

# SANDIA REPORT

SAND2013-9639  
Unlimited Release  
Printed November 2013

## Qualification for PowerInsight Accuracy of Power Measurements

David DeBonis, James H. Laros III, Kevin Pedretti

Prepared by  
Sandia National Laboratories  
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Approved for public release; further dissemination unlimited.



**Sandia National Laboratories**

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

**NOTICE:** This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government, nor any agency thereof, nor any of their employees, nor any of their contractors, subcontractors, or their employees, make any warranty, express or implied, or assume any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represent that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government, any agency thereof, or any of their contractors or subcontractors. The views and opinions expressed herein do not necessarily state or reflect those of the United States Government, any agency thereof, or any of their contractors.

Printed in the United States of America. This report has been reproduced directly from the best available copy.

Available to DOE and DOE contractors from  
U.S. Department of Energy  
Office of Scientific and Technical Information  
P.O. Box 62  
Oak Ridge, TN 37831

Telephone: (865) 576-8401  
Facsimile: (865) 576-5728  
E-Mail: [reports@adonis.osti.gov](mailto:reports@adonis.osti.gov)  
Online ordering: <http://www.osti.gov/bridge>

Available to the public from  
U.S. Department of Commerce  
National Technical Information Service  
5285 Port Royal Rd  
Springfield, VA 22161

Telephone: (800) 553-6847  
Facsimile: (703) 605-6900  
E-Mail: [orders@ntis.fedworld.gov](mailto:orders@ntis.fedworld.gov)  
Online ordering: <http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online>



# Qualification for PowerInsight Accuracy of Power Measurements

David DeBonis, James H. Laros III, Kevin Pedretti

## **Abstract**

*Accuracy of component based power measuring devices forms a necessary basis for research in the area of power-efficient and power-aware computing. The accuracy of these devices must be quantified within a reasonable tolerance. This study focuses on PowerInsight, an out-of-band embedded measuring device which takes readings of power rails on compute nodes within a HPC system in realtime. We quantify how well the device performs in comparison to a digital oscilloscope as well as PowerMon2. We show that the accuracy is within a 6% deviation on measurements under reasonable load.*



# Contents

<b>1 Qualification Report</b>	<b>9</b>
Introduction .....	10
Equipment .....	11
Methods .....	12
Data .....	15
Results & Discussion .....	18
Conclusion .....	20
<b>References</b>	<b>21</b>

# List of Figures

1.1	Configuration for Ground Truth experiment .....	12
1.2	Configuration for PowerInsight experiment .....	13
1.3	Configuration for PowerMon2 experiment .....	14
1.4	Comparison Plot for Rail Amperage at Idle .....	15
1.5	Comparison Plot for Rail Voltage at Idle .....	16
1.6	Comparison Plot for Memory Rail Reading at Idle and Under Load .....	18

# List of Tables

1.1	Comparison of PowerInsight at Idle over Power Rails . . . . .	15
1.2	Comparison of PowerMon2 at Idle over Power Rails . . . . .	16
1.3	Comparison of PowerInsight measuring the CPU Power Rails (12V isolated) .	16
1.4	Comparison of PowerInsight measuring the Memory Power Rails (5V mother- board) . . . . .	17
1.5	Comparison of PowerMon2 measuring the Memory Power Rails (5V mother- board) . . . . .	17



# Chapter 1

## Qualification Report

The following report describes experiments that were performed to verify the correct operation of the PowerInsight device that was developed by Penguin Computing in collaboration with Sandia National Laboratories. PowerInsight is an embedded system that can be added to commodity off-the-shelf computers to provide advanced power measurement capabilities. This report describes the approach that was used for the qualification experiments and the results that were obtained.

# Introduction

The purpose of this experiment is to qualify the measurement accuracy of the PowerInsight [3] device, designed and developed in cooperation with Sandia National Laboratories and Penguin Computing. Our expectation is that PowerInsight will provide accurate power measurement when the host server is under load (i.e., running an application) and reduced accuracy when the host server is idle, which results in currents outside of the range that PowerInsight was designed to measure. A key objective is to quantify this threshold as well as the percentage deviation from ground truth measurements (i.e., measurements taken with precision T&M equipment like a digital oscilloscope) both above and below this threshold. The threshold will be defined as the point at which the accuracy is within 5 percent of the actual value. A subset of the experiments will be repeated for the PowerMon2 device [1], which provides similar functionality as PowerInsight.

