

CRAY

**Cray® XC™ Advanced Power
Management Updates
CUG 2018**

Steven J. Martin, Cray Inc.



Cray[®] XC[™] Advanced Power Management (APM) Updates

Steven Martin, Greg Koprowski, Dr. Sean Wallace
Cray Inc.

{stevem, gkoprowski, swallace}@cray.com

Cray[®] XC[™] Advanced Power Management Updates
Steve J. Martin, Greg J. Koprowski, Dr. Sean J. Wallace
Copyright 2014 Cray Inc. All rights reserved.



Please hold questions for the end of the presentation!

COMPUTE

STORE

ANALYZE



- **This will be a quick update on XC APM**
 - Monitoring and control of power and energy is important in HPC
 - The Cray APM team wants to promote open communication
- **This presentation will cover the following**
 - Compute blades supported on the Cray XC50 platform
 - Cray XC50 compute blade APM features
 - Cray XC50 blade featuring Intel® Xeon® Scalable processors
 - Cray Power Management Database (PMDb) updates
 - Cray XC50 APM usage examples

Compute Blades on the Cray XC50 Platform



- **Cray APM is supported on all XC compute blades!**
- **The Cray XC50 supports the following blade types**
 - NVIDIA® Tesla® P100 PCIe GPUs
 - Intel® Xeon® Scalable processors
 - Cavium ThunderX2™ processors



COMPUTE

| STORE

| ANALYZE

Cray XC50 Compute Blade APM Features




- **Power and energy monitoring**

- Advanced out-of-band (OOB) blade level monitoring
 - Total node, aggregate CPU and Memory, power and energy
- Publishing power and energy data OOB into PMDB
- Publishing power and energy data in-band /sys/cray/pm_counters

- **Power and energy control**

- Node level power capping
- P-state and c-state limiting
- Cray advanced platform monitoring and control (CAPMC)

Cray XC50 Blade with Intel Xeon Scalable Processors



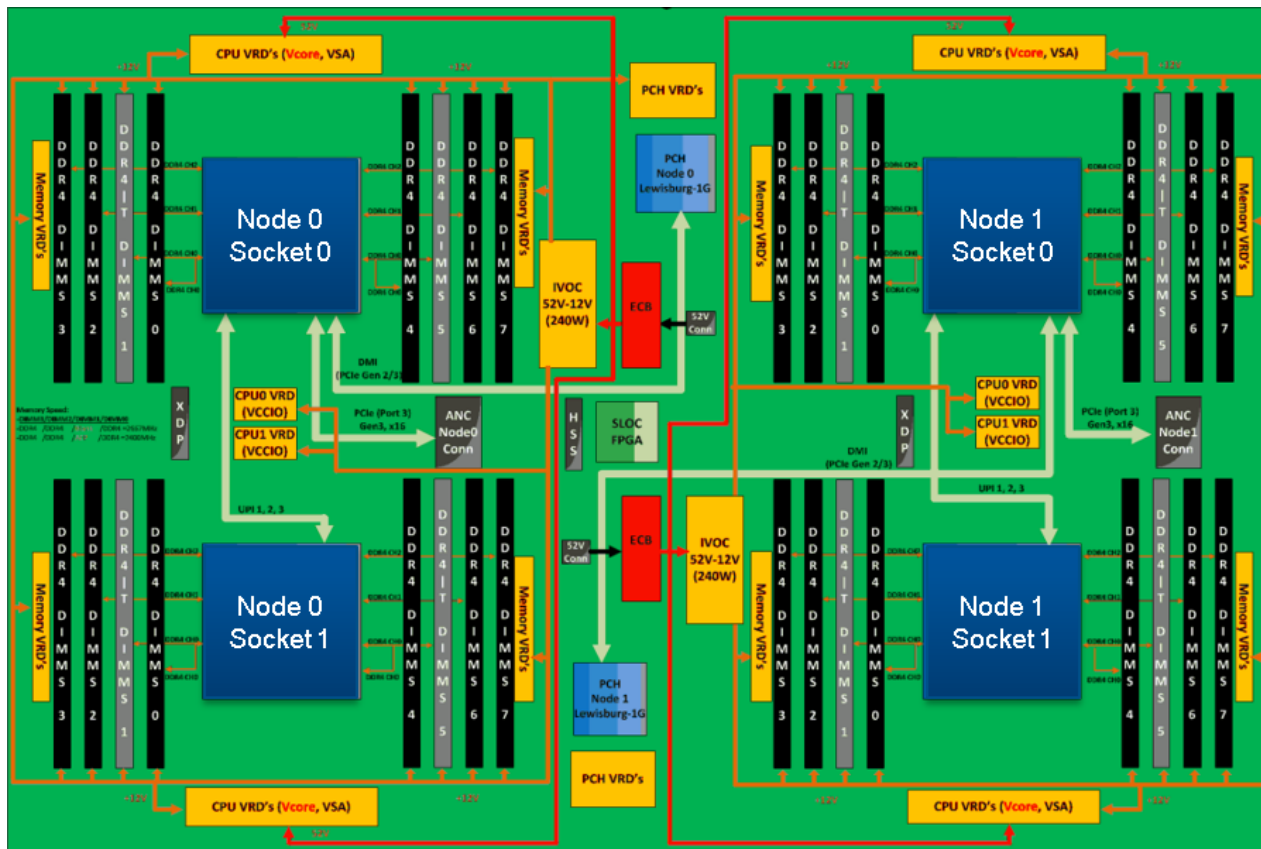
- Four 2-socket nodes per blade
- Eight DDR4 memory slots per socket
 - Six memory channels per socket
- Varying core counts and TDP up to 165 watts.

