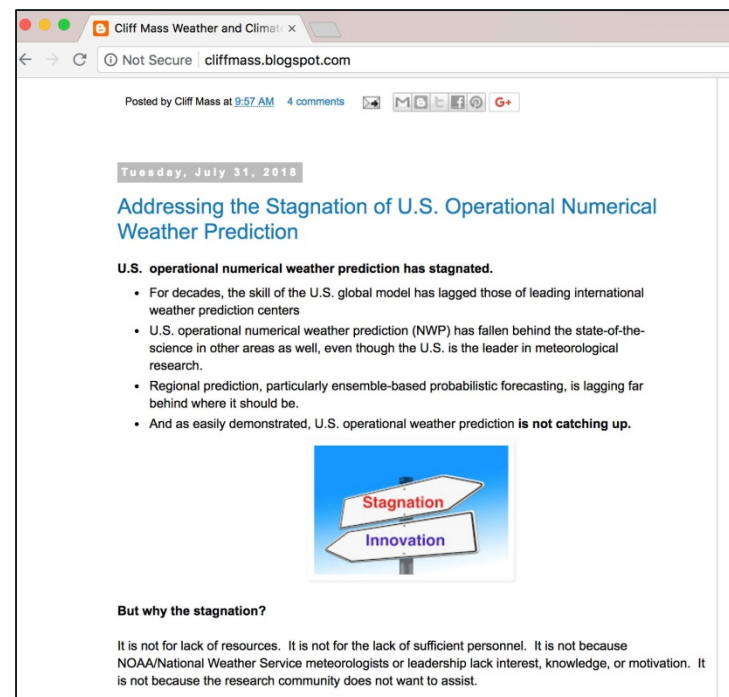





Accelerating Strategic Implementation Plan



- Building off momentum (SIP/DTC progress!)
- Addressing process deficiencies
- Committees and Governance
- Involvement / ownership
- Funding (procurement, stakeholders)
- Branding / image





NOAA

Thank You!



Agenda

2:30 – Opening Remarks (Dr. Neil Jacobs, NOAA)

2:40 – Action Item Review (Mr. Michael Bonadonna, OFCM)

2:45 – Federal Coordinator's Update (OFCM)

3:00 – Accelerating NOAA's Next Generation Global Prediction System Research To Operations. (Dr. Jacobs)

3:30 – National Earth System Prediction Capability (ESPC) High Performance Computing Discussion. (Mr. David McCarren, ESPC Staff)

4:00 – Open Discussion (All)

4:20 – Wrap-Up (Dr. Neil Jacobs, NOAA)

Purpose-Built HPC for Earth System Prediction

Dave McCarren, Project Manager
National ESPC



Background/history

- Aug '15 HPC working group self-formed
- Jan '16 Brief to NSF/NSCI
- Apr '17 NSF RFI input
- Apr '17 Published position paper
 - <https://doi.org/10.7289/V5862DH3>
- Oct '17 Briefed NSCI
- Nov '17 Supercomputing '17 Ad hoc session
- Apr '18 Info brief to FCMSSR on HPC WG activities
- Aug '18 Request Brief to ICMSSR
- Aug '18 Formal query to NSCI on Operational NWP support
- Oct '18 Request brief to FCMSSR
- Nov '18 Supercomputing '18 Birds of a Feather session

Earth System Prediction and NSCI

- NSCI strategy document identifies deployment agencies
 - NOAA/NWS operational computation time limits
 - Note: DOD agencies have similar operational constraints
- Lead agencies designing and providing metrics for future systems
 - Exascale computing project funds climate-scale (Cloud-Resolving Climate Modeling of the Earth's Water Cycle, Mark Taylor (SNL) with ANL, LANL, LLNL, ORNL, PNNL, UCI, CSU)
 - Other E3SM work (SNL, LLNL, LBNL, LANL, ORNL, etc.)
 - Does not meet operational time constraints

Agency	Lead Agencies	Foundational R&D Agencies	Deployment Agencies
Department of Defense (DOD)	X		
Department of Energy (DOE)	X		
Department of Homeland Security (DHS)			X
Federal Bureau of Investigation (FBI)			X
ODNI/Intelligence Advanced Research Projects Activity (IARPA)		X	
National Aeronautics and Space Administration (NASA)			X
National Institutes of Health (NIH)			X
National Institute of Standards and Technology (NIST)		X	
National Oceanic and Atmospheric Administration (NOAA)			X
National Science Foundation (NSF)	X		

