

Baseline Configuration

Establishing Consistency Across DoD HPC Centers

**IDC HPC User Forum
19 September 2012**

Overview

- **What Is “Baseline Configuration?”**
- **Why Is It Important?**
- **How Has It Evolved?**
- **How Does It Work?**
- **Where Do The Projects Come From?**
- **Accomplishments**
 - Existing Policies
 - What the Team Thinks
- **Summary**



Overview

- **What Is “Baseline Configuration?”**
- **Why Is It Important?**
- **How Has It Evolved?**
- **How Does It Work?**
- **Where Do The Projects Come From?**
- **Accomplishments**
 - Existing Policies
 - What the Team Thinks
- **Summary**



DoD Supercomputing Resource Centers (DSRCs)

Five Large HPC Centers



**US Air Force Research
Laboratory DSRC**



**US Army Research
Laboratory DSRC**



**US Engineer Research and
Development Center DSRC**



**Maui High Performance
Computing Center DSRC**



NAVY DSRC

The Program supports five DoD Supercomputing Resource Centers. These capabilities provide DoD's scientists and engineers with the resources necessary to solve the most demanding computational problems. Each Center is a robust complement of HPC capabilities that include: large-scale HPC systems, high-speed networking, multi-petabyte archival mass storage systems and customer support services. Additionally, a Consolidated Customer Assistance Center and two Data Analysis and Assessment Centers offer services to the DoD HPC user community. The collective set of centers:

- Support some of the world's largest computational projects
- Address expanding DoD requirements for HPC capability
- Deliver 1.075 billion processor hours of computing power in FY11

Baseline Configuration – What Is It?

- Enable **Users** to easily move between the **five DoD HPCMP centers** without learning and adapting to site-specific idiosyncrasies
- Define a common set of capabilities and functions (policies) for HPCMP systems and Centers **to facilitate User flexibility**
 - Initial focus on unclassified allocated systems and center documentation and processes
 - Goes beyond least common denominator or “standardization”
 - Provides for competition, innovation, and unique user capability
- Develop documents and toolsets to verify compliance
- Work with **Users** and HPCMP established teams to be an agent for positive impact
 - Discover, Clarify, and Resolve
 - Meticulous in communication with users on issues raised



Baseline Configuration – What Is It?

- **Guiding Principles**

- It is about user productivity
- It is about resource flexibility
- It must embrace unique features of the centers
- It must support unique workflows of seasoned users
- It must support novice user tools (i.e. ezHPC – web based GUI for HPC)
- “A foolish consistency is the hobgoblin of little minds...” - Emerson
- Leverage Lockheed Martin Technical Services staff – build upon the enterprise culture

Commonality is Our Goal
User Productivity is Our Passion



Overview

- What Is “Baseline Configuration?”
- **Why Is It Important?**
- How Has It Evolved?
- How Does It Work?
- Where Do The Projects Come From?
- Accomplishments
 - Existing Policies
 - What the Team Thinks
- Summary



Why Is It Important?

- Offers significant flexibility in distributing allocations across and within the services (Army, Navy, Air Force, Others)
- Creates a connection with major DoD HPC customers
 - Ombudsmen type role
- Establishes an enterprise culture within the DoD HPC Modernization Program
 - As opposed to disparate centers
- Documents the level of consistency required of centers
- Exposes all centers to the day-to-day challenges faced by our customers



Overview

- What Is “Baseline Configuration?”
- Why Is It Important?
- **How Has It Evolved?**
- How Does It Work?
- Where Do The Projects Come From?
- Accomplishments
 - Existing Policies
 - What the Team Thinks
- Summary



How Has It Evolved?

- **BC started seven years ago**
 - Establish a baseline requirement for all systems and centers
 - Working with HPC manufacturers would not work
- **Establish policies to enhance user experience at all centers**
 - Required and documented compliance levels
 - Semi-annual meetings with invited guests
 - Weekly telecons
 - Worked on turning projects into policies
- **Focused face-to-face meetings on customers with accounts on many of our major systems (up to ten systems for a customer)**
 - Commitment to provide feedback on issues raised
 - Our projects turned into more than just policies (tools, processes, etc.)
 - Established working relationships with key HPCMP teams



Overview

- What Is “Baseline Configuration?”
- Why Is It Important?
- How Has It Evolved?
- **How Does It Work?**
- Where Do The Projects Come From?
- Accomplishments
 - Existing Policies
 - What the Team Thinks
- Summary



How Does It Work?