COMPUTE

STORE

ANALYZE

Cray SPDC with Intel Xeon Scalable Processors



COMPUTE

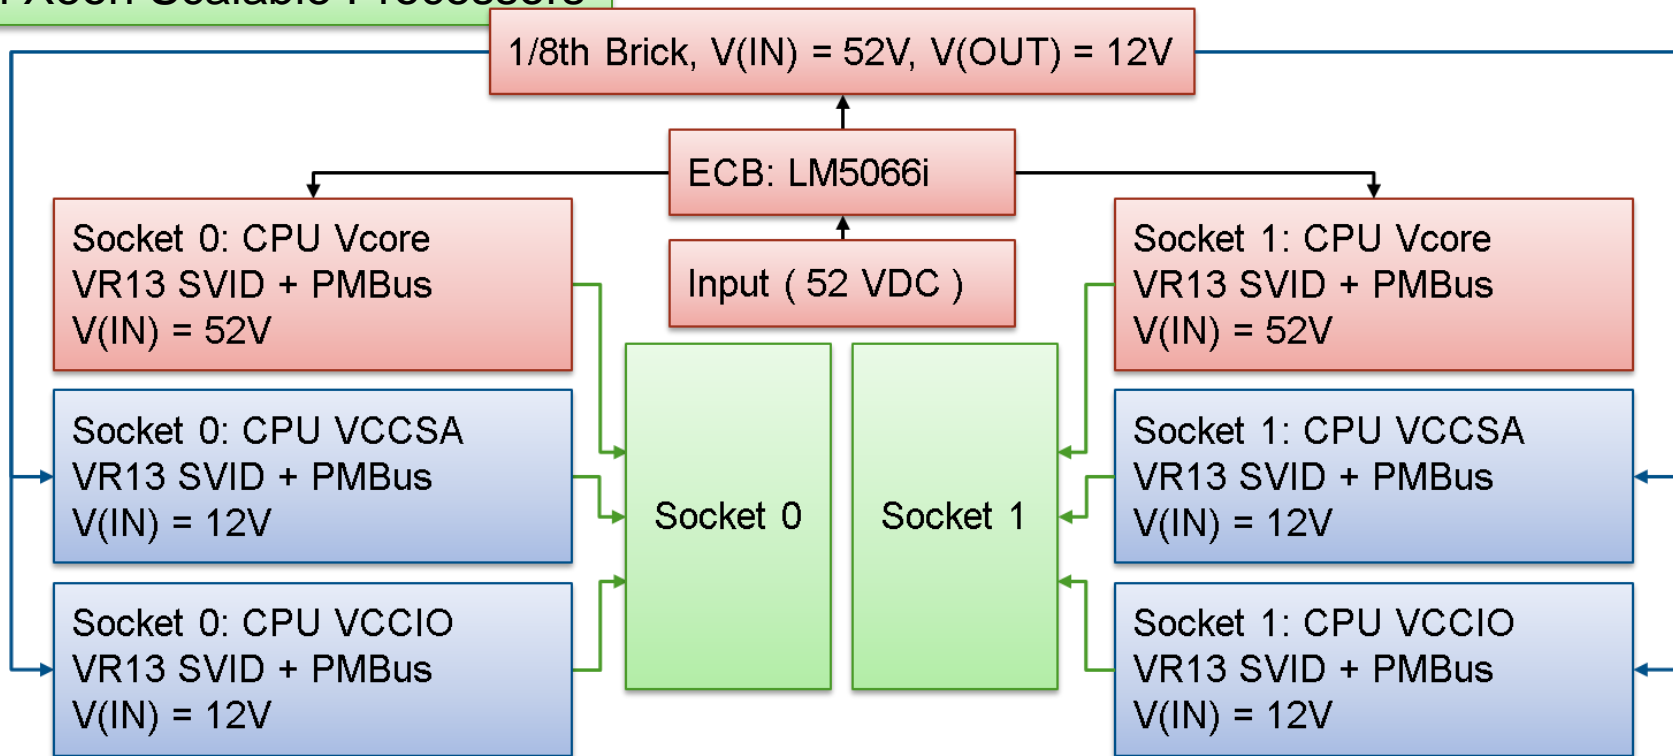
STORE

ANALYZE

Cray XC50 Vicor VCORE Power Solution



Intel Xeon Scalable Processors



COMPUTE

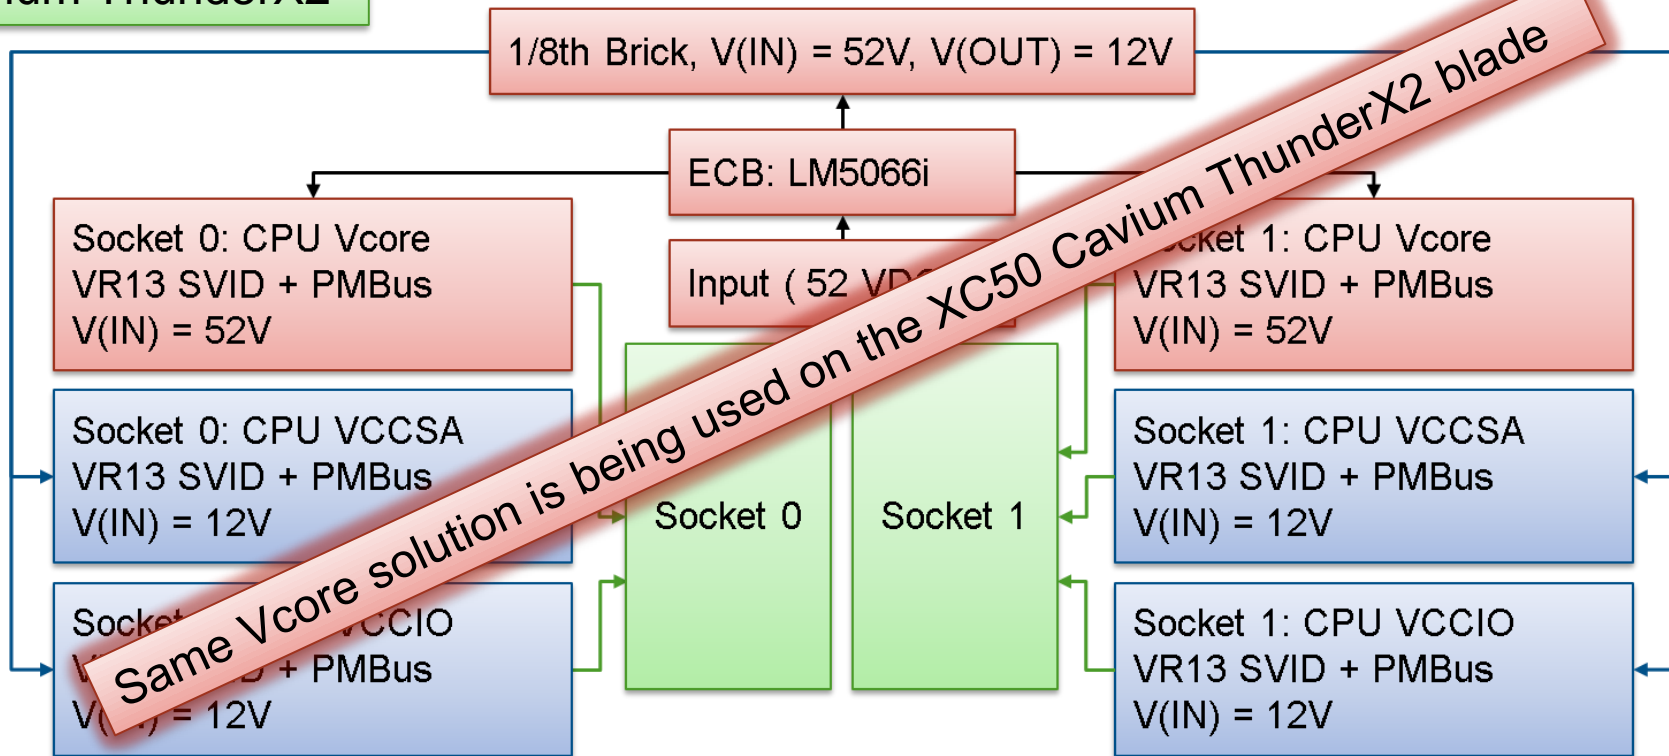
STORE

ANALYZE

