



# CORAL Technical Requirements

CORAL Vendor Meeting

May 31, 2013



U.S. DEPARTMENT OF  
**ENERGY**

LLNL-PRES-637163

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC



**CORAL**

# CORAL Time Line (Chapter 1)

- RFP Release Fall 2013 (mid-October)
- Responses due seven business weeks later
  - Holidays and SC13 not included
- Award to two diverse architectures
  - Technical requirements are flexible in order to allow this
- 2 R&D Contracts
- 3 System Build Contracts (one from each lab) with Go/NoGo points
- Delivery in 2017 of all three systems
  - Key point the features requested on the following slides are for a system with completed delivery in 2017 time frame

# Basic Terminology Definitions

**The technical requirements have these designations:**

- **Mandatory Requirement (MR):**
  - Performance features that are essential to CORAL requirements
  - Required to be **responsive**
- **Mandatory Option (MO):**
  - Performance features that each laboratory may choose to exercise
  - Required to be **responsive**
- **Technical Option (TO):**
  - Performance features that each laboratory may choose to exercise
  - Graded (1, 2 or 3) based on overall importance to CORAL
  - **Not** required to be responsive
- **Target Requirement (TR):**
  - Performance features that are important to CORAL requirements
  - Graded (1, 2 or 3) based on overall importance to CORAL
    - 1 is essential; 2 is all but essential; 3 is desirable
  - **Not** required to be responsive

CORAL

# RFP will have three Mandatory Requirements

- **Description of the CORAL System**
  - Overall system architecture and details of the interconnect
  - Detailed node architecture diagram w/all data movement paths
- **Description of the Software Stack**
  - Describe all software components and license strategy for each
  - **Strong** preference for open source software
- **Description of Program management**
  - Proposed multi-year collaboration plan
  - Plan to deliver multiple systems at the same time
  - Perceived risks, risk management plan, and potential mitigations
  - **Ability to ensure the responsiveness of its partners to the performance requirements**

# Application Performance Requirements are the highest priority to CORAL

An average “figure of merit” (FOM) improvement of 4-8X for scalable science applications and 6-12X for throughput (UQ) applications over today’s DOE systems by 2017: Mission need **(CD-0)** requirement

- The Offeror will provide actual, predicted and/or extrapolated performance results for the proposed system for the following:

- **CORAL System Performance (TR-1)**

- Average FOM over four TR-1 scalable science apps  $\geq 4.0$
- Average FOM over four TR-1 throughput apps  $\geq 6.0$
- Raw results of five TR-1 skeleton apps

- **CORAL System Extended Performance (TR-2)**

- Proposal gets a plus for presenting results for TR-2 apps

- **CORAL System Microbenchmarks Performance (TR-3)**

- Present raw results for microbenchmarks suite

**Details on applications and performance tests in next talk**

# High Level System Requirements

**It is desirable for system to meet or exceed the following:**

- **CORAL System Peak (TR-1)**

- Peak of at least 100PF (double precision floating point) (per CD-0)

- **Total Memory (TR-1)**

- Must propose the same memory capacity and configuration as is used to obtain the reported benchmark results
- At least 1 GB per MPI task. Offer can count all directly addressable memory: NVRAM, fast memory, smart memory, etc.

- **Maximum Power Consumption (TR-1)**

- Can not exceed 20 MW for system and all peripherals

- **System Resilience (TR-1)**

- 144 hour Mean time between application failure due to system fault requiring user or administrator action.
- Restarting impacted services and jobs without human intervention is allowed.

# Mandatory Options in Overall Offering

## Separately described and priced options in the response

- Scale the System Size (MO)
- Scale the System Memory (MO)
- Scale the System Interconnect (MO)
- Scale the System I/O (MO)
- Parallel File System and SAN (MO)

Because the different CORAL sites may want to procure different system configurations.

- CORAL-Scalable Unit (MO)



All CORAL sites want to procure a small version of the “other” system for app portability.

- Options for mid-life Upgrades (TO-1)



System does not have to be upgradable, but if it is, what are the options and time-line

# Site-Specific Mandatory Options

**Additional separately described and priced options in the response that reflect differences across laboratories**

- 6.1.2 Off-Cluster Connectivity (MO)
  - 9.2.2.3.1 Baseline Debugger (MO)
  - 11.1.1 24/7 Hardware Maintenance Option (MO)
  - 11.1.2 12/7 Hardware Maintenance Option (MO)
  - 11.1.5.1 FRU with non-volatile memory destroyed (MO)
  - 11.1.5.2 FRU with non-volatile memory returned (MO)
  - 11.5 On-Site Analyst Support (MO)
  - 13.5.1 Alternative Integrated Cooling Solution (MO)
  - 13.5.2 Alternative Redundant 480VAC or 600VAC Solution (MO)
- Each lab will choose one.
- Each lab will choose one.

# Additional Requirements for Discussion

## Early Access to CORAL Hardware/Software (TR-1)

Throughout the long-term collaboration, Offeror will propose mechanisms to provide the Laboratories with early access to hardware and software technology for evaluation, feedback, and application testing

## 5% Runtime Variability (TR-1)

Reproducible performance from run to run is a highly desired property.

The metric for reproducible performance is based upon the variation in job execution for the TR-1 Scalable Science and TR-1 Throughput Benchmarks

# Comments on other Chapter Requirements?

## **Chapter 4 Application Benchmarks (next talk)**

**Chapter 5** Compute Partition

**Chapter 6** IO Subsystem and burst buffer

**Chapter 7** Interconnect

**Chapter 8** OS and Middleware

**Chapter 9** Front-end Environment and programming tools

**Chapter 10** System Management and RAS

**Chapter 11** Maintenance and Support

**Chapter 12** Parallel File System (**Mandatory Option**)

Contains two internal mandatory options (12.1.1.1, 12.2.3)

**Chapter 13** Facility Requirements

**Chapter 14** Project Management

Questions?