Table 1. Summary of Agency Roles

<https://www.hpcwire.com/2016/09/07/exascale-computing-project-awards-39-8m-22-projects/>

https://www.sandia.gov/news/publications/labnews/articles/2018/11-05/E3SM.html#E3SM_Sandia

Earth System Prediction Computing: Technical Challenges

- Models do not scale up efficiently:
 - Performance wall: workload grows as 4th power of resolution, resources grow as 2nd power of resolution
 - fluid flow calculations are parallel in 3 spatial dimensions, limited by data bandwidth to memory, other supercomputer components
 - physical parameterizations are parallel in 2 spatial dimensions (parallelism in vertical is limited due to extremely fast physical coupling)
- Even those that do scale only use 6% of current CPU processor, and 1-2% of GPU processors
- Will soon result in inability to operationalize model complexity/resolution improvements available from research

ICMSSR action items

ICMSSR Action Item 2018-3.3. Ask the National Strategic Computing Initiative (NSCI) if high-resolution, fully coupled, long temporal range environmental prediction is still one of their targeted applications for the high performance computing systems that their member agencies are developing.

ICMSSR Action Item 2018-3.4. Based on the response to AI 2018-3.3, request FCMSSR approval of the following course of action:

- If “yes” ESPC will pursue a closer working relationship with NSCI to connect application agency needs (e.g. NWS, DOD) with NSCI progress and plans.
- **If “no,” or if “yes but only in a research mode, or in another mode that would not meet operational needs”**
 - ICMSSR (through OFCM) will coordinate further discussion with the NSCI on environmental prediction in HPC.
 - ESPC will coordinate a study to provide information that will enable federal agencies with environmental prediction responsibilities to map an effective path forward towards high-resolution, fully coupled, long temporal range environmental prediction.

NSCI Response: project goals/resources

1) What are the current goals/objectives of the environmental prediction project, such as resolution and speed of the computation system, which environmental parameters would be predicted, and how far ahead would the prediction system look?

