

Language for Power Monitoring

Revision date: 20160208

WHAT: This document enumerates requirements for power data collection on computer systems, reporting, and documentation.

WHY: Power data is collected and reported in order to be sent to other display and analysis systems outside of the local system referred to by the Statement of Work

Notes to Project Directors:

1. The language in this document is the starting point for a discussion with vendors.
2. In evaluating vendor responses and the ability to report power per rack and system power data, the compute racks typically consume more power than cooling and CDUs. Thus, collecting data from compute racks is more important than the cooling and CDU power. Scores for comparison of vendor solutions should reflect this bias. The scoring could be weighted by the relative power load expected per rack type.
3. Requirement 5 specifies a fixed-sized database for data storage. The size of the database, number of entries, specific format, and fields are intentionally unspecified at this time. If there are more specific requirements, they will be supplied to the Contractor.
4. Ganglia is referenced in this document. Which software is used may change in future requirements.
5. If any of the definitions below (rack, system, measurement, reporting) are defined elsewhere, these definitions shall take precedence with respect to power monitoring capability.

Terms used herein are defined as follows:

System: A group of interacting, interrelated, or independent components forming an integrated and complex whole. This includes all computer equipment for normal operations—compute, networking, management, storage, cooling—that the Contractor has procured or built, delivered, installed, and configured per the System’s contract or contracts. The system is composed of one or more racks.

Rack: Also referred to as a cabinet. A Rack is often the largest replicable unit in a System. A rack is typically composed of some combination of compute, cooling, network, storage, or management devices.

Measured value: The sampled or determined value of a time-varying parameter at a specific instant in time. Measurement occurs at least as often as reporting.

Reported value: The value of a parameter calculated from or related to one or more measured values of that parameter. Reported values characterizing historical behavior are based on one or more measured values.

Instantaneous Power Draw (IPD): Reported and measured values at a specific instant in time.

Power monitoring infrastructure: the rack and system measurement, collection, and reporting software and hardware.

Power Monitoring Infrastructure – WHAT to report

1. The Contractor shall implement a capability for measuring Instantaneous Power Draw (IPD) at the per-rack and system levels.
2. The Contractor shall implement a capability for reporting the following power values. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - a. Maximum and minimum power over the previous reporting interval, in watts.
 - b. Average power draw over the previous reporting interval, in watts. Average power is defined as the sum of measured values divided by the number of measured values
 - c. Energy consumption over the previous reporting interval, in kWh.

Power Monitoring Infrastructure – HOW to report

3. Reported power values
 - a. shall be reported in watts unless otherwise specified by the Government. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - b. shall have a resolution of at least one watt. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - c. shall be date- and time-stamped with a resolution of at least one second. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - d. shall have date and time-stamps with accuracy of at least one second. If time accuracy requirements are specified elsewhere in this requirements document, then those requirements shall take precedence. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - e. shall be date and time-stamped with a time zone. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - f. shall have time zone of GMT. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - g. shall be accurate to within +/- 2% of tested value. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
4. Reporting power data
 - a. interval shall be at least once per minute. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.
 - b. rate shall be no more than 100 times per second.

Power Monitoring Infrastructure – HOW to store

5. All power data shall be stored in a fixed-size database with age-off of data readable by Ganglia. The database shall retain reported power data (both system and per-rack) for at least three days. If the Contractor is unable to comply, the Contractor shall report to the Government what can be done.

Power Monitoring Infrastructure – HOW to collect

6. The Contractor shall report Rack and System power figures using one or more of the following methods:
 - a. Plain-text output in response to a command-line query on at least one of the System's management servers.
 - b. Plain-text output in response to a command-line query issued from at least one of the System's management servers to a networked, in-Rack PDU.
 - c. Plain-text output in response to a command-line query issued from at least one of the System's management servers to a networked, in-Rack Automatic Transfer Switch.
 - d. A Government-approved reporting mechanism. The Contractor shall propose to the Government an alternative implementation.

Constraints on Implementation

7. Sustained, maximum usage of the Power Monitoring Infrastructure shall not negatively impact any contractually obligated performance parameter below its threshold or its minimally acceptable value.
8. The Power Monitoring Infrastructure shall not negatively impact contractually obligated network or computer security requirements.

Documentation

9. Where applicable, the Contractor shall physically label all System power cables inside each rack and between racks—at both ends of each cable—indicating which device is plugged into which outlet.
10. The Contractor shall provide documentation describing architecture, functionality, usage, and support of the Power Monitoring Infrastructure and its software and firmware components.
11. The Contractor shall develop, document, and execute a test plan for the Power Monitoring Infrastructure.
12. The Contractor shall document and submit a Power Monitoring Infrastructure test report.
13. The Contractor shall provide education and training of Government personnel on Power Monitoring Infrastructure best practices and common use cases.

Power Monitoring Infrastructure Deliverables (suggested)

DID Number	Title	Description	Delivery
	Test Plan	Power Monitoring Infrastructure Test Plan (what to test) <ul style="list-style-type: none"> • Test command-line • Test Ganglia compatibility • Validate Power values 	Initial and Interim Submissions: As agreed upon at contract kickoff. Final: NLT 1 month before first Test Plan procedure.
	Test/Inspection Report	Power Monitoring Infrastructure Test Report – show results from the testing	Submissions: NLT 5 days after completion of each Test Plan procedure.
	Training Project Plan (TPP)	Power Monitoring Infrastructure Training Plan	Initial and Interim Submissions: As agreed upon at contract kickoff. Final: NLT 1 month before first Training Plan event.
	Technical Report	Power Monitoring Infrastructure Architecture Document	Initial and Interim Submissions: As agreed upon at contract kickoff. Final: NLT 1 month prior to Government Acceptance Test.
	Technical Report	Power Monitoring Infrastructure Usage and Support Document	Initial and Interim Submissions: As agreed upon at contract kickoff. Final: NLT 1 month prior to Government Acceptance Test.

To validate power monitoring, the government acceptance team will be testing the following:

- Plot power as a function of time with live data.
- Visualize data using an open source tool (i.e. Ganglia).
- Visually verify plots and ability to get live data.
- Plot shall include max, min, average values per reporting interval.
- Plot power data during idle state.
- Plot power data during benchmark runs.
- Verify that power during benchmark is greater than power during idle.
- Verify power data exists for both per-rack and full system.
- Verify power reporting interval
- Verify data is plain text via command line.