NASA Blazes a Different Path to Energy-Efficient Supercomputing

Overview

For years, NASA had a very straightforward process for replacing high-performance computing hardware: over a three-year period, when it became more expensive to operate an older suite of hardware than it did to replace it with new products that could accomplish the same work, we simply replaced the old hardware. For NASA's High-End Computing Capability (HECC) Project, that process changed when we reached the limits of our facility's power, cooling, and floor-loading capacity, becoming a strategy of decommissioning the least productive hardware and replacing it with more capable counterparts. The impact was that we provided our users with less supercomputing capability than we would have without the limitations.

Additionally, with 25% of our total power consumption going to cool our systems and 50,000 gallons of water per day being evaporated, we wanted a solution that would expand our compute facility while being sensitive to the impact on our environment.

Project Details

The facility housing HECC systems is located at NASA's Ames Research Center, near the southern end of the San Francisco Bay. The area's temperate climate opened up solutions that are well suited to this location and allow for extremely energy-efficient approaches. Working with NASA Headquarters, which commissioned a study to evaluate possible expansion solutions, we determined the best path forward was to deploy a prototype module system with a goal of reducing both power and water utilization. Using 25% of total power consumption for cooling equates to a 1.33 Power Usage Effectiveness (PUE) for the primary NASA Advanced Supercomputing (NAS) facility; our goal was to deploy a system with a 1.06 PUE and to all but eliminate water utilization.

We deployed the first prototype module in late 2016, and it has been operational for about a year. It houses a 1.2-petaflop system, called Electra, that uses about 400 kilowatts (kW) of power. The module exceeded our optimistic goals. Over the third quarter of 2017, the PUE for this system was 1.025, beating our target of 1.06; water usage for the entire year was about 55,000 gallons, a reduction of more than 97% over the primary facility.

We successfully used outside air to cool the compute system throughout the year, only running water for cooling when needed. The outside air is pulled into the module through large particle filters and enters a mixing chamber, where it can be mixed with air from the module's hot aisle to lower the humidity or raise the temperature. It then passes through small particle filters and a mesh media that can contain water to cool the air (if needed) before passing into the cold aisle and through the computers.

Results and Impact

The success of the first prototype led to the acquisition of a second, highlighting another advantage of this approach: expansion on demand. By this time, new technologies became available that were not an option with the first prototype. The second module can hold twice the number of nodes, and can remove more heat load—allowing us to increase total power from 500 kW to 1.2 megawatts. The new system uses the same air/water cooling technology as the first, but also has direct-to-chip heat transfer to a cold-water loop. Early testing is in line with the first module with a PUE value of 1.03, exceeding our goal of 1.08 for the new system. When fully populated, the new Electra system—comprising the nodes in both modules—will reach a peak performance of over 8 petaflops.

What's Next

Next year, we will begin construction on a new modular facility based on the technologies proven in the prototypes. This new facility will enable additional expansion to meet NASA's future challenges.

NASA Blazes A Different Path To Energy Efficient Computing

William Thigpen
Advanced Computing Branch Chief
NASA Advanced Supercomputing

NASA's HEC Requirements: Capacity

HEOMD (engineering-related work) require HEC resources that can handle large numbers of relatively-low CPU-count jobs with quick turnaround times.

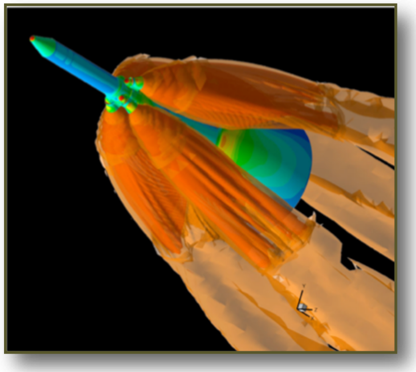

Over 1500 simulations utilized ~ 2 million processor hours to study launch abort systems on the next generation crew transport vehicle

The formation of vortex filaments and their roll-up into a single, prominent vortex at each tip on a Gulfstream aircraft

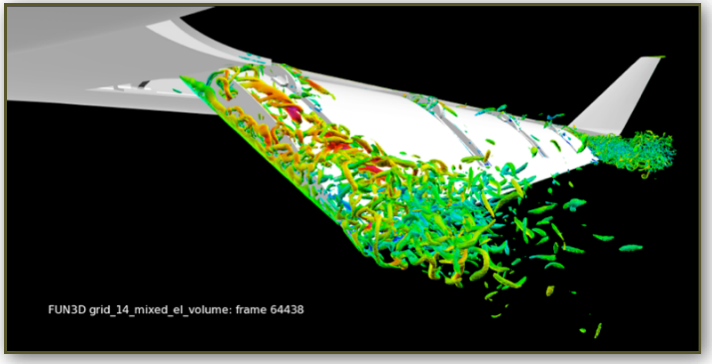

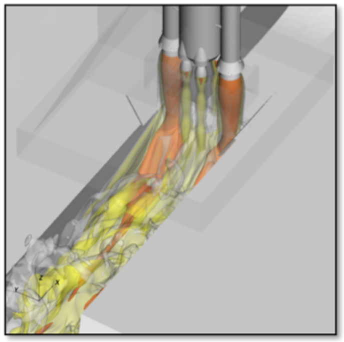

Over 4 million hours were used over a 4 month project to evaluate future designed of the next generation launch complex at the Kennedy Space Center

NASA's HEC Requirements: Capability

ARMD and SMD (aeronautics and science related work) require HEC resources that can handle high fidelity relatively-large CPU-count jobs with minimal time-to-solution. Capability enables work that wasn't possible on previous architectures.

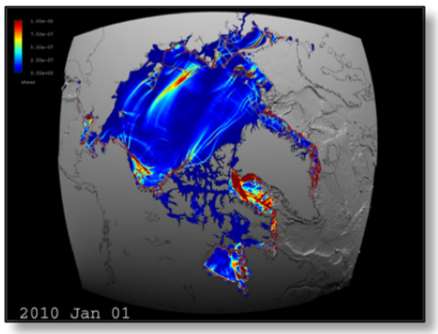

NASA is looking at the oceans, running 100's of jobs on Pleiades using up to 10,000 processors. Looking at the role of the oceans in the global carbon cycle is enabled by access to large processing and storage assets

For the first time, the Figure-of-Merit has been predicted within experimental error for the V22 Osprey and Black Hawk helicopter rotors in hover, over a wide range of flow conditions

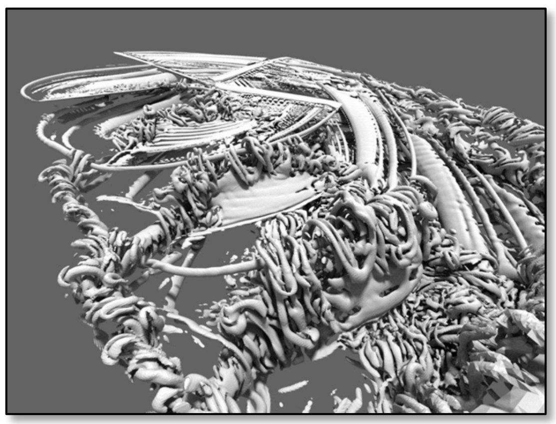

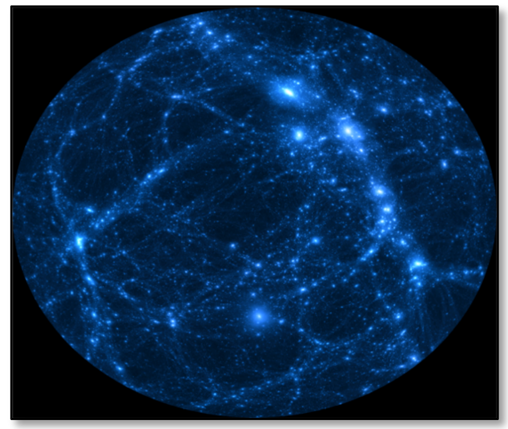

To complete the Bolshoi simulation, which traces how the largest galaxies and galaxy structures in the universe were formed billions of years ago, astrophysicists ran their code for 18 straight days, consuming millions of hours of computer time, and generating massive amounts of data

NASA's HEC Requirements: Time Critical

NASA also has need for HEC resources that can handle time-sensitive mission-critical applications on demand (maintain readiness)

