Flow of Ideas

- New Paradigm for C4I & Training Systems
  - Dr. Roger Smith, US Army PEO-STRI
  - Dr. Dave Pratt, SAIC
- Projects Underway
  - COL Craig Langhauser, US Army RDECOM STTC
  - Sergio Tafur, UCF IST

Roger Smith
US Army, PEO-STRI

High Performance Computer

- HPC = Supercomputer
- A computer which, among existing general-purpose computers at any given time, is **superlative**, often in several senses:
  - highest computation rate,
  - largest memory, or
  - highest cost.
- Predominantly, the term refers to the **fastest number crunchers**, that is, machines designed to perform numerical calculations at the highest speed that the latest electronic device technology and the state of the art of computer architecture allow.
- Processors: Teraflops (10^{12})
- Storage: Petabytes (10^{15})
- Network: Gigabits/sec (10^{9})

	Single	Double	Quad
1X	2 Gbit/s	4 Gbit/s	8 Gbit/s
4X	8 Gbit/s	16 Gbit/s	32 Gbit/s
12X	24 Gbit/s	48 Gbit/s	96 Gbit/s

Top 10 Super Computers

2007 List

Rank	Rmax Rpeak (Tflops)	Name	Computer Processor cores	Maker	Site Country, Year
1	1026 1375.8	Roadrunner	IBM BladeCenter QS22/LS21 122400 (Cel	IBM	Los Alamos National Laboratory
2	478.2 596.4	Blue Gene/L	eServer Bl 212992 (Pov		
3	450.3 557.1	Intrepid[1]	Blue Gene 163840 (Pov		
4	326 503.8	Ranger	Sun Conste 62976 (Opte		
5	205 260.2	Jaguar	Cray XT4 30976 (Opte		
6	167.3 222.8	JUGENE	Blue Gene 65536 (Pow		
7	126.9 172.0	Encanto	SGI Altix IC 14336 (Xeor		
8	117.9 170.9	EKA	Cluster Pla 14240 (Xeor		
9	112.5 139.3		Blue Gene 40960 (Pow		
10	106.1 122.9		SGI Altix IC 10240 (Xeor		

2008 List

Rank	Rmax Rpeak (Tflops)	Name	Computer Processor cores	Maker	Site Country, Year
1	1105 1456.7	Roadrunner	BladeCenter QS22/LS21 129600 (Cell/Opteron)	IBM	Los Alamos National Laboratory United States, 2008
2	1059 1381.4	Jaguar	Cray XT5 150152 (Opteron)	Cray	Oak Ridge National Laboratory United States, 2008
3	487.01 608.83	Pleiades	SGI Altix ICE 8200EX 51200 (Xeon), InfiniBand	SGI	NASA/Ames Research Center United States, 2008
4	478.2 596.4	Blue Gene/L	eServer Blue Gene Solution 212992 (Power)	IBM	Lawrence Livermore National Laboratory United States, 2007
5	450.3 557.1	Intrepid[1]	Blue Gene/P Solution 163840 (Power)	IBM	Argonne National Laboratory United States, 2007
6	433.2 579.38	Ranger	Sun Constellation System 62976 (Opteron), Infiniband	Sun	Texas Advanced Computing Center United States, 2008
7	266.3 355.51	Franklin	Cray XT4 38642 (Opteron)	Cray	NERSC/LBNL United States, 2008
8	205.0 260.2	Jaguar	Cray XT4 30976 (Opteron)	Cray	Oak Ridge National Laboratory United States, 2008
9	204.2 284.0	Red Storm	Cray XT3 38208 (Opteron)	Cray	NNSA/Sandia National Labatories United States, 2008
10	180.6 233.47	Dawning 5000A	Dawning 30720 (Opteron)	Dawning	Shanghai Supercomputer Center People's Republic of China, 2008

Who has the Power?