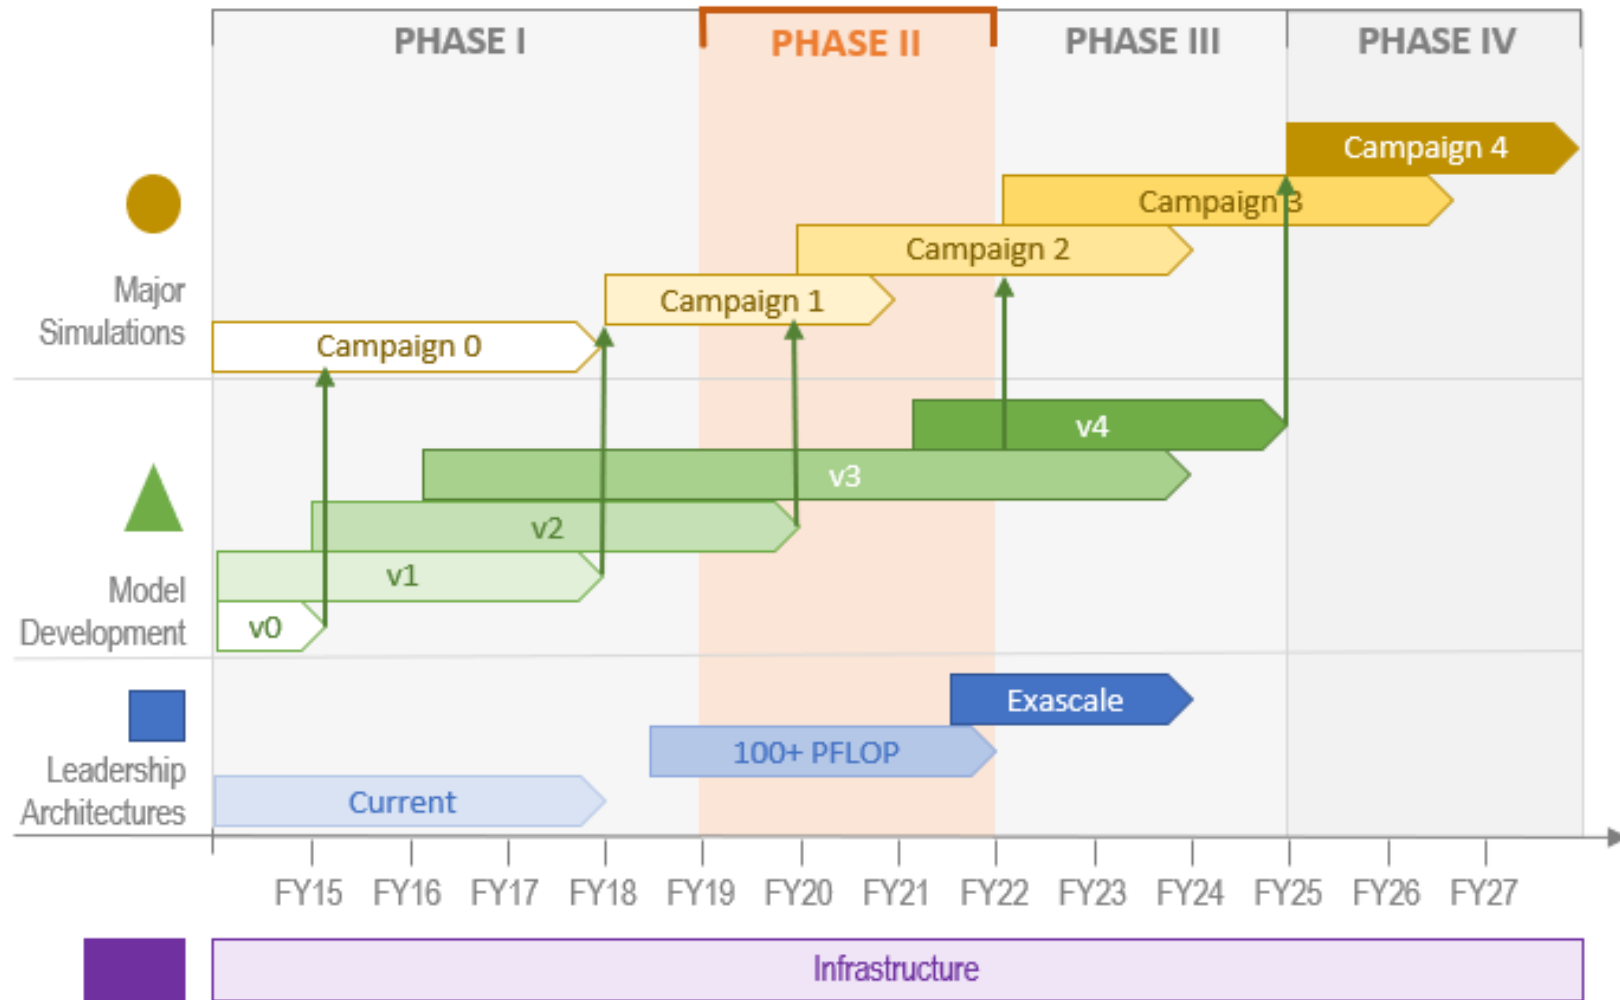
E3SM has a goal of fully-coupled decadal and longer term climate simulations using computational and scientific advances for convection-permitting and convection-resolving simulations with horizontal grid spacings between 1 and 5 km, in the atmosphere and 5-30 km (eddy-resolving) resolution in the ocean and sea-ice.

3) Which agencies are contributing manpower, computational facilities, and/or funding resources to this effort? Is the outlook that sufficient resources are available, or are shortfalls expected?

US Department of Energy/Office of Science

(timeline for #2 on next slide)

NSCI Response: Project Timeline



We request support for: Interagency Study on Purpose-Built HPC

- The National ESPC HPC working group advocates for an interagency study investigating:
 - the **widening gap between earth system application requirements and currently evolving HPC**
 - a supercomputing system designed with the singular purpose of running **exascale earth system prediction models in operational time constraints**
- This study will:
 - help identify the current needs of earth system prediction models
 - determine whether or not a purpose-built earth system prediction computer is feasible from several perspectives, including cost and efficiency
- Birds of Feather session at SuperComputing 2018 may discuss this study with the broader community

We also request support for: Earth System Prediction computation at the National Strategic Computing Initiative

- Ask the NSCI PM to brief at the next FCMSSR meeting?
- Clarify/demonstrate communication path for operational requirements to NSCI?
- Inform your representatives of these issues.