Cray XC50 Vicor VCORE Power Solution



Cavium ThunderX2




COMPUTE

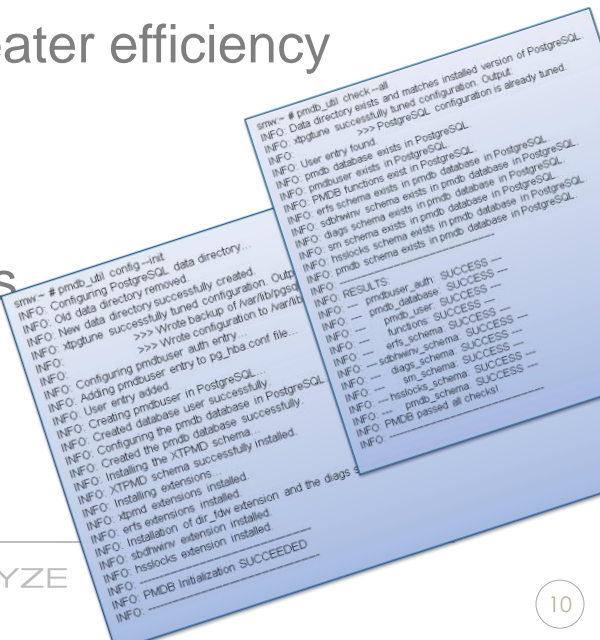
STORE

ANALYZE

Cray PMDB Updates in SMW 8.0.UP06



- **PMDB time series data changed to use TimescaleDB**
 - TimescaleDB is a fully integrated PostgreSQL extension
 - Replaces previous Cray partition based implementation
 - Integration with SQL query planner results greater efficiency
- **New PMDB configuration utility**
 - New pmdb_util utility replaces previous utilities (pmdb_auto_migrate, init_pmdb, etc.)
 - Provides ability to configure and now check PMDB in single place