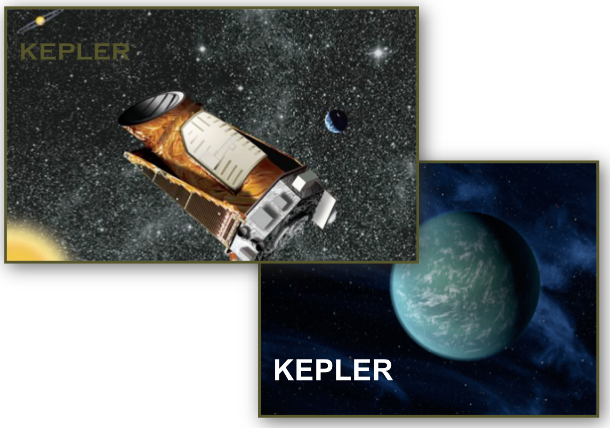

HECC enables the enormous planetary transit searches to be completed in less than a day, as opposed to more than a month on the Kepler SOC systems, with significantly improved accuracy and effectiveness of the software pipeline

ReEntry

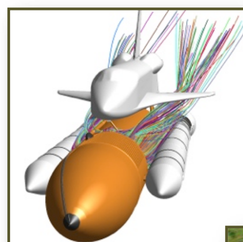

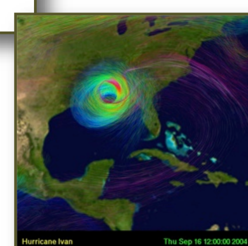

Storm Prediction

UAVSAR produces polarimetric (PolSAR) and interferometric (repeat-pass InSAR) data that highlight different features and show changes in the Earth over time

HECC Assets

HECC Traditional Computer Floors

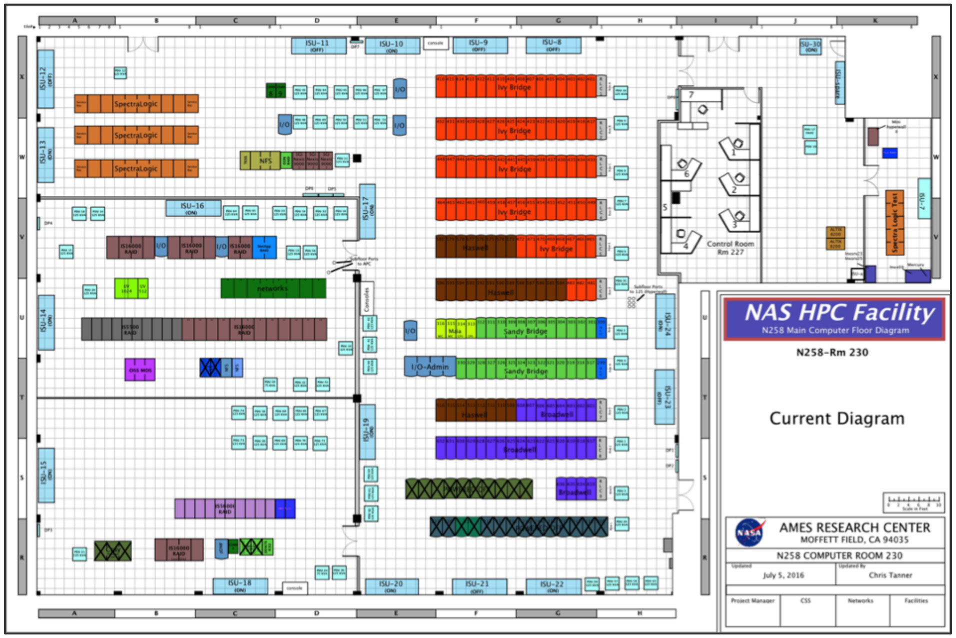

230
1,560 m²

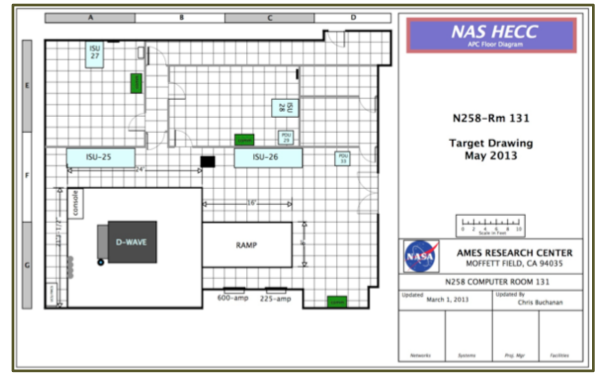

131
193 m²

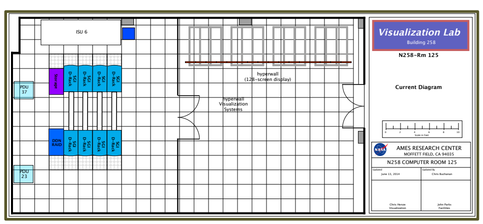

125
118 m²

190
640 m²

189
250 m²

HECC Modular Computer Floors

- **HECC Prototype Facilities**
- **Modular Supercomputer Facility**

Concrete Pad	**362 square meters / 2.4 MW / will hold 2 adjacent DCoD-20 modules**
• DCoD-20 Module 1	90 square meters / 40 square meters computer floor / 500 KW
• Custom Module 2	90 square meters / 86 square meters computer floor / 1,200 KW

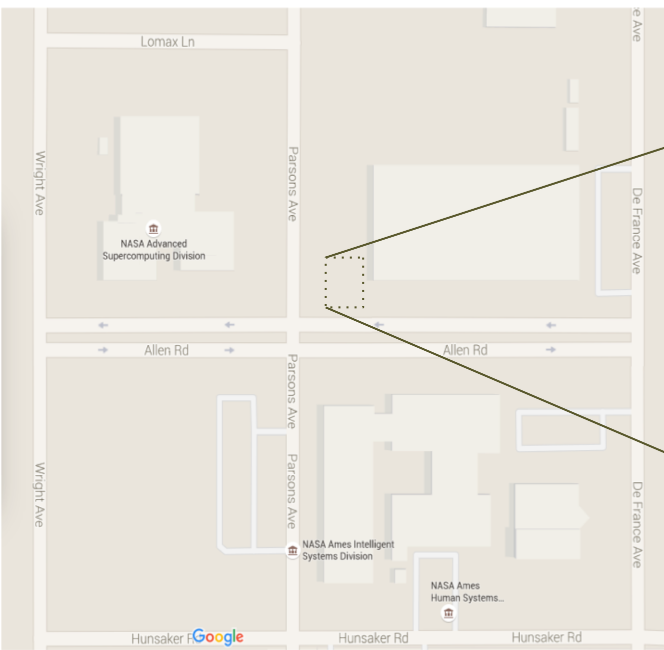

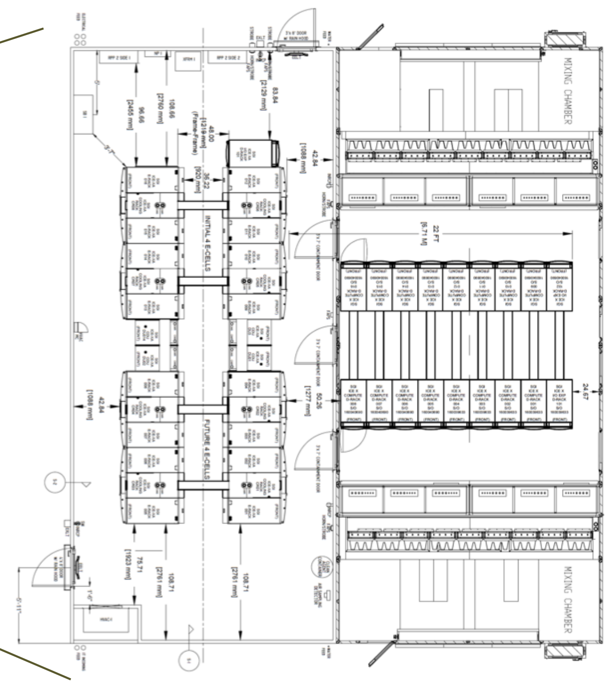

HECC Compute Portfolio

HECC Assets

4 Compute Clusters

• Pleiades	161 Racks / 11,340 nodes / 7.57 PF / 32,230 SBU/hr
• Electra	20 Racks / 2,304 nodes / 4.78 PF / 11,566 SBU/hr
• Merope	56 ½ Racks / 1,792 nodes / 252 TF / 1,792 SBU/hr
• Endeavour	3 Racks / 2 nodes / 32 TF / 140 SBU/hr
1 Visualization Cluster	245 million pixel display / 128 node / 703 TF
7 Lustre File Systems	39.6 PB
6 NFS File Systems	1.5 PB
Archive System	490 PB