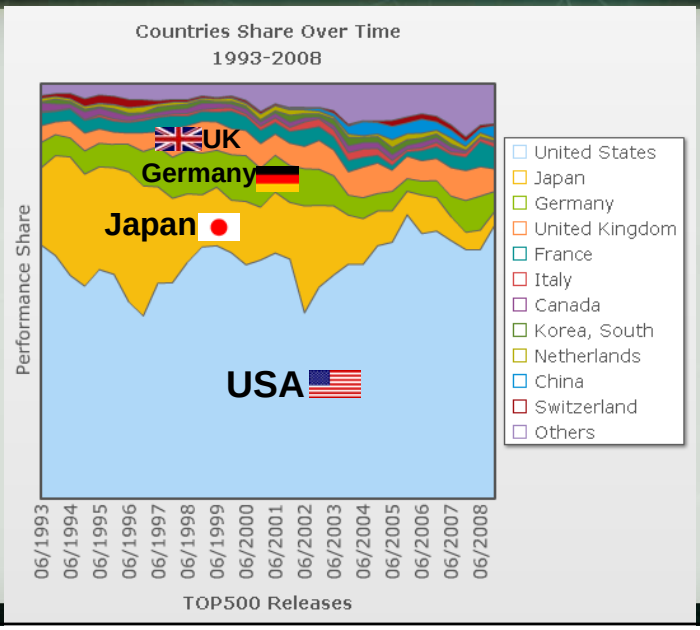

HPC as the Backend of Training

- High Compute Power in Professionally Managed Centers
  - Scalable to large exercises and large numbers of exercises
  - On-demand access to resources
- Power to Model
  - Finally put the “Reality” in “Virtual Reality”
  - Tighter system connections reduces lag
- Server-side Computing for LVC
  - Provide modeling for all types of exercises and experiments
- Heterogeneous System-of-HPC-Systems
  - Multi-site collaboration