TimescaleDB in Cray PMDB



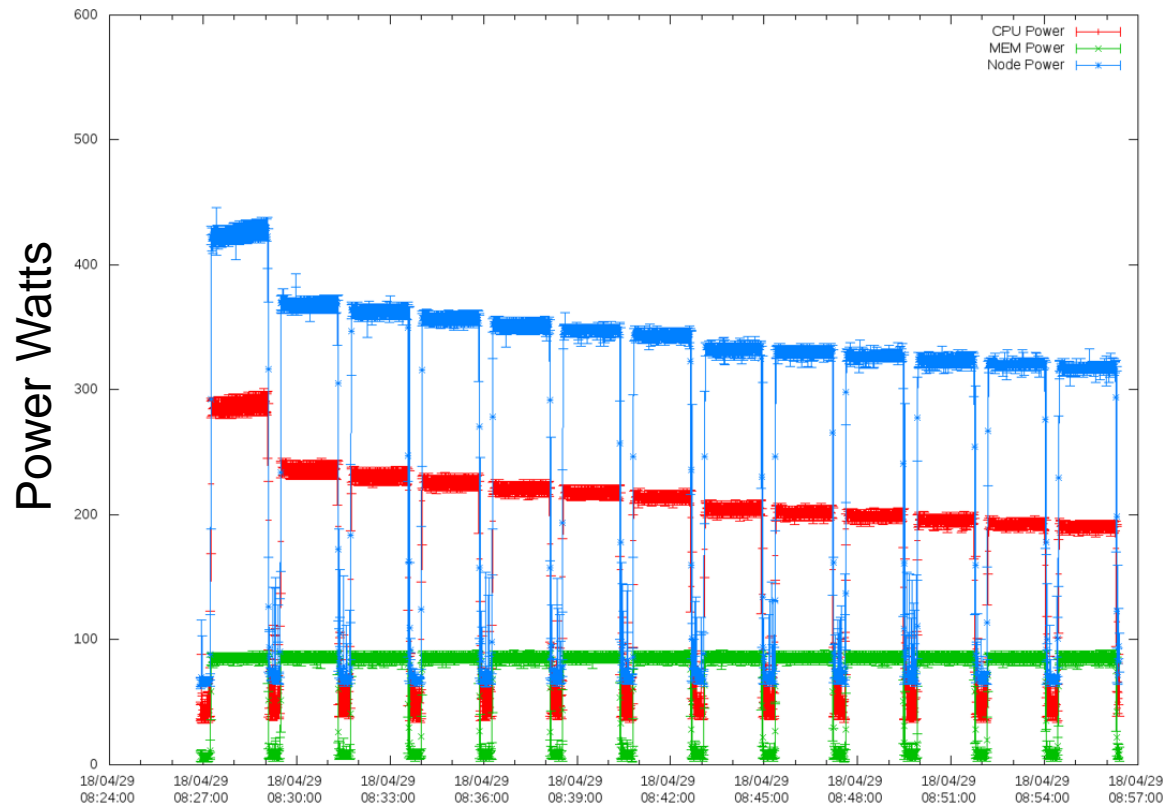
- **Abstracts idea of a single continuous table**
 - Top level “hypertable” is sharded into many partitions/chunks
 - Treated like standard PostgreSQL table
- **Transparently partitions based on time/space**
 - Each chunk contains a known time interval
- **Allows aggressive query optimization**
 - Only chunks that can contain data are ever scanned
 - Speedup of at least 2x with most queries