Open Discussion

Wrap-Up

- OFCM will document any new Action Items and provide the meeting Record of Action within two weeks.
- Next FCMSSR meeting proposed for **April 2019**
- **Wrap-Up (Chair)**

Backup material

1340 Meteorologist Series Qualification: Revised Proposal

A. Degree in Meteorology or Atmospheric science

OR

B. A degree in other natural science major that included:

- At least 24 semester hours of credit in meteorology/atmospheric science including a minimum of:
 - Six semester hours of atmospheric dynamics and thermodynamics;*
 - Six semester hours of analysis and prediction of weather systems (synoptic/mesoscale);
 - Three semester hours of physical meteorology; and
 - Two semester hours of remote sensing of the atmosphere and/or instrumentation.
- Six semester hours of physics, with at least one course that includes laboratory sessions.*
- Three semester hours of ordinary differential equations.*
- At least nine semester hours of course work appropriate for a physical science major in any combination of three or more of the following: physical hydrology, statistics, chemistry, physical oceanography, physical climatology, radiative transfer, aeronomy, advanced thermodynamics, advanced electricity and magnetism, light and optics, and computer science.

** There is a prerequisite or corequisite of calculus for course work in atmospheric dynamics and thermodynamics, physics, and differential equations. Calculus courses must be appropriate for a physical science major.*

OR

C. Combination of education and experience – at least 18 semester hours of meteorology or atmospheric science and at least three semester hours in differential equations (or an equivalent course), plus appropriate experience or additional education.

Evaluation of Education: Courses acceptable towards meeting the meteorology course requirement in paragraph C must include at least three of the following: atmospheric dynamics and thermodynamics, analysis and prediction of weather systems, physical meteorology, remote sensing and/or instrumentation.

Study Objectives

- Performance measurement and modeling to systematically collect and characterize detailed, quantitative requirements from the earth system modeling community;
 - Corresponding detailed measurement and characterization of current and roadmap technologies for processor, memory system and network technologies;
 - Gap analysis to determine if custom design or manufacture of components would be cost-effective for a system focused on PDE solution, including the level of customization and spanning the processor, interconnect, memory, and other essential parts of a computing system;
 - Determine if a PDE-solving supercomputing platform would benefit from specific (and custom) software such as compilers, libraries, programming models or domain-specific languages;
 - Estimation of a rough order of magnitude of investment needed for such a custom-built supercomputer
-

Priority: Share results with Vendors

What is Required for an Interagency Study?

- Planning and coordination across the involved agencies
 - Identify common objectives
 - Promote cross-agency visibility for understanding the current state of HPC platforms for earth system prediction
 - Involve HPC hardware & software experts
- Identify deliverables and estimate costs
- Agree to funding commitments
- Options for Agency-funded study:
 - Agency PMs fund
 - NOPP study - invited hardware vendors and/or HPC research firms
 - NRC study - funded by ESPC agencies