Experimental Quantum D-Wave 2
- System with 1,097 qubits

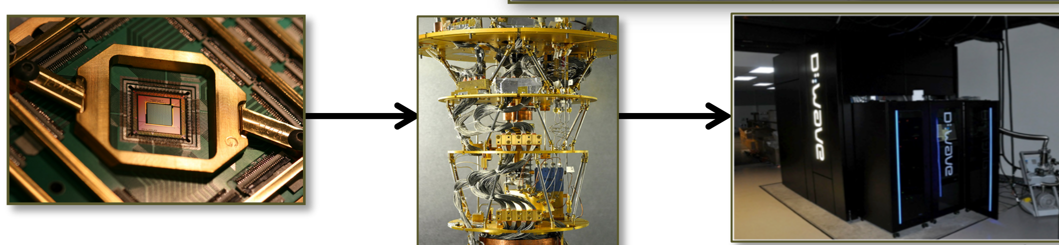

HECC Growth

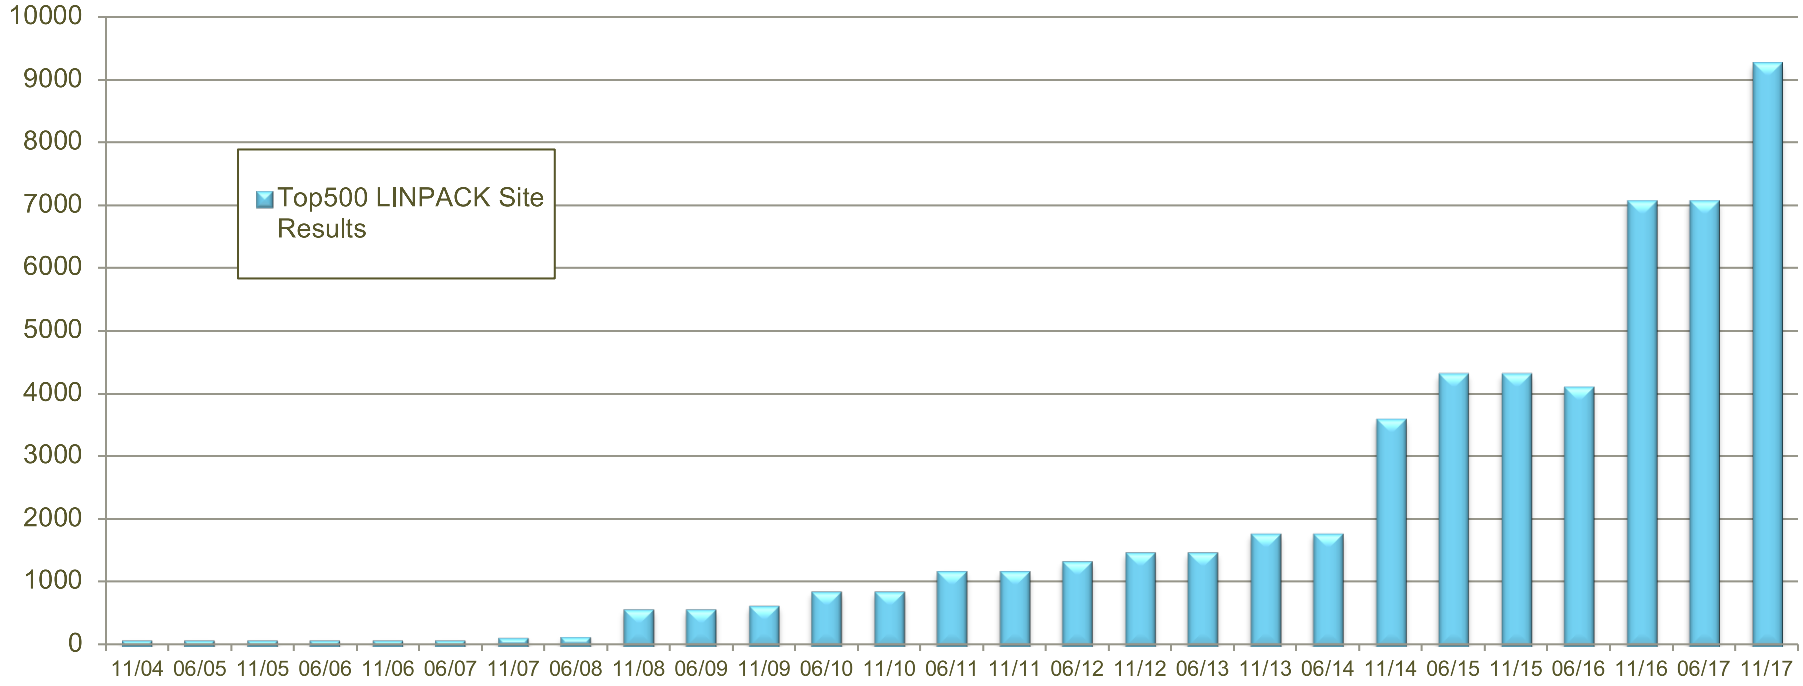

10000
9000
8000
7000
6000
5000
4000
3000
2000
1000
0
Top500 LINPACK Site Results
11/04 06/05 11/05 06/06 11/06 06/07 11/07 06/08 11/08 06/09 11/09 06/10 11/10 06/11 11/11 06/12 11/12 06/13 11/13 06/14 11/14 06/15 11/15 06/16 11/16 06/17 11/17

HECC Conducts Work in Four Major Technical Areas

Supercomputing Systems

Provide computational power, mass storage, and user-friendly runtime environment through continuous development of management tools, IT security, systems engineering

Application Performance and User Productivity

Facilitate advances in science and engineering for NASA programs by enhancing user productivity and code performance of high-end computing applications of interest

Data Analysis and Visualization

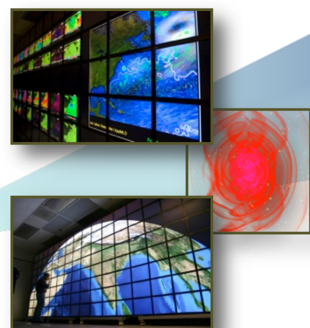

Create functional data analysis and visualization software to enhance engineering decision support and scientific discovery by incorporating advanced visualization technologies

Networking

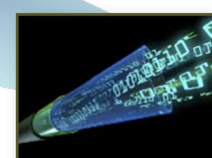

Provide end-to-end high-performance networking analysis and support to meet massive modeling and simulation distribution and access requirements of geographically dispersed users

Supporting Tasks

- **Facility, Plant Engineering, and Operations:** Necessary engineering and facility support to ensure the safety of HECC assets and staff
- **Information Technology Security:** Provide management, operation, monitoring, and safeguards to protect information and IT assets
- **User Services:** Account management and reporting, system monitoring and operations, first-tier 24x7 support
- **Internal Operations:** NASA Division activities hat support and enhance the HECC Project areas