- **User Experience Interface – Stay grounded!**
 - Unvarnished criticism from experienced customers
 - Commitment to communicate progress and outcomes
- **Center Involvement – Assess feasibility**
 - All centers experience personal contact with customers
 - Quick turnaround from centers on potential solutions
- **Leadership Commitment – Part of being a Service Organization**
 - Making things easier for our customers
 - Understanding of the challenges that creates for center staffs
 - Implementing many of these ideas may look easy, but heavy lifting is often required
- **People – Good People Make Good Things Happen**



Baseline Configuration Team

Jeff Graham - Team Lead, Technical Director at AFRL DSRC

Aram Kevorkian - Co-team lead, HPC Chief Scientist at SSC Pacific

Vacant – Administrative Executive

Steve Finn – UAG member/rep – Principal Systems Analyst

Christine Cuicchi – HPCMPO Liaison, Deputy Associate Director for HPC Centers

Sunita Allwerdt – Alternate, HPCMPO

Matt Koop - PETTT Advanced Computational Environment (ACE)

Rhonda Vickery – Alternate, PETTT ACE

Tracey Smith – Consolidated Customer Assistance Center Lead, AFRL DSRC

Darryl Dawson - HPCMPO Security

Jackie Steele – Alternate, HPCMPO Security



Baseline Configuration Team

Center Representatives

ARL – George Petit, Software Support and Team Lead

Steve Thompson – Alternate

AFRL – Robert Jones, Operations Lead

Don Cable – Alternate

ERDC – David Sanders, HPC Service Center Specialist

TBD – Alternate

NAVY – Robert Lunceford, System Administration Lead

Lee Whatley/Mike Greer – Alternates

MHPCC – Steve Gima, System Administration Staff

Lance Terada – Alternate



Overview

- What Is “Baseline Configuration?”
- Why Is It Important?
- How Has It Evolved?
- How Does It Work?
- **Where Do The Projects Come From?**
- **Accomplishments**
 - Existing Policies
 - What the Team Thinks
- **Summary**



Face-to-Face Meeting Outcome

Baltimore - December 2011

Invited Speakers/Talks

Dr. Mark Cowan: Recent Benchmark Team Activities

Dr. Larry Davis: HPCMP Security

Dr. Doug Post: OSD Review of HPCMP by IDA-Led Panel

Dr. Danny Weddle: HPCMP User Satisfaction Survey

Dr. Dave Cronk: Workload Management



Invited Customers/Affiliation (Allocation, accounts/systems/centers)

Dr. Edward Byrd & Dr. Bill Mattson: Energetic Materials Science, ARL (11M, 6/3/3)

Dr. Robert Doney: U.S. Army RDECOM (69M, 7/7/3)

Bradford Green: NAVAIR (9M, 17/8/4)

B. Christopher Rinderspacher: U.S. Army RDECOM (5M, 5/5/4)



Face-to-Face Meeting Outcome

Baltimore - December 2011

Dr. Edward Byrd & Dr. William Mattson, Energetic Materials Science, ARL

- Consolidated Customer Assistance Center (CCAC) helpdesk and maintenance issues
- General DSRC issues
- Code and library issues

B. Christopher Rinderspacher, U. S. Army RDECOM

- Batch and queuing issues
- Code development issues

Brad Green, NAVAIR

- Inconsistencies, issues and recommendations

Dr. Robert Doney, U.S. Army RDECOM

- Documentation is scattered
- Major Differences in DSRC setups
- Negative user experiences



Face-to-Face Meeting Outcome

Dayton - May 2012

Chris Martin : Air Force Institute of Technology

- Data disappearing too quickly – can users extend life?

Dr Rajiv Berry: AFRL Materials Directorate

- Why so frequent shutdowns?

Dr. Oleg Kontsevoi: DTRA, ONRDC


- Documentation is scattered, incomplete, difficult to navigate
- Default user quotas are so low, they are a joke

Dr. Chris Woodward: AFRL Materials Directorate

- \$HOME data inaccessible when system is down
- Lots of “bad behavior” – queue stuffing in particular

Dr. Doug Davis: AFRL Propulsion Directorate

- Moving data from workspace to archive through the CWFS is causing problems



Overview

- What Is “Baseline Configuration?”
- Why Is It Important?
- How Has It Evolved?
- How Does It Work?
- Where Do The Projects Come From?
- **Accomplishments**
 - Existing Policies
 - What the Team Thinks
- Summary