# Equipment

We will use an instrumented node of one of the Advanced Architecture Testbeds [2] at Sandia National Laboratories for our tests. The nodes are 2-U computing systems with AMD A10-5800K Fusion APU containing four Piledriver x86 cores and a 384-core 800MHz Radeon GPU, 32GB memory, QLogic and GigE NIC. The instrumentation is the PowerInsight device which is being qualified, consisting of a BeagleBone embedded processing platform with a TI Arm-Cortex processor and custom ADC cape and sensor harness designed and manufactured by Penguin Computing.

The instrumentation that will be used to quantify ground truth for current and voltage measurements will be an Owon SDS7102 Deep Memory Oscilloscope with an Amek SL261 Amp Probe. An additional custom IV tap is needed to put inline with the power rails to allow for accessibility to measurements.

## Methods

The experiment will consist of three configurations, each building on the other to establish a foundation for comparison. These experiments will be conducted with the testbed node in idle state. Next, PowerInsight will be tested at both idle state and under heavy load for both CPU and memory. CPU load will be tested using the LAMMPS application in serial mode with Rhodo input and memory testing will be with the hpcCG application using single node OpenMP with a matrix size of 125x125x125.

The first experiment (see Figure 1.1) will be with the power supply of the Teller node tapped for measuring voltage and current manually using the oscilloscope mentioned above. Readings over the course of 10 seconds will be taken at a sampling frequency of 500 samples per second, with a calculated average for both voltage and amperage.

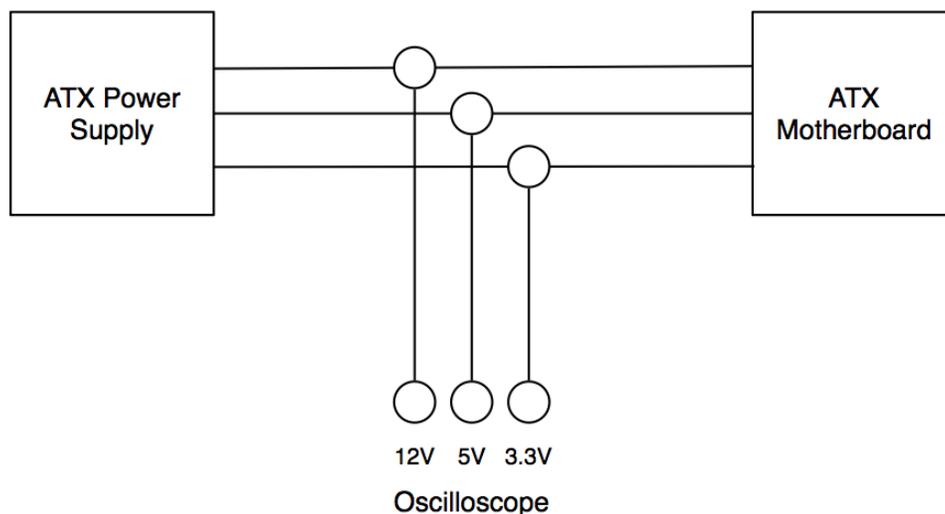


Figure 1.1: Configuration for Ground Truth experiment

The second experiment (see Figure 1.2) will be with both the oscilloscope (as configured in the first experiment) and the PowerInsight device inline between the power supply and motherboard. Reading over the course of 10 seconds will be taken at a sampling frequency of 500 samples per second using the PowerInsight device, gathering both voltage and current, with a manually calculated average.

The third experiment (see Figure 1.3) will be using the PowerMon2 device and oscilloscope in series between the power supply and motherboard. Readings over the course of 10 seconds will be gathered at a sampling frequency of 500 samples per second using the PowerMon2 device, gathering raw voltage and current, manually calculating the average for each.