Resource Utilization

User Location 2017

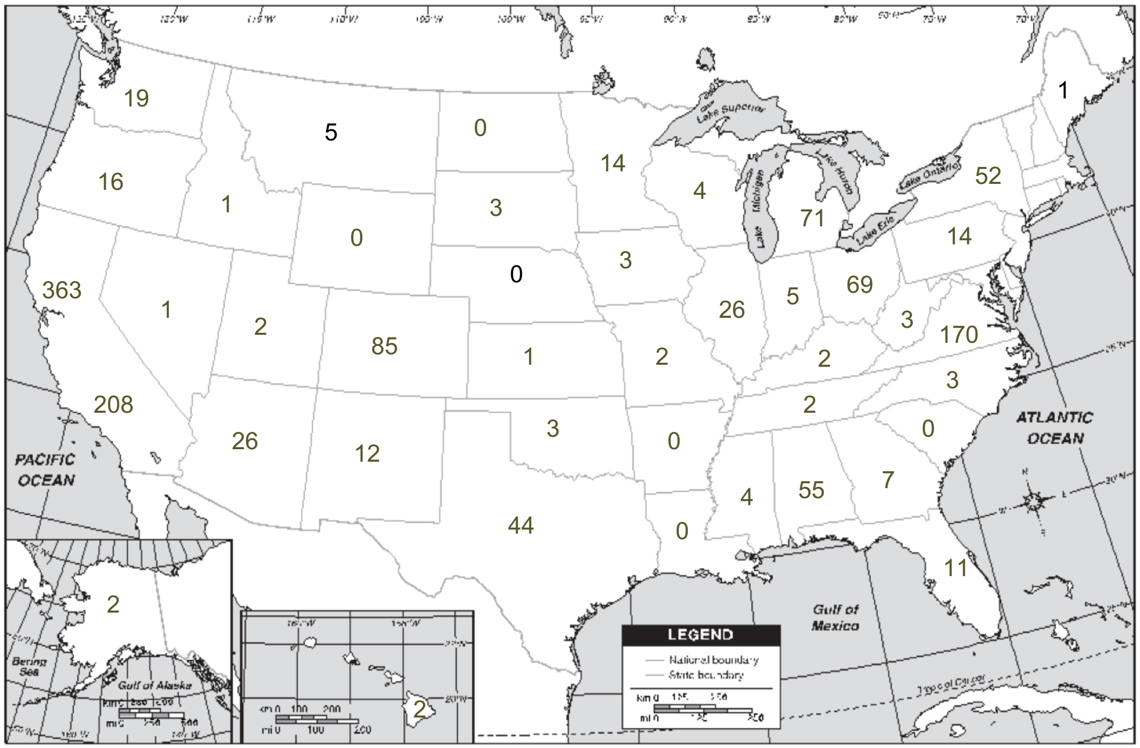

CT – 5
DE – 3
HI – 2
MA – 69
MD – 155
NH – 5
NJ – 26
RI – 2
VT – 1

Offshore – 68

Quarterly Utilization Over 10+ Years

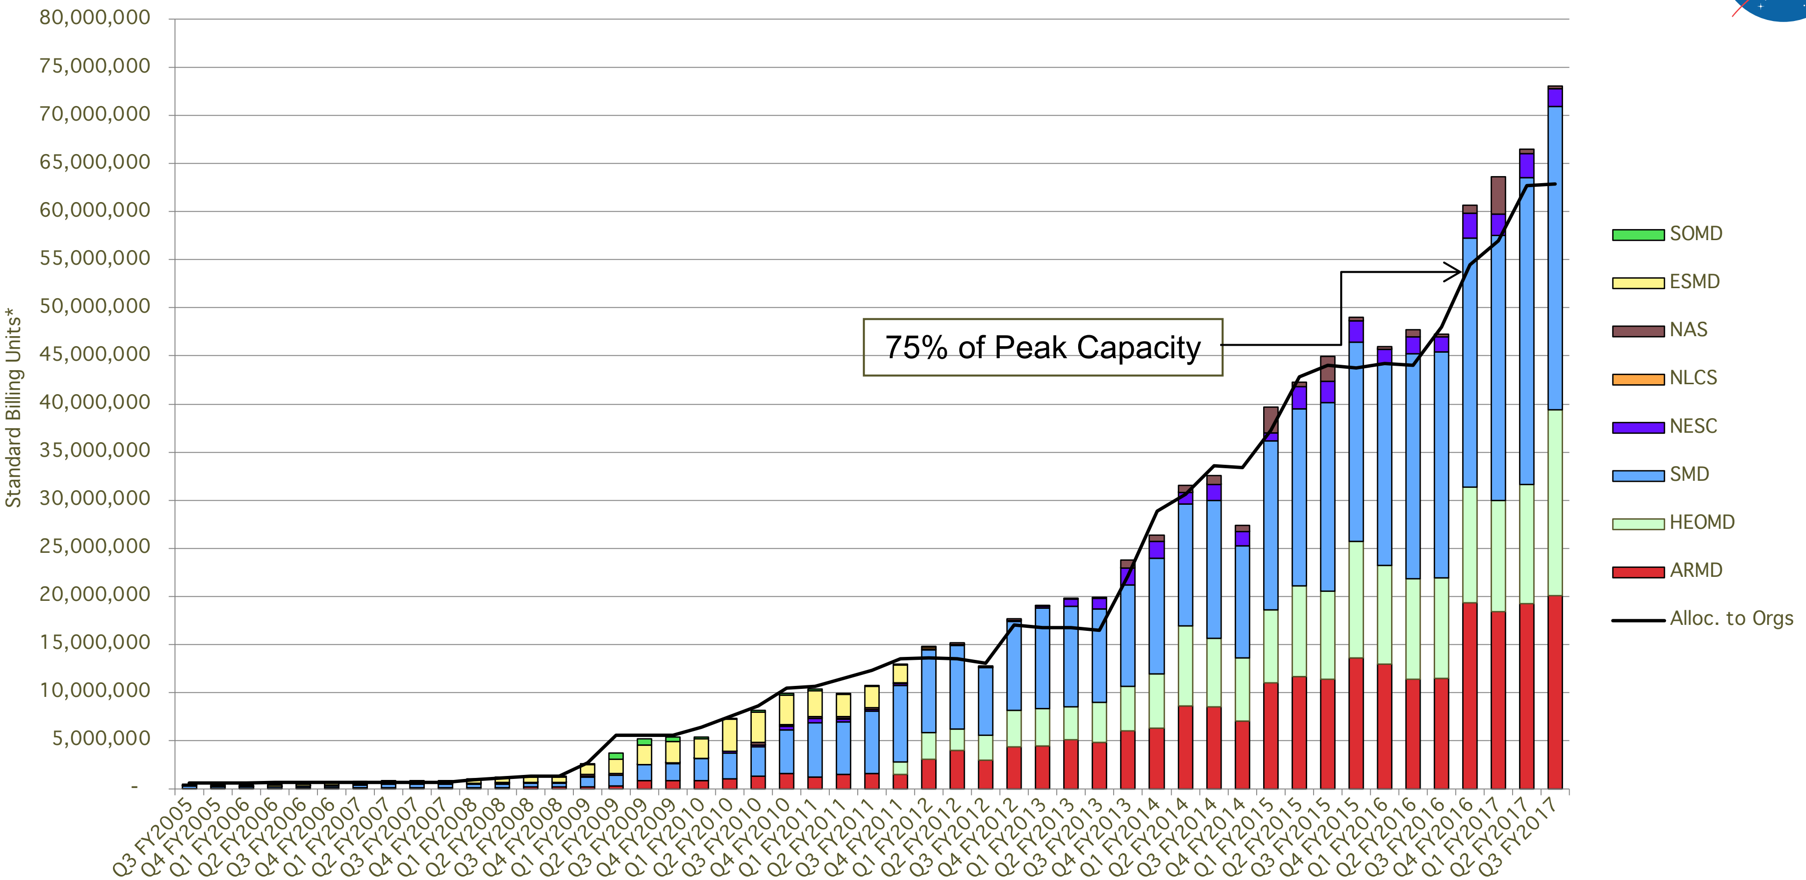

Aeronautics Support (55,958,567 SBUs)

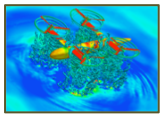

Advanced Air Vehicles
# of projects: 74
# of SBUs used*: 21,212,681

✦HECC is used to develop concepts and technologies for dramatic improvements in the noise, emissions, and performance of transport aircraft.
✦HECC is used to develop concepts and technologies to increase rotorcraft speed, range and payload, and decrease noise, vibration and emissions.
✦HECC is used to develop advanced computer-based prediction methods for supersonic aircraft shape and performance and to develop technologies that will help eliminate today's technical barriers (such as sonic booms) to practical, commercial supersonic flight.
✦ HECC is used to develop computer-based tools and models and scientific knowledge that will lead to significant advances in our ability to understand and predict flight performance for a wide variety of air vehicles.

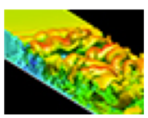

Transformative Aeronautics Concepts
# of projects: 48
# of SBUs used*: 40,017,897

✦HECC is used to develop and utilize Reynolds-averaged Navier-Stokes (RANS) and Large Eddy Simulation (LES) methods, and hybrid RANS-LES techniques to improve calculation methods for propulsion flows dominated by turbulent boundary layers and mixing.
✦HECC is used to assess natural laminar flow concepts, to elucidate the physics and control of boundary layer transition in swept wing flows and drag reduction concepts for compressible boundary layers.
✦HECC is used to validate chemistry, chemistry-turbulence and spray models being developed under the National Jet Fuels Combustion program.

Airspace Operations and Safety
# of projects: 3
# of SBUs used*: 2,173,135

✦HECC is used for developing reliable computational tools for predicting and analyzing stability & control characteristics of aircraft prior to or while encountering loss-of-control flight conditions characterized by abnormal flight (e.g., stall), abnormal vehicle conditions (e.g., damage, jammed control surfaces), external upsets (e.g., wake vortex, wind shear, gusts), and icing.
✦ HECC is used to develop methods for computing aerodynamic performance degradation associated with ice accretions on swept wing geometries.
✦HECC is used to produce real-time icing impact fields for flight planning and post mission analysis.

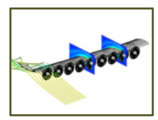

Integrated Aviation Systems
# of projects: 9
# of SBUs used*: 14,004,676

✦HECC is used for accurate prediction of airframe noise from a full scale aircraft and evaluation of flap and landing gear noise reduction concepts in flight environments.
✦HECC is used to develop technology for compact, high-power-density electric motors to power an all-electric general-aviation aircraft or helicopter, a hybrid turbine-electric regional airliner or a large transport with many small engines distributed around the aircraft.
✦HECC used for parametric studies conducted to optimize size, shape and placement of an array of fluidic actuators for maximizing the lift for control surfaces on an aircraft, which would help reduce the size of control surface and the weight of an aircraft.