37 Policies Completed



Compliance Matrix (12 September 2012)

| http://www.ccac.hpc.mil/consolidated/bc/policy.php | | Participating Shared Resource Centers | | | | | |
|---|--|---------------------------------------|-----|------|-------|------|-----|
| Project # | Policy Topics | AFRL | ARL | ERDC | MHPCC | NAVY | ORS |
| FY05-01 | Multiple-Version Software | G | G | G | G | G | G |
| FY05-02 | Minimum Scratch Space Retention Time | G | G | G | G | G | G |
| FY05-03 | Kerberos Ticket Life | G | G | G | G | G | G |
| FY05-04 | Environment Variables | G | R | R | G | G | G |
| FY05-05 | Queue Names | R | G | G | G | G | G |
| FY05-06 | Login Shells | G | G | G | G | G | G |
| FY06-01 | Open Source Math Libraries | R | G | G | G | G | G |
| FY06-02 | Kerberos/SSH Access | G | G | G | G | G | G |
| FY06-04 | Editors and Scripting Tools | R | R | G | G | G | G |
| FY06-05 | Debuggers | G | G | G | G | G | G |
| FY06-06 | Allocation Expiration Notice | G | G | G | G | G | G |
| FY06-09 | Archival Retrieval | G | R | G | G | R | G |
| FY06-10 | Standardized User Notices | G | G | G | G | G | G |
| FY06-11 | Announcing and Logging Changes | G | G | G | G | G | G |
| FY06-12 | Common Change Control Rollback | G | G | G | G | G | G |
| FY06-14 | Cross-Center File Transfers | R | G | G | G | G | G |
| FY06-15 | Sample Code Repository | G | G | G | G | G | G |
| FY06-16 | MPI Test Suite | G | G | G | G | G | G |

New column coming for Utility Servers

Compliance Matrix (12 September 2012)

| Project # | Policy Topics | Participating Shared Resource Centers | | | | | |
|-----------|---|---------------------------------------|-----|------|-------|------|-----|
| | | AFRL | ARL | ERDC | MHPCC | NAVY | ORS |
| FY06-17 | Multiple-Version Software Access via Modules | G | G | G | G | G | G |
| FY06-19 | Open Source Utilities | G | G | G | G | G | G |
| FY07-02 | Open Source Performance and Profiling Tools | G | G | G | G | G | G |
| FY07-03 | Common Set of Open Source Compilers | G | G | G | G | G | G |
| FY08-01 | ezVIZ Visualization Tool | R | G | G | G | G | G |
| FY09-01 | Advance Reservation | G | G | G | G | G | G |
| FY10-01 | New/Returning User Welcome Letter | G | G | G | G | G | G |
| FY10-02 | Common Open Source High Productivity Languages | G | G | G | G | G | G |
| FY10-03 | Consistency in Number of Characters in a Project Identifier | G | G | G | G | G | G |
| FY10-04 | Sustained System Performance (SSP) Test Job Execution | G | G | G | G | G | G |
| FY10-05 | Debug Queue Limits | G | G | G | G | G | G |
| FY10-06 | Common DSRC Commands and Tools | G | G | G | G | G | G |
| FY10-07 | Common Location to Maintain Codes | R | G | G | G | G | R |
| FY11-01 | Interactive Job Submittal Scripts | G | R | R | G | R | G |
| FY11-02 | HPC Round Robin Implementation | R | G | G | G | G | G |
| FY11-03 | Common User Group Naming | R | G | R | R | G | R |
| FY11-04 | Common Modules Environment | R | G | R | G | G | R |
| FY11-05 | Common Compiling Environments | R | G | R | G | G | R |
| FY12-01 | Minimum Home Directory Size and Backup Schedule | R | R | R | G | R | R |

Accomplishments - What The Team Thinks

- Round Robin Logins
- Consistency in number of characters in Project Identifier
- Sample Home Repository (\$SAMPLE_HOME)
- Mini Sustained System Performance Benchmarks
- Consistency in Announcing Changes and Outages
- Enterprise Environment Variables
- Common DSRC Commands and Tools
 - Check_license, node_use , qpeek, qview, show_queues, show_storage, show_usage
- Multiple Version Software Access via Modules
- Consolidated Hardware Information website
- But the most prevalent response as our greatest accomplishment was not a specific policy and takes us to the next slide



Accomplishments - What The Team Thinks

- Building a “program-wide culture of considering the centers as part of an enterprise comprising the entire program and not stand-alone centers. We have flexibility for accommodating local conditions at a center such as a particular user base, unique systems, or center focus but still retain core consistency
- Establishing “unique channels of communication with the DoD user community” – “Improving communications in both directions”
- “I think it is a big accomplishment to get buy-in from each center”
- “Rather than have 5 centers develop the same solutions independently, we have often had a solution developed at one center and then shared with the rest of the centers
- “If it weren’t for the BCT, people at the centers would still be clinging to the we’re different and our way is better concept. I’m glad that is gone now.”



Comprehensive Audit of All Policies

- Assigned a POC for each of the 34 existing policies, and underwent a comprehensive audit – looking at relevance, importance, and modernization
- The process was completed in 2 months (March & April 2012)

Common Change Control Rollback

BC Project: FY06-12

Date of Policy: 17 Apr 2006

Last Updated: 23 Mar 2012 ([see Revision Log](#))

This policy defines a Change Control Rollback policy to be followed consistently across participating centers. This policy applies to scheduled system changes that impact the user experience. These scheduled system changes will include a recovery plan, outlined in advance of the change, in case the change is determined to be unsuccessful and rollback to a previous configuration is necessary.