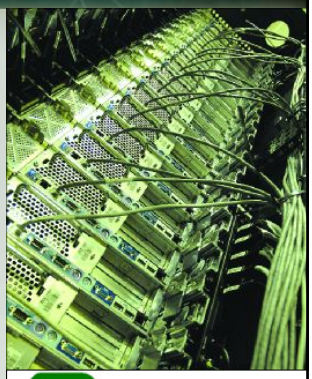

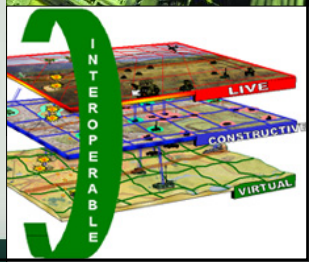

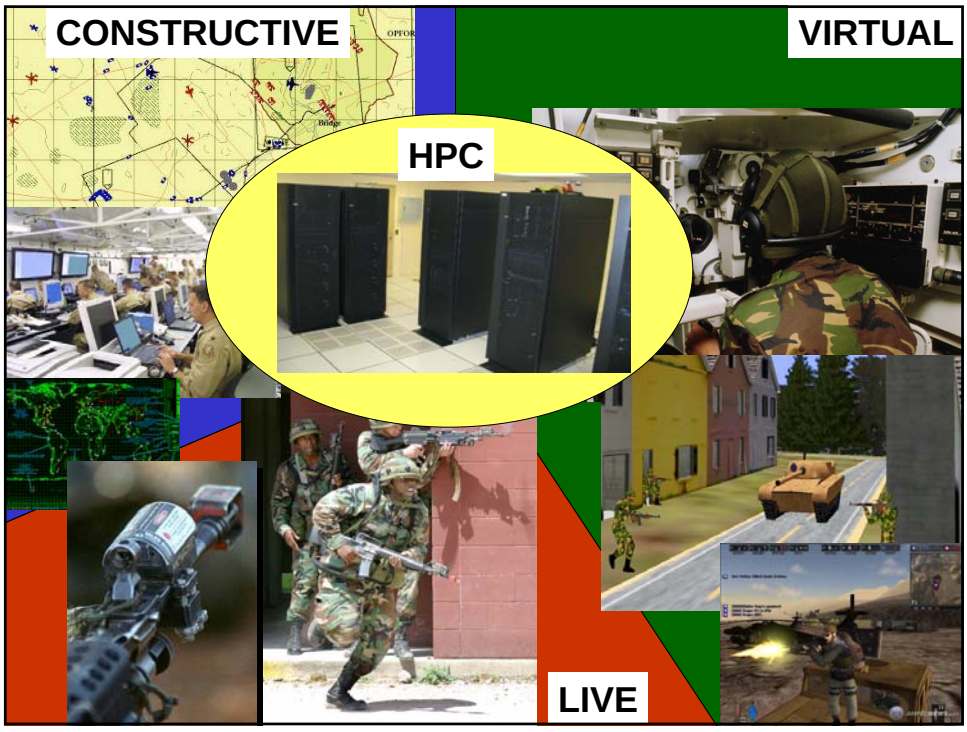

Simulation as an IT Service

- Scalable Simulation Services provided to remote customers on the customers' schedule
- Break the 1-to-1 relationship between equipment and events
- Light simulation client as an application on any military system
  - Browser-based
  - Generic Sim Engine & Tools
  - Flexible Game Engine
- Evolving Services at the core
  - Computation on Demand

[Commercial SOA Diagram]

Success = Customer Access

Dave Pratt
SAIC

Simulation Fidelity

- Model the large-scale contemporary operating environment (COE) to provide real-time, operationally relevant, Course Of Action Analysis (COAA)
  - Dynamic environment modeling
  - Individual, organizational, cultural modeling
  - Infrastructure/effects-based model
  - Weapon, communication, sensor system modeling
  - Learning/evolutionary systems
  - Data management and mining

All acronyms are spelled out on the last slide

The Big Idea

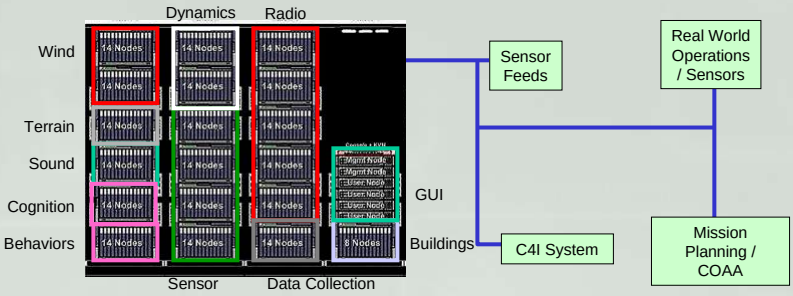

- Integrate high-fidelity emulation of the physical and electromagnetic environment with behavioral, organizational, and cultural simulation and real-world sensors and events to provide real-time seeding for mission planning/rehearsal

All acronyms are spelled out on the last slide

Sample Functionality Deltas

Functionality	Current	Desired
➢ Seismic and Chem/Biio Sensors Are Positioned Throughout a Neighborhood ❖ Command post monitors	➢ No Seismic Sensor/Emitters	➢ Real-Time Sensing
➢ Seismic Sensors Detect Truck Entering Area ❖ Patrol is sent out ❖ Truck is monitored	➢ No Seismic Sensor / Emitters	➢ Moving Sources
➢ Passengers Remount Truck ❖ Leave behind “device” and sniper	➢ No Chem/Bio Play	➢ Chem/Bio Sensors
➢ Passengers Dismount from Truck ❖ Picked up by the sensors ❖ Patrol is redirected to dismount area	➢ No Seismic Sensor/Emitters	➢ Multiple Types/Multiple Moving Sources
➢ “Device” Release ❖ Sensors detect plume ❖ Patrol is redirected up wind ❖ Locals “feel the effects”	➢ Uniform Static Wind Flow ➢ Limited Dispersion Model	➢ Chem/Bio Sensors ➢ Dynamic Urban Wind Flow ➢ Physically Correct Dispersion
➢ Patrol Contains Area ❖ Use smoke to signal aircraft ❖ Engage sniper	➢ No Acoustical Play ➢ No Obscurants	➢ Acoustic Sensor/Emitters ➢ Obscurant -based Transport And Line of Sight