*October 1, 2016 to September 30, 2017

Human Exploration and Operations & Safety Support (51,658,239 SBUs)

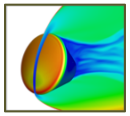

Multi-Purpose Crew Vehicle
# of projects: 8
# of SBUs used*: 3,563,459

✦HECC is used to support the creation of hundreds of computational solutions that model the flow field around the Crew Module and Launch Abort System for all flight regimes to be used as input for the aerodynamic databases.
✦HECC is used to run computational fluid dynamics simulations to study the aerodynamic and aerothermal environments for the Multi-Purpose Crew Vehicle.
✦HECC is used to develop and deploy a prototype system for rapid aerodynamic performance database generation and to use it on real-world problems faced by the Human Exploration and Operations mission directorate.

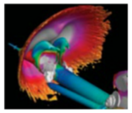

Space Launch Systems
# of projects: 14

# of SBUs used*: 45,945,964

✦ HECC is used for computational fluid dynamics simulations of Space Launch Systems ascent to assess aerodynamic performance, protuberances, stage separation, and plume effects (such as plume-induced flow separation) for evolving vehicle designs.
✦ HECC is used for computational fluid dynamics analysis of Advanced Booster development efforts in the combustion stability areas.
✦HECC is used for prediction of the launch induced environment for the Space Launch System including liftoff acoustics, ignition over-pressure, separation environments, debris, Launch Pad Abort Environments and hydrogen entrapment.
✦HECC is used to simulate tanks and main propulsion system components (including feedlines, valves, manifolds, ducts, and pogo accumulators) for evaluation of criteria such as flow uniformity and component pressure drop.

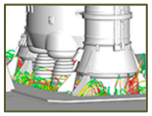

HEOMD - Space Flight Operation & General
# of projects: 37
# of SBUs used*: 11,895,099

✦HECC is used to simulate the effect of larger solid rocket boosters and new propulsion systems on the launch facility at Kennedy Space Center, such as investigating whether ignition overpressure waves generated during liftoff are fully suppressed by the existing water suppression system.
✦HECC is used to evaluate visiting-vehicle induced loads on the International Space Station (ISS) during mated and rendezvous operations and to evaluate crew Extra-vehicular Activity/Intra-vehicular Activity and attitude control loads on ISS.
✦HECC is used in developing a combustion response model to investigate combustion instability in hydrocarbon-fueled rocket engines.
✦HECC is used for technology development for entry, descent and landing systems.

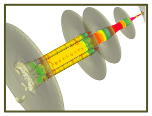

NASA Engineering & Safety Center
# of projects: 7
# of SBUs used*: 7,425,622

✦HECC is used for simulations to provide guidance to the Space Launch System advanced booster designers by providing aerodynamic loading implications for various potential advanced booster geometric configurations.
✦HECC is used to improve the capability to predict combustion stability in liquid rocket engines to increase NASA engineers' capability to more confidently and efficiently identify and mitigate combustion stability issues in engine development programs.
✦HECC is used to used for studies of large eddy simulations of oblique-shock / supersonic hot jet interaction, aimed at prediction of plume-induced vibroacoustics.

*October 1, 2016 to September 30, 2017

Science Support (94,488,707 SBUs)

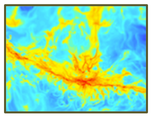

Astrophysics
# of projects: 105
# of SBUs used*: 48,143,291

✦ HECC is used by the Kepler mission to find Earth-sized planets around other stars and to fully analyze the Kepler data to find any undiscovered planets still "hiding" in the data.
✦HECC is used to understand the physics of high redshift galaxy formation and make detailed predictions that can be used to guide NASA observations of the first galaxies.
✦HECC is used for quantifying the redistribution of matter in galaxies when supernova energy is deposited; exploring the growth of black holes and the impact of active galactic nuclei on galaxy evolution; and determining whether the ultraviolet light from stars in galaxies can "escape" to re-ionize the universe.

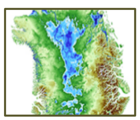

Earth Science
# of projects: 126
# of SBUs used*: 23,309,412

✦HECC is used to combine observational data with numerical simulations of the global ocean circulation to provide vital information for understanding climate change and its impact on land and sea ice, ocean ecology, and the global carbon cycle.
✦ HECC is used for high-resolution cloud resolving model simulations to provide unique and detailed insights into the processes that form tropical clouds and cloud systems, which account for approximately two-thirds of global rainfall.
✦ HECC is used explore the feedback mechanisms between polar ice sheets and atmosphere circulation in order to determine how global temperature changes translate into increased sea level rise.
✦HECC is used to improve the understanding of the current balance of carbon in the Arctic and to provide a framework for early detection of future carbon destabilization.

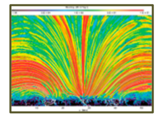

Heliophysics
# of projects: 92
# of SBUs used*: 24,538,939

✦ HECC is used for modeling solar magneto-convection in order to understand how magnetic fields emerge through the sun's surface, heat the sun's outer atmosphere, and produce sunspots, spicules, and flares.
✦HECC is used for realistic multi-scale simulations to understand the complicated physics of the turbulent convection zone and atmosphere of the sun and for analyzing and interpreting observations from the NASA space missions.
✦HECC is used to simulate small-scale magnetic fields generated by turbulent dynamo action just beneath the solar surface in order to accurately predict space weather events that impact the Earth environment.

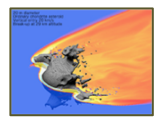

Planetary Science
# of projects: 87
# of SBUs used*: 24,115,485

✦HECC is used to decipher the structure of the lunar interior to understand the origin and thermal evolution of the moon and to extend this knowledge to other bodies in the inner solar system.
✦HECC is used to model the origin and evolution of Kuiper belt objects to determine how their properties constrain our current models of planet formation.
✦HECC is used perform modeling and simulation of asteroid entry, breakup, airburst, blast propagation, and tsunamis to assess the risks that potentially hazardous asteroids could pose to populations and infrastructure in the event of an Earth strike.