1. Each Center will have a Contingency/Recovery Plan for rollback to a previous configuration prior to making hardware or software changes.
2. Each center will utilize any applicable Test and Development System or other non-production system to test changes prior to implementation. To the extent feasible, appropriate test methods that represent daily system workload should be used to enhance risk mitigation of adverse consequences of changes.
3. Customer assistance staff will have access to a description of the scheduled system changes to facilitate identification of problems encountered by users following hardware/software changes.

BC Common Capabilities

Location

CCAC Web site under the Resources section

<http://www.ccac.hpc.mil/consolidated/bc/commonCapabilities.php>

Complete List of Current BC Common Capabilities

Open Source Math Libraries

Open Source Performance and Profiling Tools

Open Source High Productivity Languages

Source Code Repositories

MPI Test Suite Results (New)

Consolidated Hardware Information (New)

Consolidated Programming Environment (New)

Consolidated Hardware Information

<http://www.ccac.hpc.mil/commonHardware.php>

Instructions

Please check what fields you would like to display. Site and System Name will automatically be displayed.

Select [All](#) | [Consolidated Hardware](#) | [Compilers](#) | [MPI Implementation](#) | [Math-Scientific Libraries](#) | [None](#)

System Information

Vendor Model Decommission Date Average Capability (# of HABUs) Login Node Path Login Syntax Login File Total # of Cores System Speed

Login Node

Operating System Core Type Core Speed # of Cores Memory Type
 Memory Size

Compute Node

Operating System Core Type Core Speed # of Cores Memory Type
 Memory Size

Compilers

Intel PGI PathScale GNU SGI

MPI Implementation

IntelMPI OpenMPI SGIMPI xt-mpt
 mpt

Math-Scientific Libraries

LAPACK ScaLAPACK FFTW

[Submit Request](#)

[Clear Form](#)

Consolidated Programming Environment

<http://www.ccac.hpc.mil/commonProgram.php>

Instructions

Please check what fields you would like to display. If no "Systems" are selected, then results for ALL "Systems" will be returned.

Systems

Chugach DaVinci Diamond Einstein Garnet Harold Hawk Jade Mana Raptor

Compiler Vendor

Intel PGI PathScale GNU SGI

MPI Implementation

IntelMPI OpenMPI SGIMPI xt-mpt
 mpt

Compiler Type

C C++ FOR77 FOR90 FOR95
 FORTRAN Fortran 2003

[Submit Request](#)

[Clear Form](#)

Projects added in FY11/12

- **Eight New BC Policies**

- **FY10-06: Common DSRC Commands and Tools (Mar 2011)**
- **FY10-07: Common Location to Maintain Codes (Sep 2011)**
- **FY11-01: Interactive Job Submittal Scripts (Sep 2011)**
- **FY11-02: HPC Round Robin Implementation (Aug 2011)**
- **FY11-03: Common User Group Naming (Oct 2011)**
- **FY11-04: Common Modules Environment (Mar 2012)**
- **FY11-05: Common Compiling Environment (Mar 2012)**
- **FY12-01: Minimum Home Directory Size & Backup Schedule (Aug 2012)**



HPC Round Robin Implementation

Date of Policy: 22 Aug 2011

Last Updated: 02 Apr 2012

This policy provides a simple method for accessing login nodes on HPC systems, and thus, contributes to an improved, and more consistent customer experience for HPCMP users. The adopted method provides the functionality using round robin cycling through all available login nodes. All new systems, post Technology Insertion - 2010 (TI-10), will incorporate this functionality. It is recommended that the method be implemented as an alias on existing systems in order to create a baseline for user access.

The naming convention for login nodes will be:

systemnameNN

where NN ranges from 01 to 99.

By this policy, users will be able to login by specifying the system name instead of systemname-l01, or other.

Overview

- **What Is “Baseline Configuration?”**
- **Why Is It Important?**
- **How Has It Evolved?**
- **How Does It Work?**
- **Where Do The Projects Come From?**
- **Accomplishments**
 - Existing Policies
 - What the Team Thinks
- **Summary**



Summary

- **Allowing our customers to leverage the strength of the enterprise – expanding the culture**
- **Team energized by expanding connectivity**
- **Maintaining focus on user productivity**
- **Steady incremental progress – harder than it looks**
- **Team works well with other groups**
 - Focus is on leveraging strengths of other groups/process
 - Avoid duplication of effort through effective communication and working relationships



QUESTIONS?