Current System Limitations

- Computer architecture
  - CPUs, memory, networks
- Network-based architectures
  - Very high latency between nodes
- Software architecture and capabilities
  - Environmental modeling
  - Physiological modeling
  - C4ISR modeling
- Data management and analysis
  - Distributed data generation
  - Semi-centralized analysis and interpretations
- Entity-based decomposition
  - Consistency management is a major consumer of resources
  - Uniprocessor -based codes
    - Can only model things that fit on one processor

All acronyms are spelled out on the last slide

Current Systems Can't Cut It

- Single processor lack of computational power
  - Simplification/elimination of physical reality
- Network of workstations (NOW) limitations
  - Communication delays and overhead
  - System management costs/difficulties
  - World consistency overhead
  - Scalability limitations
- Clusters move bottleneck
  - Still require extensive consistency management
  - High latency between processors
- Co-processors are not flexible enough for operational scenarios
  - Shared bus architectures limit memory bandwidth
  - Multiple concurrent codes require reloading pipeline
  - Local computing results need to be propagated

All acronyms are spelled out on the last slide

Previous Simulation Experiments

- JWARS parallelization
  - HPCMO, SAIC®, CACI®
- WARSIM Port
  - HPCMO, SAIC, ARL
- Millennium Challenge Exercise Clutter, using JointSAF
  - JFCOM, MHPCC/AFRLMSRC, USC, Alion®, LM
- Physics-based Environment for Urban Operations, using OneSAF®
  - HPCMO, STRI, SAIC, NRL, CERDEC / OpCoast®
- PEO-C3T C4ISR On-the-Move (OTM) program includes OneSAF
  - CERDEC, HPTi®, SAIC, HPCMO

All acronyms are spelled out on the last slide

OneSAF Porting

- HPCMO funded
- Team Orlando coordination
  - STRI, STTC, HPCMO, SAIC, UCF IST
- Porting OneSAF® to UCF Stokes HPC
  - IBM xSeries® machine running Red Hat® Enterprise

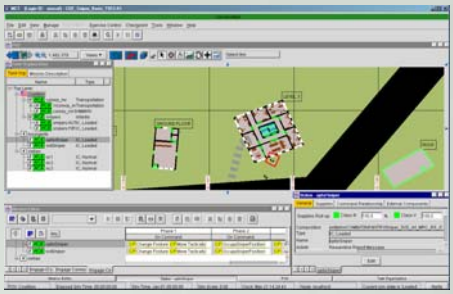

All acronyms are spelled out on the last slide

OneSAF HPC Research Problems

- Porting
  - Host OneSAF® Sim Core and MCT* on HPC
- Computational distribution
  - Efficiency of thread distribution in HPC environment
  - Function of Java™ virtual machine (JVM), node/process/core availability
- MCT interface
  - Internal to HPC with VNC video exported
  - External with efficient network communications
- Light interface
  - Operate via light GUI outside of HPC (e.g. browser interface)
- Infiniband™ network
  - Multiple instances using Infiniband versus Ethernet to communicate

All acronyms are spelled out on the last slide

Technical Challenges

1. **Interactive HPC** – exploring bandwidth sufficiency from the computational elements to multiple external users
2. **HPC I/O structure** – HPC structure that best supports interactive users
3. **Simulation as an IT service** – using HPC as the server-side of a ubiquitous software service
4. **Fault tolerance** – auto-restarting a job when a processor dies and doing so without losing the partial data that was in the works
5. **Processing hierarchy** - introduction of a processing hierarchy in the logic of simulation architecture design
6. **Organizational acceptance** - technical and organizational challenges of using a shared resource for interactive simulation, rather than distributed commodity hardware
7. **Parallel programming** – training the simulation industry in parallel programming techniques versus the network programming that has dominated for 20 years
8. **Cloud compute environments** – load-balancing and task assignment in a network of HPCs and traditional workstations
9. **Interactive user security** – verification of users communicating with jobs on open ports

All acronyms are spelled out on the last slide

U.S. Army Research, Development and Engineering Command

RDECOM

TECHNOLOGY DRIVEN. WARFIGHTER FOCUSED.