Commercial Cloud Trade Study


Goal: *Evaluate the suitability of commercial clouds for HPC Applications*

Approach

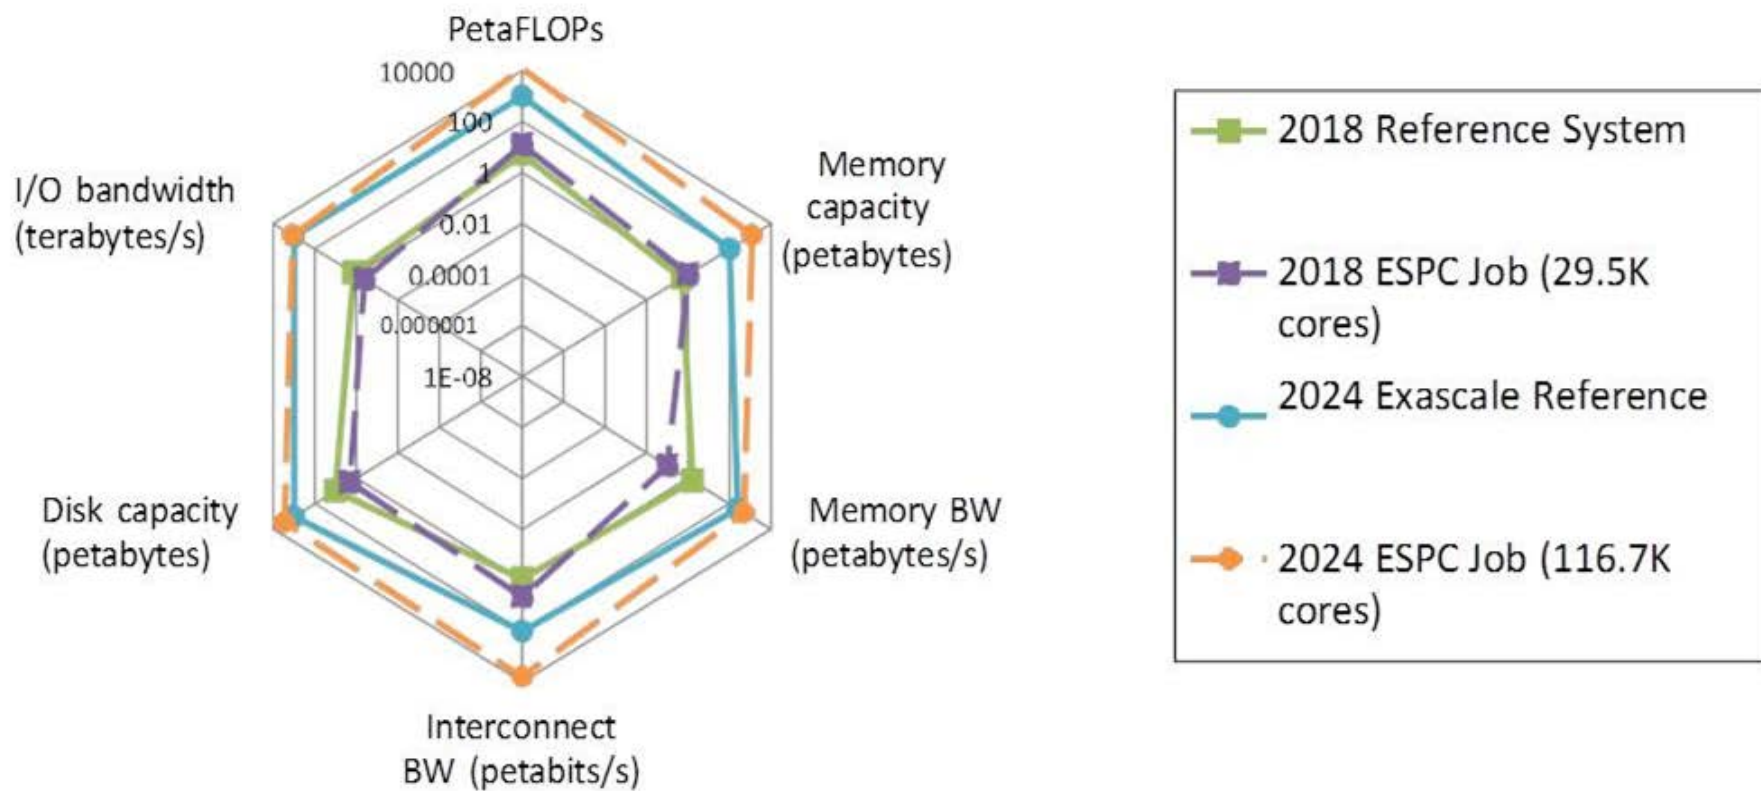
- **Workload:** NPBs, six full-sized applications (ATHENA++, ECCO, ENZO, FVCore, WRF, OpenFOAM)
- **Systems:** HECC systems Pleiades and Electra, Amazon Web Services (AWS), Penguin-on-Demand (POD)
- **Cost Basis:** HECC – full cost of running (hardware, software, power, maintenance, staff, and facility costs); AWS and POD – only the compute costs from published rates and any publicly-known discounts (spot pricing, lease price, etc.)
- **Key Findings:** *Commercial clouds currently do not offer a viable, cost-effective approach for replacing in-house HPC resources for NASA HPC applications. However, there may be use cases where a commercial cloud is a viable alternatives, e.g., specialized hardware*
- **Actions:**
 - Continually evaluate the suitability of commercial clouds
 - Develop an environment to support bursting to commercial clouds (on a full-cost recovery basis) for S&E projects – Phase 1 pilot project available September end.

Evaluating the Suitability of Commercial Clouds for NASA's High Performance Computing Applications: A Trade Study: Chang et al, NAS Technical Report NAS-2018-01, May 2018
https://www.nas.nasa.gov/assets/pdf/papers/NAS_Technical_Report_NAS-2018-01.pdf

Suitability of Cloud for HPC - from Navy Study (Preliminary)