*October 1, 2016 to September 30, 2017

Return on Investment

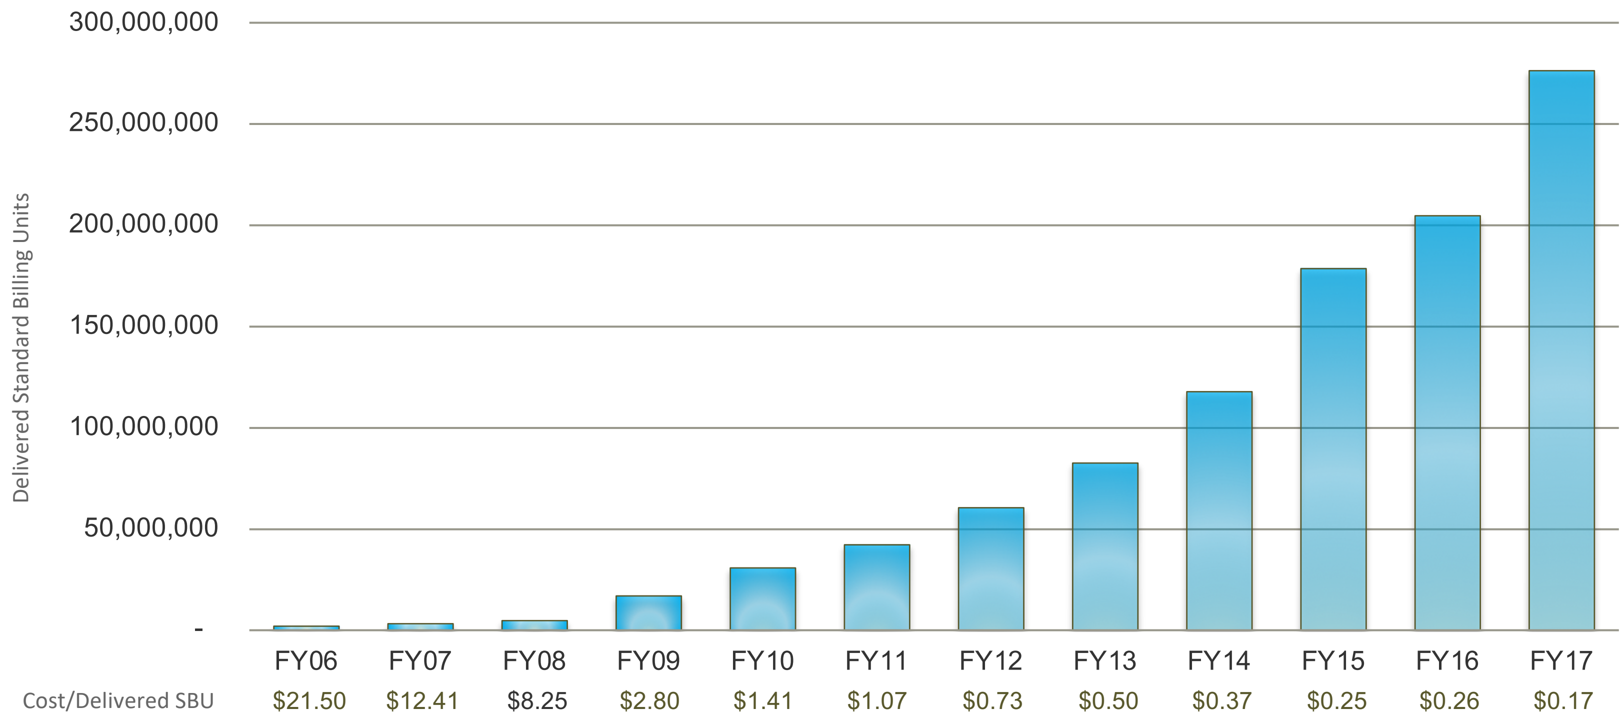

NAS Facility Expansion

Why are We Doing This

- **The calculation used to be very simple…**
  - When the cost of maintaining a group of nodes for three years exceeded the cost to replace those nodes with fewer nodes that did the same work, we replaced them.
- **Now, not so much…**
  - We look at the total computing our users get and procure new nodes within our budget and remove enough nodes to power and cool the new nodes.
  - This means that we are not able to actually realize all of the expansion we are paying for.

But That’s Not All

- **Our computer floor is limited by power and cooling**
- **Our Current Cooling System**
  - Open Air Cooling Tower with 4 50HP pumps
  - 4 450 Ton Chillers
  - 7 pumps for outbound chilled water
  - 4 pumps for inbound warm water
- **Our Electrical System**
  - Nominally the facility is limited to 6MW
  - 20% - 30% is used for cooling
  - 4MW – 5MW for computing

N258 Cooling Flow Chart

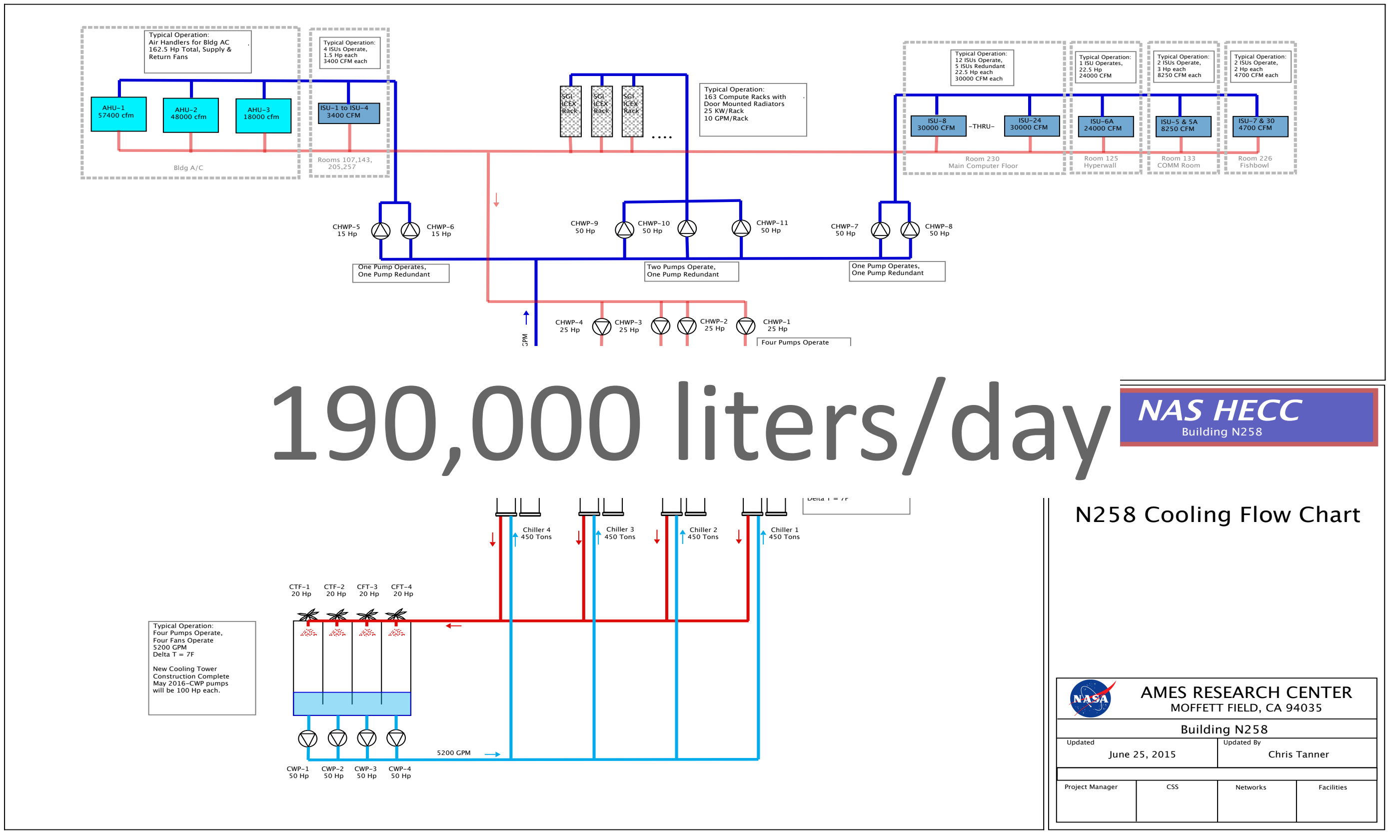

190,000 liters/day

What Could We Have Done?

- Do nothing
- Augment the existing facility
- Upgrade a different facility on the center
- Build a new traditional center using latest technologies
- Build a modular center using latest technologies
- Expand at a remote NASA computer center
- Expand at a remote Government computer center
- Expand at a private sector computer center

What We Did

- **Based on a NASA HQ study conducted by IDC in 2015 we**
  - Engineered and deployed a 2 phase prototype modular facility
  - Completed initial design of NAS Facility Expansion