COL Craig G. Langhauser
Director STTC

TECHNOLOGY DRIVEN. WARFIGHTER FOCUSED.

STTC HPC Vision

- Interactive High Performance Computing
  - Leverage the Power of HPC as a Server Farm for Interactive Simulations
  - Enable Multiple Simultaneous Exercises Supported from a Single Simulation Center
  - Create Physics-based Objects, Weather, Terrain Modeling, and Multitudes of Dynamic Entities
  - Facilitate Tight Network Connections Between Applications to Eliminate Lag

TECHNOLOGY DRIVEN. WARFIGHTER FOCUSED.

STTC Partnerships

- PEO-STRI
  - Negotiated Transfer of Funds from HPCMO
  - Provided Background to Allow Funding for a New HPC
- STTC
  - Obtained and Awarded a Congressional plus-up to UCF
  - Managed Effort
- UCF
  - Bought and Built Stokes
- Productivity Enhancement and Technology Transfer (PET) Program
  - Funded by the High Performance Computing Modernization Office (HPCMO) to Help get Code Running on HPC assets
- PM OneSAF
  - Target Transition Agency

OneSAF on a HPC

- Accomplishments
  - Compiled and Ran OneSAF on the Cluster
  - OneSAF {Built on a Build Node (Intel)} was Ported to the Virtualization Node (AMD) via Symlinks to the Same Libraries
- Challenges
  - Porting issues with Compilers and Libraries
  - Requires Java 1.5 Rel 14
- Current Insights
  - Reducing Overall Enterprise Lifecycle Costs by using Centralized Computational Platforms is Feasible While Increasing
    - Soldier/Unit Access to Training
    - Exercise Reliability and Availability
    - Model Fidelity and Synchronization
  - It is Still not Write Once, Run Anywhere

Potential HPC Future Efforts

- Benchmark OneSAF v2.0 on Different HPC to Gauge Machine Performance
- Thin Client/Location Transparent Simulations
- Using High Fidelity Models on a Centralized HPC to Create a Functionally Distributed Training Model

TECHNOLOGY DRIVEN. WARFIGHTER FOCUSED.

HPC Utilization

- Looking for Partners
  - DARPA RAID Program
  - RDECOM MATREX Program Use of HPC OneSAF Version
  - Synthetic Natural Environments on HPC

TECHNOLOGY DRIVEN. WARFIGHTER FOCUSED.

HPC Challenges

- Software, Software, Software
  - No Commercial Killer App
  - Hardware too Cheap
  - No Competitive Advantage
- DoD or Academia Will Have to Lead
  - Artificial Intelligence
  - Dynamic Real Time Modeling & Simulation

HPC Prediction

- Assertions:
  - Multi-core processors are the Standard NOW
  - DoD Software Will be Reengineered
  - HPC Costs are Dropping all the Time
- Prediction:
  - HPC will be the Simulation Technology of Choice within the next 2 Tech Refresh Cycles.