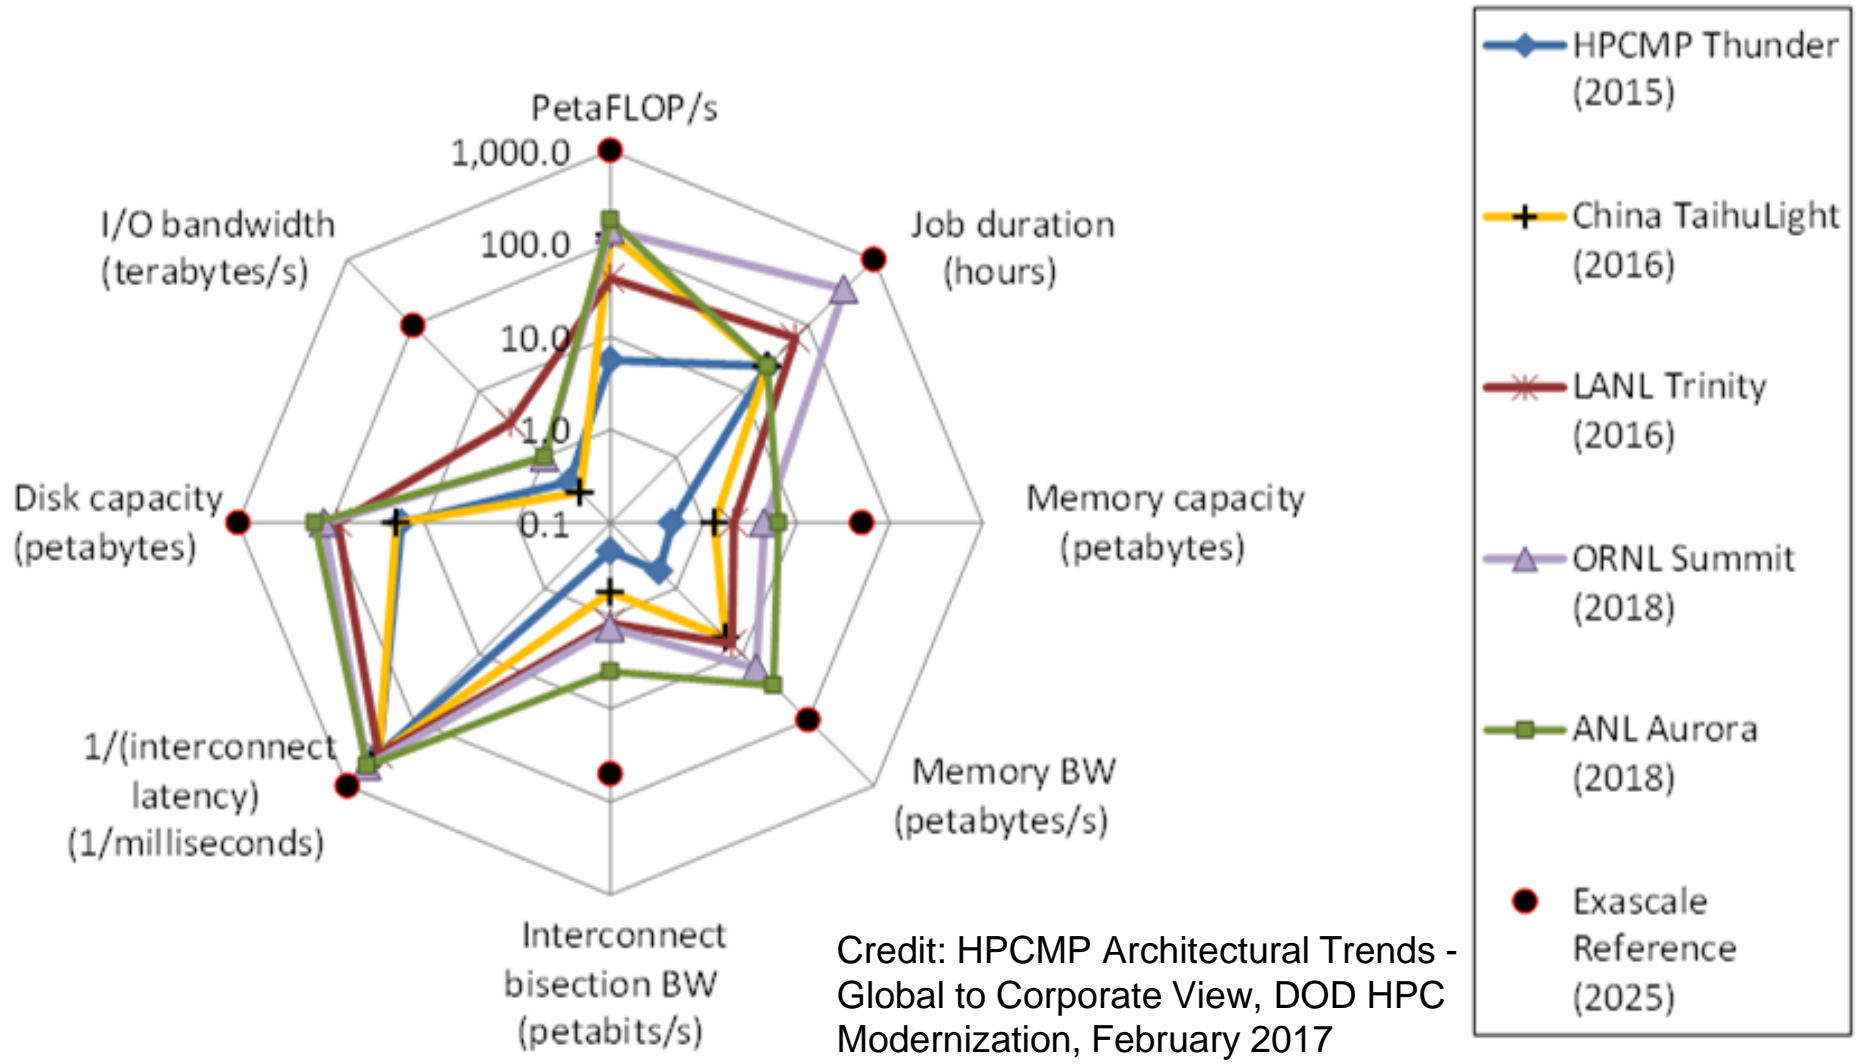


HPC Requirements for Earth System Modeling





Internal report: *The Future of DoD Climate, Weather and Ocean High Performance Computing Requirements*, 15 Aug 2016, Figure 24

HPC Outlook



Maximizing Peak Performance

- Without a way around memory-bandwidth issues, even models that scale strongly will not be able to harness the full power of the hardware.



* Results from recent publica8on by Mueller-Kopera-Marras-Isaac-Wilcox-Giraldo


Scalability of NUMA on Mira (IBM BG/Q) using the full machine: at 3 million MPI ranks, the model scales perfectly.

Computation efficiency measured in percent of peak performance. Red curve is around 12%; rest of the code is below 6%.


- Today, most NWP codes will be around the 5-7% peak; this number needs to go up.
- Recommendation: even after using hardware-agnostic languages, we still need to optimize

Memory vs. Compute Bound

- Current models are memory-bandwidth bound.
- Here we show roofline plots for the NUMA model on Titan (Nvidia K20 GPUs) on the left and on one node of Mira (IBM BG/Q) on the right.
- The sloped line shows the peak memory-bandwidth of the hardware and the flat line shows the peak computational performance. Note that all the different parts of the code are near the memory-bandwidth line (we are at the mercy of the communication speed of the hardware because we are moving way too much data). We desperately need to get around this barrier.



* Results from recent publication by Abdi et al. (2016)



* Results from recent publication by Mueller et al (2016).

Possible Solutions to Future HPC Challenges

- Two Approaches
 - Hardware-optimized: Different compute-kernels for each computer.
 - e.g., CUDA/OpenCL or OpenACC for GPUs and Intel Cilk or OpenMP for Xeon Phi
 - Hardware-agnostic: Write compute-kernels in one language, then write translators for each platform.
 - This is the idea behind OCCA* (Virginia Tech), Kokkos* (Sandia National Laboratory), Stella* (ETH), PSyclone (UK Met Office), and OpenACC*(NOAA) hardware-agnostic languages.
- Main Metrics
 - Time-to-solution (wallclock time)
 - Percentage of computer required
- A common modeling or computing technology would simplify this effort, but may not be possible.