PHASE 1 OF THE PROTOTYPE EXCEEDED EXPECTATIONS

Site Layout

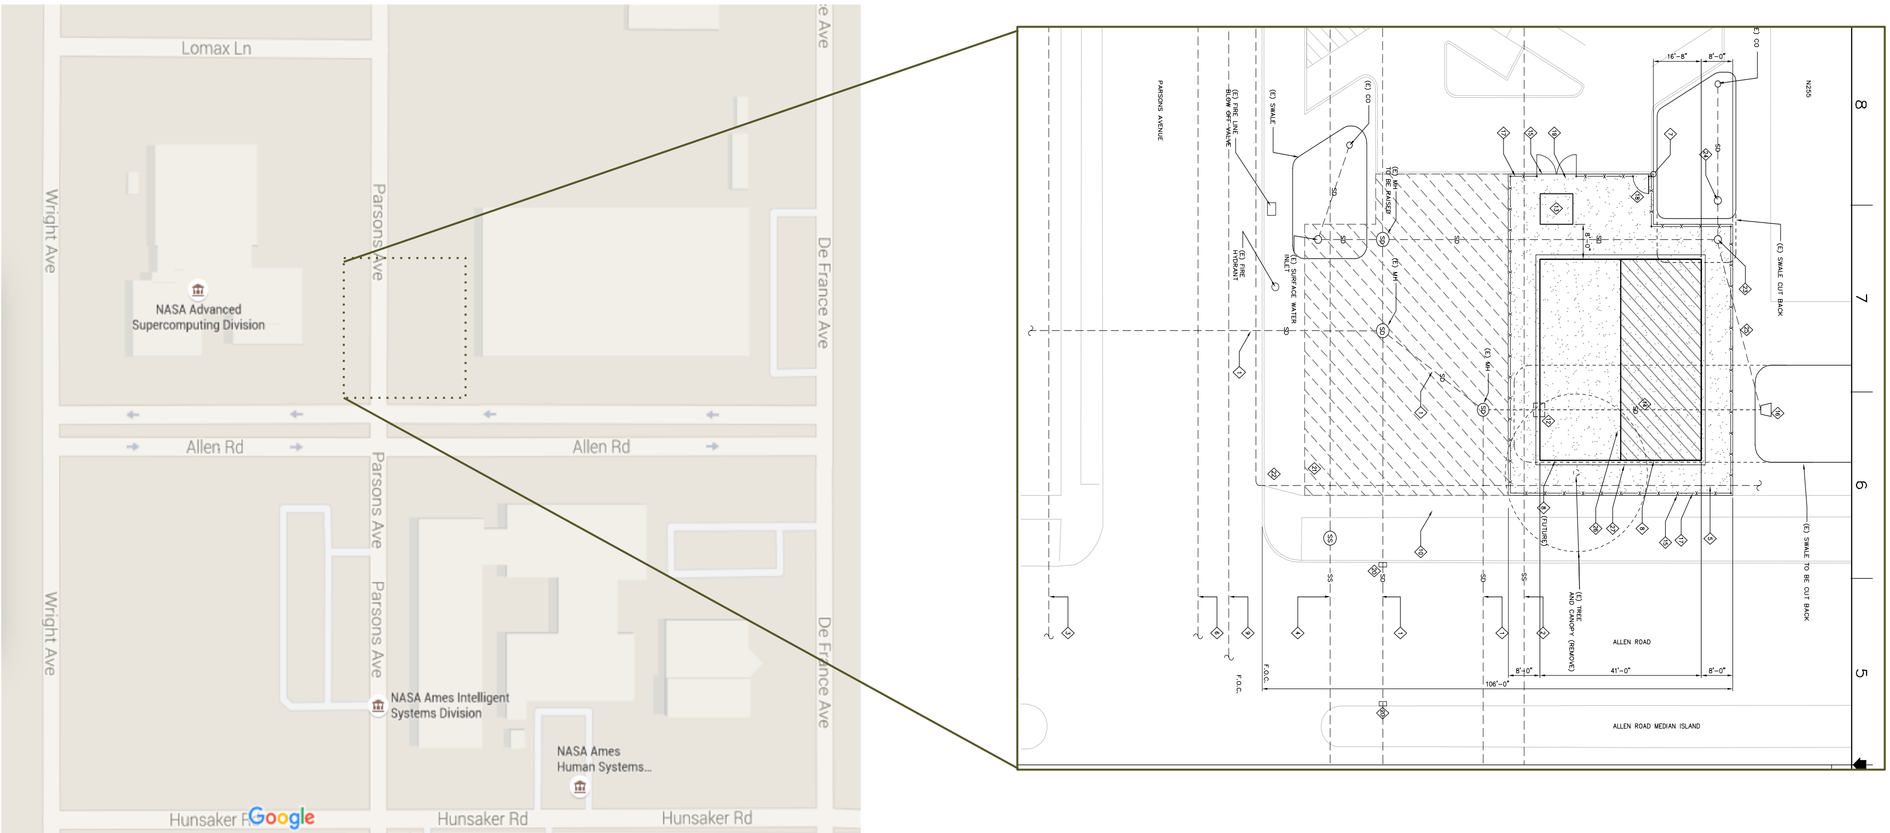
Lomax Ln
Wright Ave
Parsons Ave
De France Ave
NASA Advanced Supercomputing Division
Allen Rd
NASA Ames Intelligent Systems Division
NASA Ames Human Systems...
Hunsaker Rd
Google
PARSONS AVENUE
ALLEN ROAD
ALLEN ROAD MEDIAN ISLAND

DCU-20

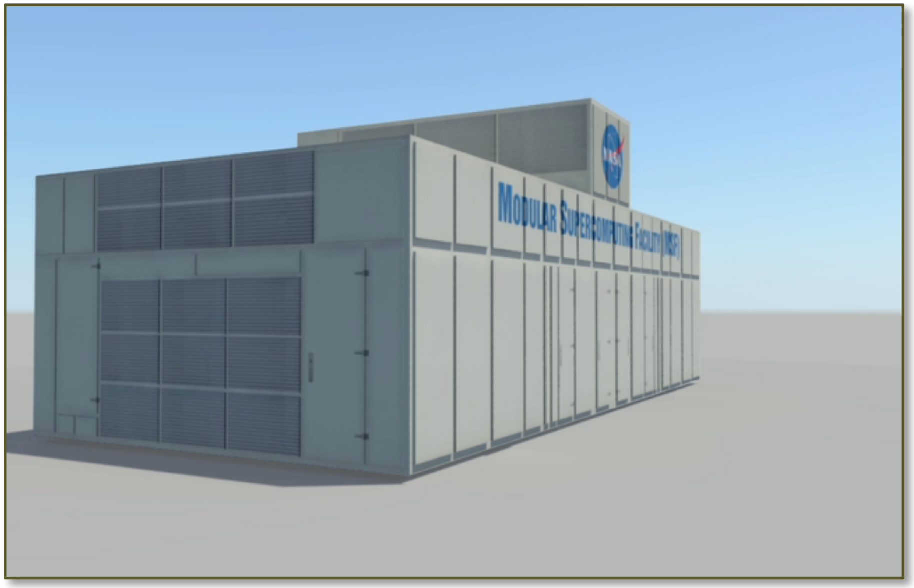

Concrete Pour

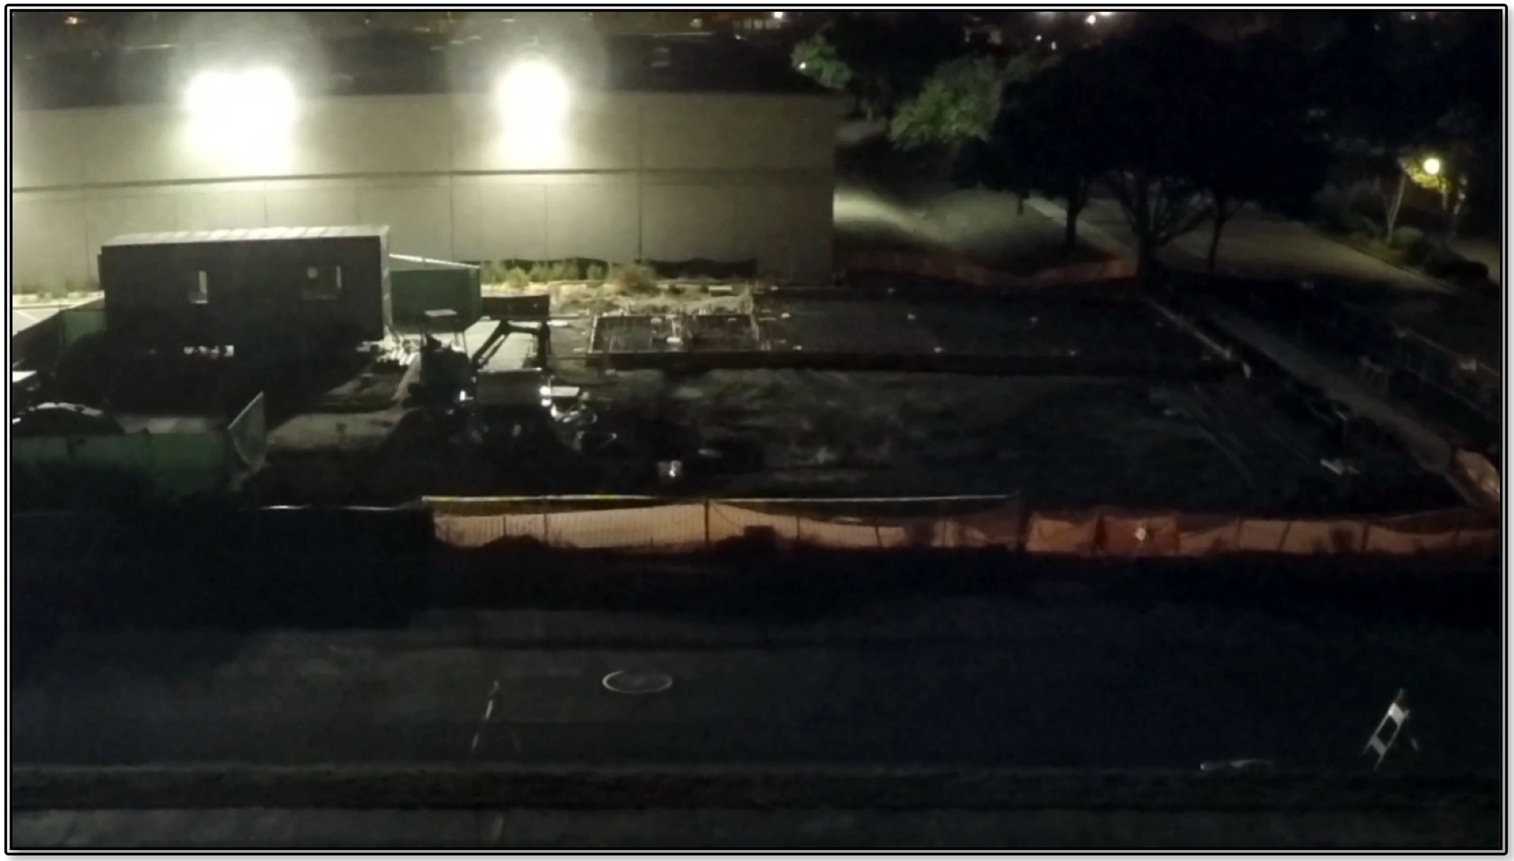

Module 1 Assembly

Annual Energy Impact

16 Computer Racks (1152 Nodes)	Existing Facility	DCoD-20 Facility	% Savings
Water Utilization Per Year	5,527,000 L*	208 L**	96%
Electricity per year	3,728,256 kwh°	2,873,280 kwh°°	22.9%(overall) 92.4%(cooling)