Cray XC50 APM Usage Examples

Stream runs with supported p-states


Xhpl run with power capping

Stream Runs with Supported P-States



ALPS aprun --p-state=<kHz> at job launch over supported frequencies

Stream Runs Normalized Performance




Stream reported “Best Rate MB/s” normalized over all runs

COMPUTE

STORE

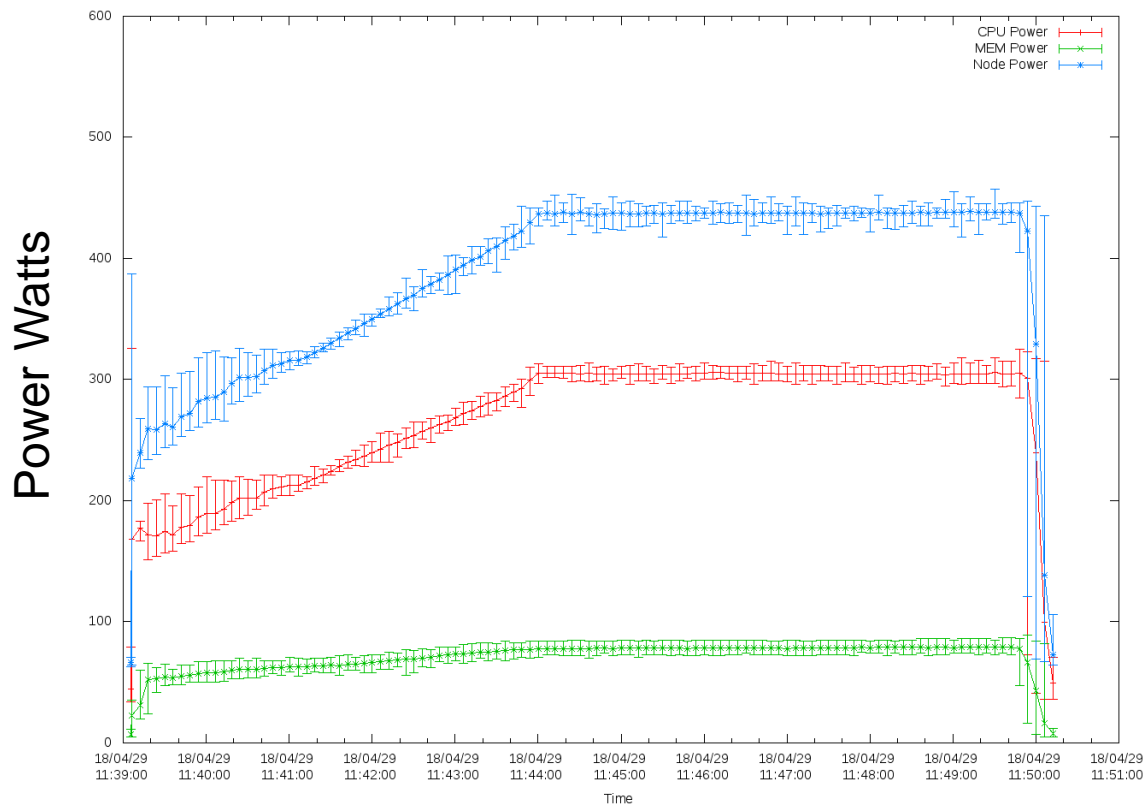
ANALYZE

Stream Run with P-States Sweep



Dynamic p-state limiting with CAPMC of running stream job

Xhpl with Power Capping (Ramp Up)