*OCCA: <http://libocca.org/>

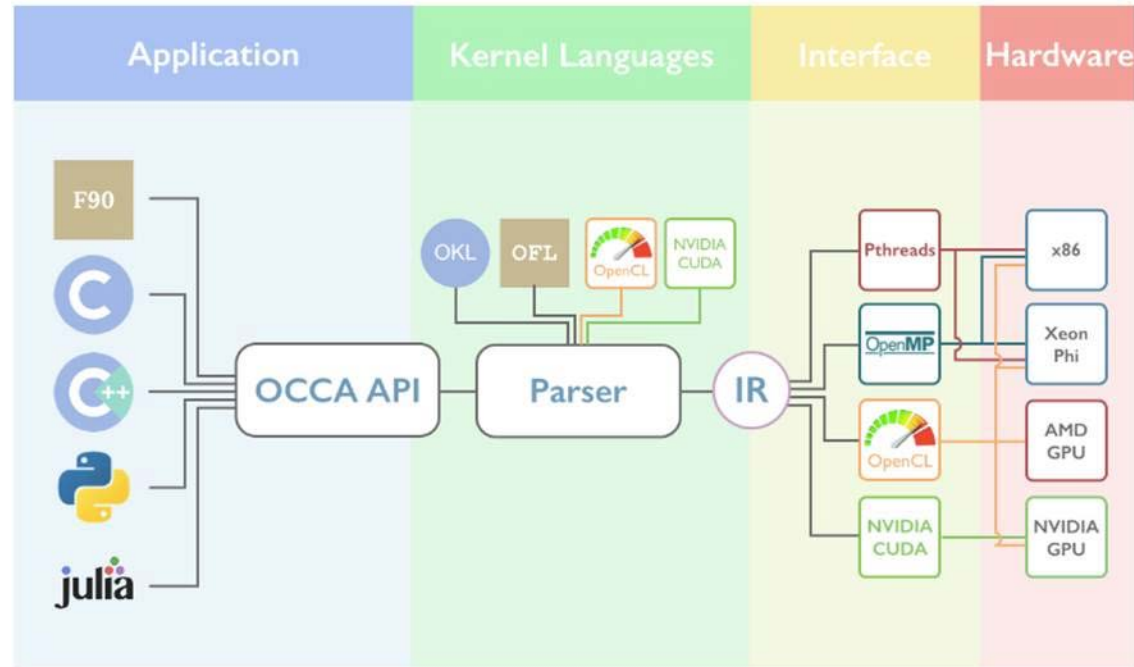
*Kokkos: <https://github.com/kokkos>

One option: Hardware-optimized code

- Allows greatest efficiency of machine use by code, which may be critical for realizing exascale performance
- Machine constraints (dimensionality of problem, bandwidth) still apply
- Requires extensive model redesign and re-coding for every hardware type, and hardware update

Another option: Hardware-Agnosticism

- For discussion, take OCCA as hardware-agnostic language (there are many other options).
- The computer model codes are written in a **language** of modeler's choice.
- Software engineers pick a **specific kernel language**.
- The **library interface** translates to the language best suited for the **hardware**.
- May be difficult to optimize multiple algorithms with specific computational characteristics



Earth System Modeling Requirements

- ***“HPC architectures are developing in the wrong direction for state-heavy, low computational intensity (CI) Earth system applications.” - ESPC HPC White Paper***

– Top500 (June 2018, <https://www.top500.org>):

Rank	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	Summit - IBM Power System AC922, IBM POWER9 22C 3.07GHz DOE/SC/Oak Ridge National Laboratory United States	2,282,544	122,300.0	187,659.3	8,806
			<i>Developed for 9 flop/byte application</i>		
2	Sunway TaihuLight - Sunway SW26010 260C 1.45GHz, NCRCP National Supercomputing Center in Wuxi China	10,649,600	93,014.6	125,435.9	15,371
			<i>Developed for 25 flop/byte application</i>		
3	Sierra - IBM Power System S922LC, IBM POWER9 22C 3.1GHz DOE/NNSA/Lawrence Livermore National Laboratory United States	1,572,480	71,610.0	119,193.6	--
			<i>Developed for 9 flop/byte application</i>		

– Exascale systems will require applications providing upwards of **50 flops/byte** [Goodacre, J., Manchester U., ECMWF Oct. 2016]

- Most computationally intense components in **today’s Earth system models rarely reach two operations per byte and typically run less than one operation per byte** over the full application. (Carman et al. 2017 <https://doi.org/10.7289/V5862DH3>)

Earth System Prediction Computing Needs

- Predict hazards at short time ranges *and* enable decision making in weather-to-climate overlap
 - Weather predictions:
 - Strict time requirements ($1 \text{ model day} \leq 8 \text{ min wall time}$)
 - Seasonal through decadal predictions:
 - Short run times for evaluation, development, reforecasting
- Future computing needs will exceed 1000 times of today's existing computing and possibly require custom built hardware & software
 - Need accurate forecasting of local floods at catchment level and to resolve hurricane structure/rainbands.
 - Significant investment will be needed to port our models to exascale systems.
- **White paper** (Carman, et al. "Position Paper on High Performance Computing Needs in Earth

System Prediction." National Earth System Prediction Capability (ESPC) program. April 2017.

<https://doi.org/10.7289/V5862DH3>