If You are not Planning for HPC now…
You are BEHIND the Power Curve

University of Central Florida

Sergio Tafur
UCF Institute for Simulation & Training

University of Central Florida

High Performance Computing at the University of Central Florida

30

Some perspective: Computing Power and Capabilities

The Hans Moravec vision: **What $1000 will buy**

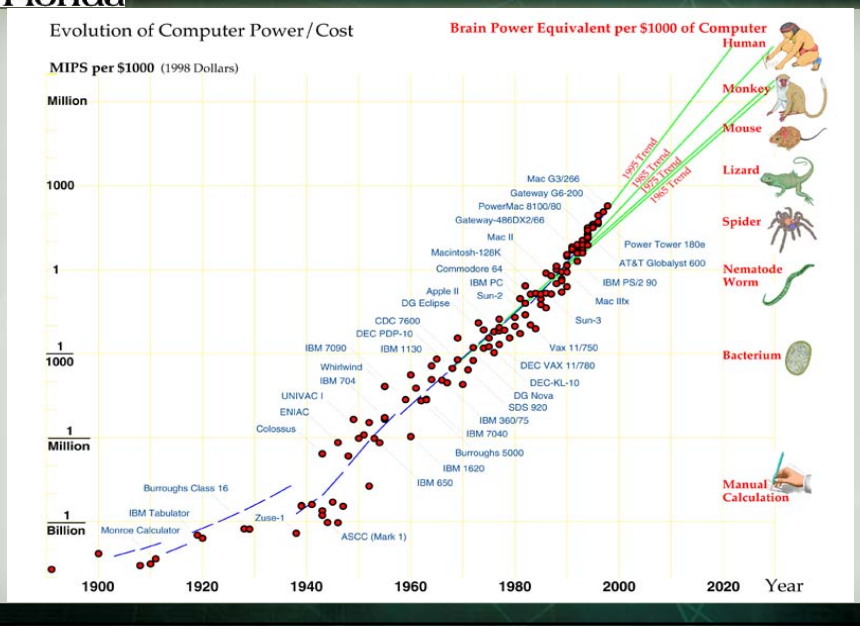

31

Computing Power and Capabilities

The Hans Moravec vision: **What you can do at any cost**

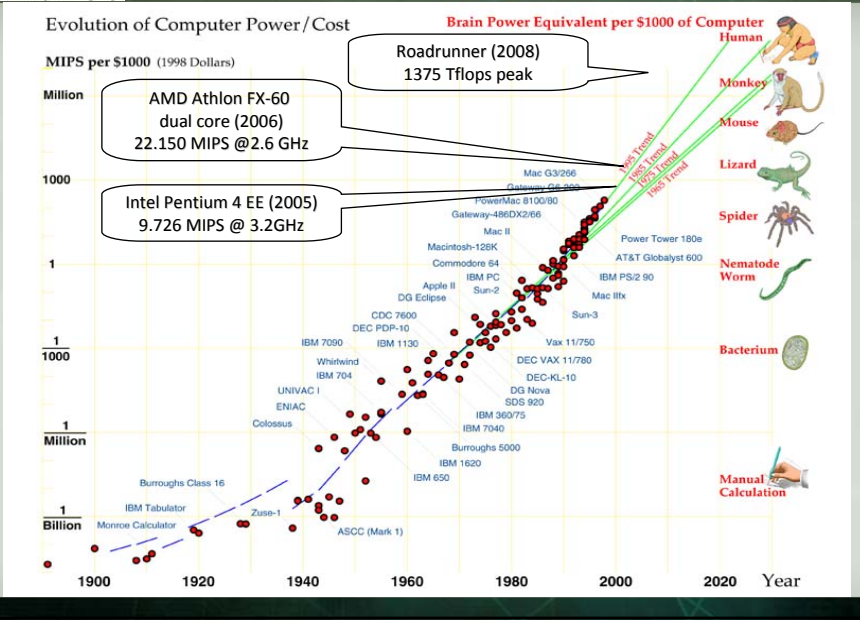

32

HPC Today - Desktop Tomorrow

- Multi-core is already in your computer, you now have a supercomputer of 10 years ago on your desktop.
- In ten years you will have today's supercomputer capability on your desktop.
- Pretty much nothing that runs on your current system can take advantage of this power.
- You need to be thinking about this issue in every procurement.