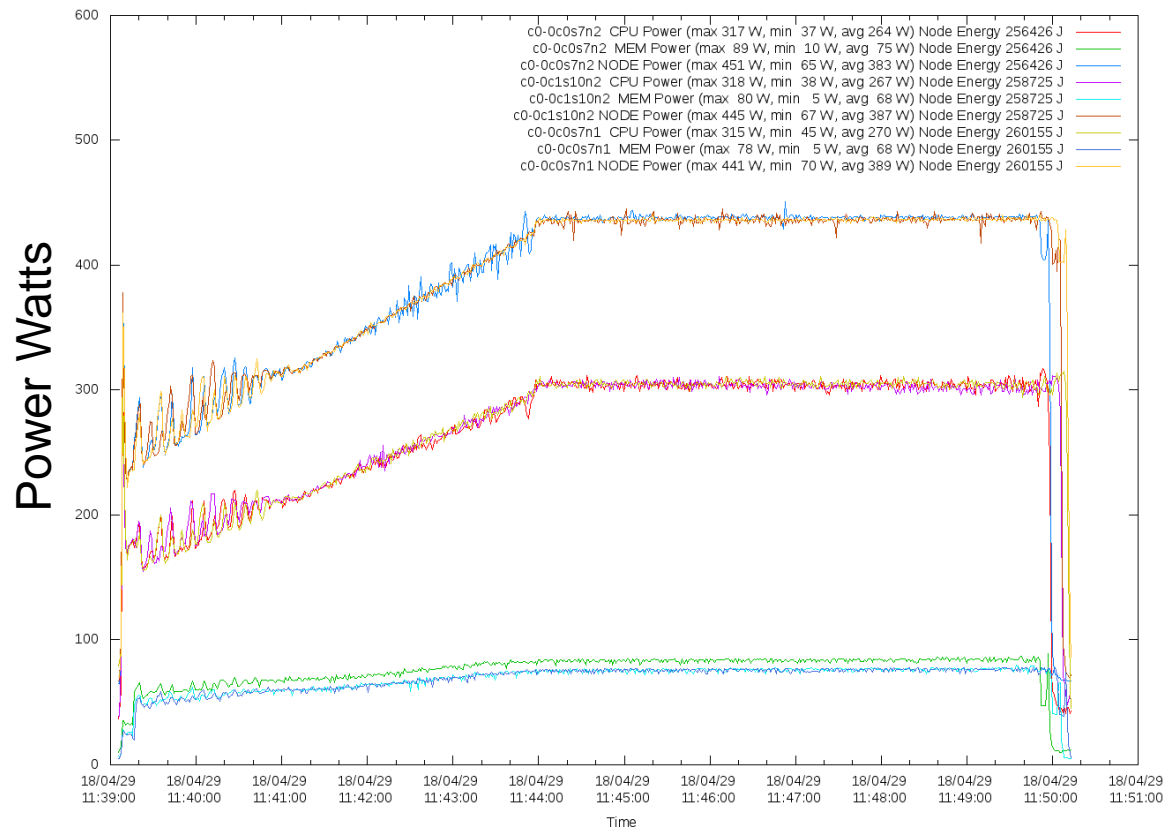


COMPUTE

STORE

ANALYZE

Xhpl with Power Capping (Ramp Up)




COMPUTE

STORE

ANALYZE

Xhpl (Not Power Capped)



COMPUTE

STORE

ANALYZE

Cray APM Supported on All XC Compute Blades



- **Advanced monitoring and control**
 - Telemetry data published into PMDB and CLE:/sys/cray/pm_counters
 - CAPMC enables dynamic monitoring and control
- **New XC50 blades using new 52VDC → VCORE solution**
 - Smaller footprint, full PMAX power
 - More accurate monitoring , and reduced power distribution losses
- **Cray PMDB Updates in SMW 8.0.UP06**
 - PMDB now using TimescaleDB (fully integrated PostgreSQL extension)
 - Utility pmdb_util replaces previous utilities to manage & configure PMDB

Q&A

A scenic view of a historic city skyline, likely Copenhagen, with colorful buildings and a prominent church spire, reflected in the water. The sky is clear blue, and the water is calm, showing a clear reflection of the city.

Steven J. Martin
stevem@cray.com

Legal Disclaimer



Information in this document is provided in connection with Cray Inc. products. No license, express or implied, to any intellectual property rights is granted by this document.

Cray Inc. may make changes to specifications and product descriptions at any time, without notice.

All products, dates and figures specified are preliminary based on current expectations, and are subject to change without notice.

Cray hardware and software products may contain design defects or errors known as errata, which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Cray uses codenames internally to identify products that are in development and not yet publicly announced for release. Customers and other third parties are not authorized by Cray Inc. to use codenames in advertising, promotion or marketing and any use of Cray Inc. internal codenames is at the sole risk of the user.

Performance tests and ratings are measured using specific systems and/or components and reflect the approximate performance of Cray Inc. products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance.

The following are trademarks of Cray Inc. and are registered in the United States and other countries: CRAY and design, SONEXION, URIKA and YARCDATA. The following are trademarks of Cray Inc.: CHAPEL, CLUSTER CONNECT, CLUSTERSTOR, CRAYDOC, CRAYPAT, CRAYPORT, DATAWARP, ECOPHLEX, LIBSCI, NODEKARE, REVEAL. The following system family marks, and associated model number marks, are trademarks of Cray Inc.: CS, CX, XC, XE, XK, XMT and XT. The registered trademark LINUX is used pursuant to a sublicense from LMI, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis. Other trademarks used on this website are the property of their respective owners.