* Assumes 16 racks represent 8% of facility load
** Year 1 usage
° 1.33 PUE (3rd quarter actuals)
° ° 1.025 PUE (3rd quarter actuals)

PHASE 2 OF THE PROTOTYPE IS DEPLOYED

Technologies Advanced in a Year

- **Rack density improved, doubling from D-Racks to E-Cells**
- **Chip performance and memory improved**
- **Facility advances allow us to deploy over a megawatt of computing in a single module effectively doubling the number of nodes we can deploy in a given footprint**
- **The second prototype can host eight HPE E-Cells with 2,304 nodes**

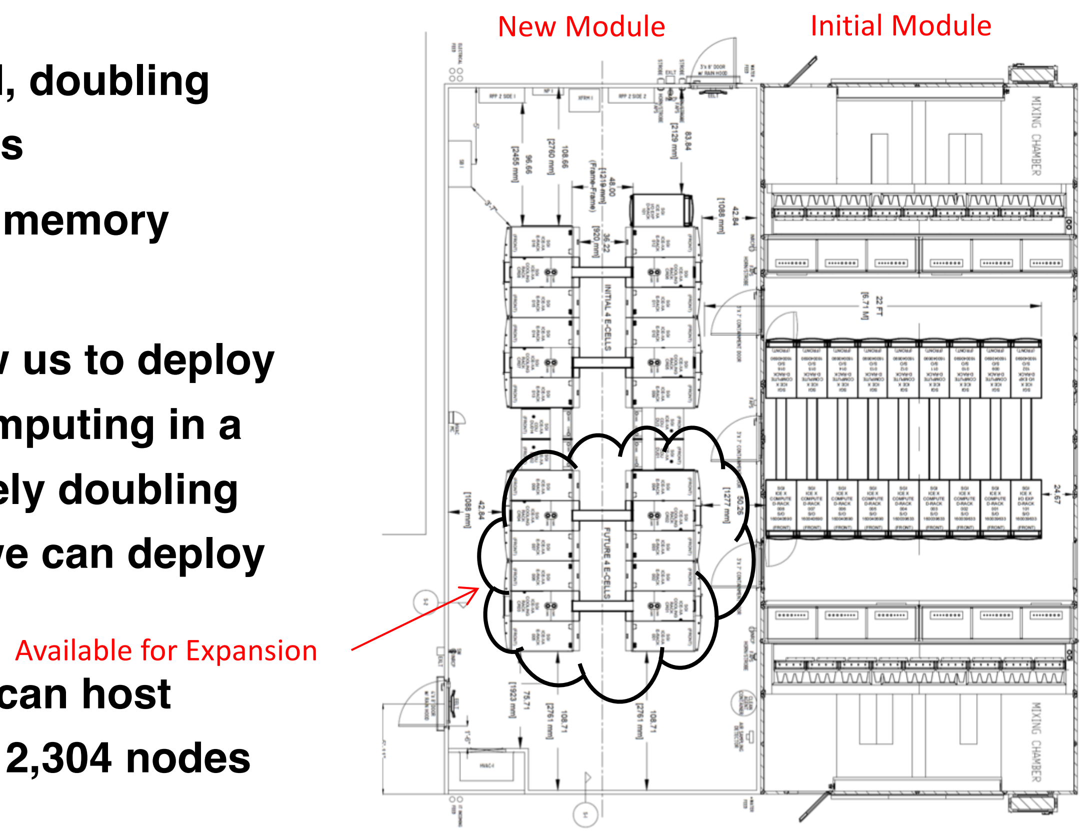

Module 2 Assembly

NAS FACILITY EXPANSION FY18 ->

We're Approved to Deploy a Facility Expansion

- **Our goal is to provide the infrastructure for a site that could deploy a fully-connected system with 1,000,000 cores.**
- **It needs to be flexible to handle advances in technology**
- **It needs to be expandable**
- **It needs to be energy efficient**

Full Site Deployment Concept

- 8 Compute modules house 96 tightly coupled E-Cells providing 84.9 PF
- 5 Data modules house 420 PB of formatted storage protected by dual generators and battery UPS
- System joins existing HECC assets with shared file systems and data archive
- Project deployed on site currently being constructed and available in early FY19
- Project fully operational in FY19

Site Location

Prepared Site

- 250 ft x 180 ft
- 3 ½ ft of Vertical, Engineered Fill (1½ ft Above DeFrance Road)
- Site Surrounded by De France Rd or Fire Access Road
- Ramp to Top of Elevated Site from DeFrance Road
- 25 kV Switchgear yard at Southwest Corner of Site
- Water Main Point of Connection at Southwest Corner of Site

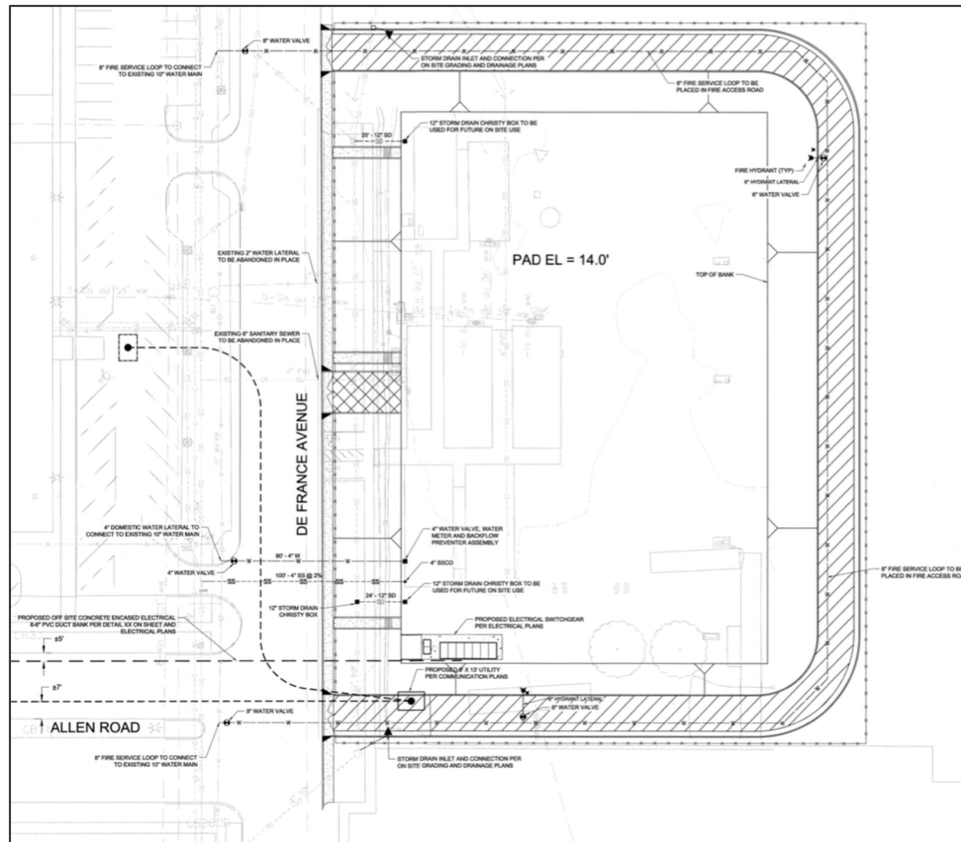

Prepared Site Utilities

Electrical at Site

- 25 kV Switchgear Yard (40' x 12')
  - Four 25 kV Vacuum Circuit Breakers will Distribute up to 15 MVA at 24.9 kV to Step-Down Transformers used in Improved Site
  - Power Meters installed on each Vacuum Circuit Breaker
  - Site Low Voltage Power
    - 1 Additional 25 kV VCB for Site Power
    - 150 kVA Transformer, 24.9kV/208V, 3 Phase, 4 wire
    - 400 A Panelboard

Water at Site

- 4-inch Water Main capable of 200 GPM at 40 psi at Point of Connection
  - RPZ Backflow Preventer & Water Meter Installed in 4-inch Water Line
- Sewer & Storm Drain Piping Installed to edge of Prepared Site

Communications at Site

- Data to N258 will be Provided by Conduits & Manholes
- Communication Conduits will Terminate at Comm Manhole in Center of Prepared Site
- Fiber Optic Procurement & Installation by NASA personnel

Questions

http://www.nas.nasa.gov/hecc