The next big thing: Interactive Simulation on HPC

- Needs
  - Real time capability using fast processors and high-speed interconnects
  - High fidelity
  - Low latency/High bandwidth interconnects
  - Real time I/O
  - Connection to real world assets
  - Fixed frame rates (some apps)
- Interfaces with sensory processors (e.g., interactive visualization, haptics, …)
- Scalability in terms of HPC architecture and simulation entities

Things that need to be investigated

- Multi-core Programming for M&S Applications
  - Tight Timing Constraints
  - Low Latency
  - I/O Bound
- Use of Cell Processor for M&S
- Multi-World Systems
- LVC Implementations/Experimentation
- Terrain Correlation
- Granular Propagation Mitigation Methods
- Multi-scale Simulations
- Benchmarks
- De-coupling SAF Models
- ????

Important Considerations

- Let's remember the 'human factor'
  - How will a user interact with an HPC?
  - How will multiple users interact with an HPC & maintain coherence of I/O?
  - How will interim results be gathered?
  - How can timely and relevant HF experiments be developed to influence the design?
- Get developers involved…

Human Centered M&S Research

- Apparent Parallelizable Systems (SAF/Games)
  - Approaches to Parallelization
  - Spatial & Temporal Coherency
  - Performance Assessment & Optimization
- Interactive & Visualization
  - Review Lit in Sci Vis & Comp Steering
  - Leverage Existing Software (e.g., OLIVE, DCV)
  - Consider & Baseline Different Approaches
- LVC Modeling

Central Florida HPC Assets
UCF/IST HPC System (Stokes)

- Processors, Xeon 3 GHz, 64b
  - ~6.6 Tflops (1000 billion floating point operations/sec.)
- 648 Cores
  - 4 Visualization Nodes
- 1.424 TeraBytes Memory (1000 billion numbers)
- 42+ TB Storage
- System Interconnects
  - IB 20 Gigabits/sec
  - Gigabit Ethernet

Team Orlando HPC "Stokes"

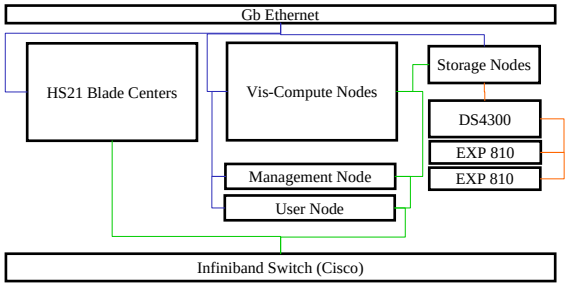

Management server

7: chassis
14: blades in each chassis
8: cores in each blade
784: total number of cores

BladeCenter® Cluster

4: servers
8: cores in each server
32: total number of cores

Visualization & Compute Servers

43: number of disk drives
750: disk capacity in GB
32,250: raw capacity in GB

Storage Subsystem

39

IBM HS21 BladeCenter® Cluster

	Installed in Each Blade
Intel Xeon Processor	2 quad-core E5450 (Harpertown) 8 cores @ 3.0 GHz
L2 Cache	2 X 2 X 6144 KiB
Memory	8 GB, 667 MHz, DDR2
Front Side Bus	1333 MT/s
internal disk	73 GB, 10K RPM SAS
Power	80 W
Ethernet	1 Gb Ethernet
InfiniBand	Single-port 4X DDR IB PCI-E HCA (Cisco)
Linux OS	Red Hat V5
Compilers	GCC Intel Fortran V10.1 Intel C++ V10.1 PGI V7

Orlando HPC: 85 Blades, 680 cores

40

Usage Groupings

- **Science Based M&S Usage**
  - Nano Technology
  - Civil Engineering
  - Physics
- Batch Processing
- Existing Programs (e.g., MatLab)
- New Data
- Large Runs
- Segue to Larger Systems

- **Human Centered M&S Usage**
  - IST
  - Army
  - Partnering Industry
- Interactive
  - Human in the Loop
  - Modeling Human Activity
- Multi-modal I/O
- Multi-user
- No Existing HPC Programs or Data

Current Users

- IST
- Physics
- Mathematics
- Chemistry
- Nanoscience
- Civil Engr
- Mech. Engr
- Industrial Engr
- Electrical & Computer Engr
- CREOL

- SAIC
- Forterra
- <insert your name here>

Potential for Getting Involved

- Relevance to UCF Interests
  - UCF M&S (Fully Supported)
  - Other UCF (Partially Supported)
  - Other Entities (Profit and Non-Profit)
    - With UCF M&S (Fully Supported)
    - With Other UCF (Partially Supported)
- Other Users (Lower Queue Priority)
  - University/Non-Profits (Case by Case)
  - For Profit Proprietary
    - Provide Funds for Staff
    - Constrained Use of Software
- Joint Proposals

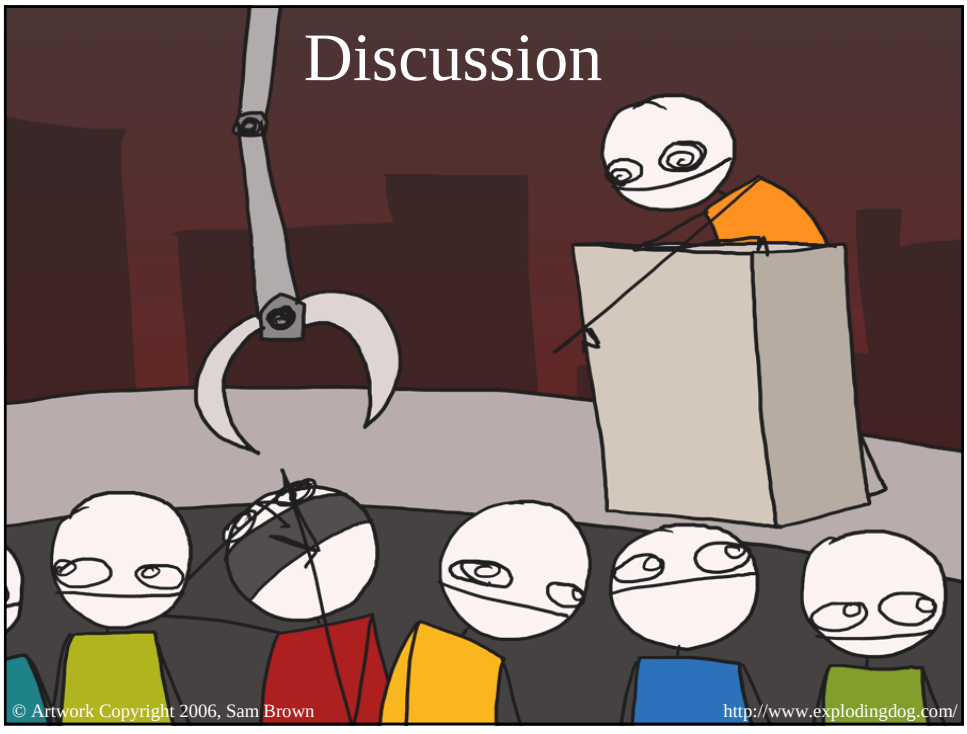

ACRONYMS

- ARL – Army Research Laboratory
- C4ISR – command, control, communications, computers, intelligence, surveillance and reconnaissance
- CERDEC – Communications-Electronics Research, Development, and Engineering Center
- GUI – graphical user interface
- HALLE – HPC Army Laboratory for LVC (live/virtual/constructive) Experimentation Project
- HPC – high-performance computer
- HPCMO – High Performance Computing Modernization Office
- JFCOM – Joint Forces Command
- JWARS – Joint Warfare System Now called the Joint Analysis System (JAS)
- LM – Lockheed Martin
- LVC – live/virtual/constructive
- MCT – Management and Control Tool
- MHPCC/AFRLMSRC– Maui High Performance Computing Center/Air Force Research Laboratory Major Shared Resource Center
- NRL – Naval Research Laboratory
- PEO STRI – Program Executive Office Simulation, Training and Instrumentation
- STTC – Simulation & Training Technology Center
- UCF – University of Central Florida
- UCF IST – University of Central Florida Institute for Simulation and Training
- USC – University of Southern California
- VNC – virtual network computing
- WARSIM – Warfighters' Simulation