The above three scenarios will be used to gather measurements over the three main power rails off of the ATX harness; 3.3V, 5V, and 12V. Once this trial is done, a second pass will be

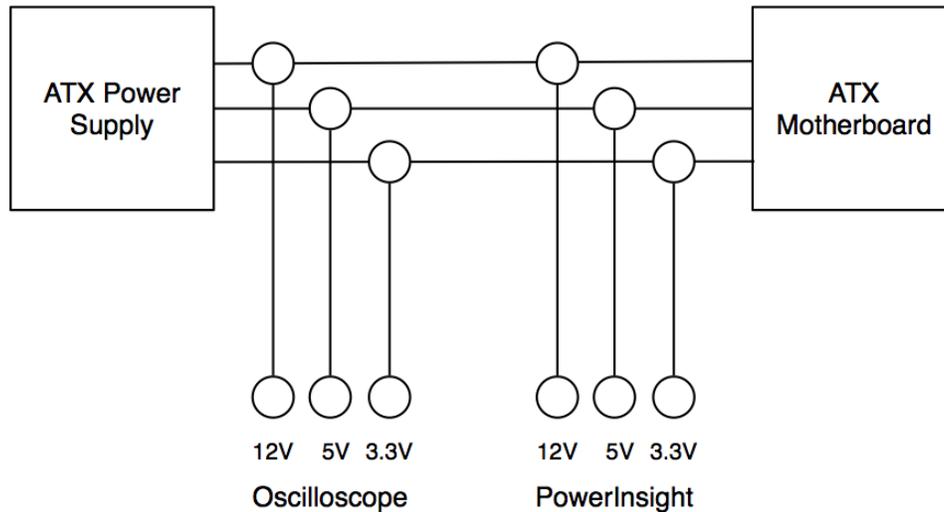


Figure 1.2: Configuration for PowerInsight experiment

performed on another testbed node instrumented in the same fashion to examine differences between nodes and devices.

The experiments up to now will be used to verify that each instrument is gathering statistically accurate readings independent of whether other devices are inline with the sensors. The reason for this is that the PowerMon2 device uses a shunt resistor which results in a slight load on each of the rails whereas the PowerInsight is decoupled from the circuit through a hall-effect sensor.

Once these experiments have been completed, all readings shall be taken in parallel on all test equipment. This will allow us to avoid any variability between runs since all devices will be monitoring the same events at the same time. After again gathering data for the low-amperage idle power rails of the motherboard, we will proceed in testing the CPU and memory rails at both idle and under load (as mentioned above) while monitoring the PowerInsight (and PowerMon2 for memory only) and comparing it to our oscilloscope readings. For these tests both devices and oscilloscope will be configured to sample for 10 seconds at a rate of 500 samples per second. We would have also collected CPU readings from PowerMon2 but we did not have the proper cable for this particular motherboard to do this (the P1 connector was a four pin connector on PowerMon2 but modern ATX motherboards use an eight pin connector). The memory power rails are the same as the 5V motherboard power rails, while the CPU power rails are isolated 12V power rails and independent of the 12V motherboard power rails used for the earlier experiments.

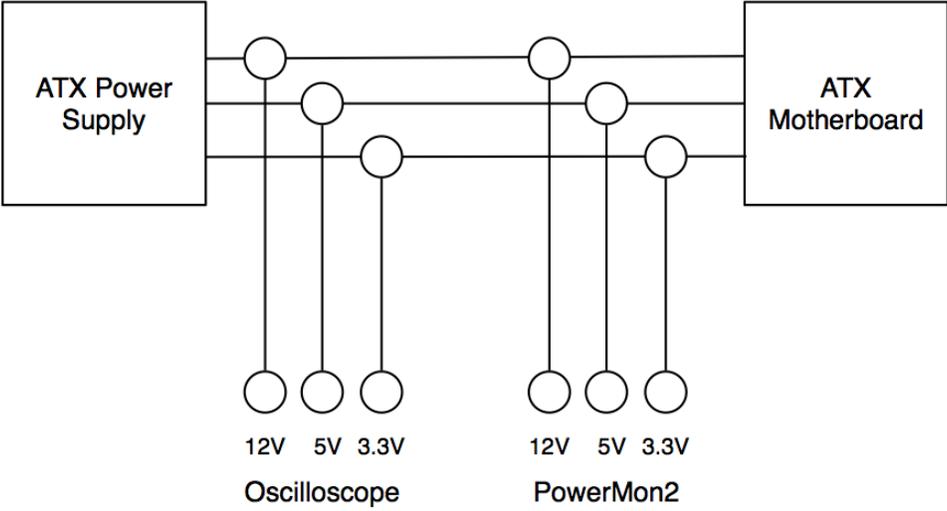


Figure 1.3: Configuration for PowerMon2 experiment

# Data

Figure 1.4 through 1.5 give a graphical summary of the data collected for the idle states of the three power rails of the motherboard while the system is at idle. These numbers show readings from all three measurement devices; Oscilloscope, PowerInsight, and PowerMon2.

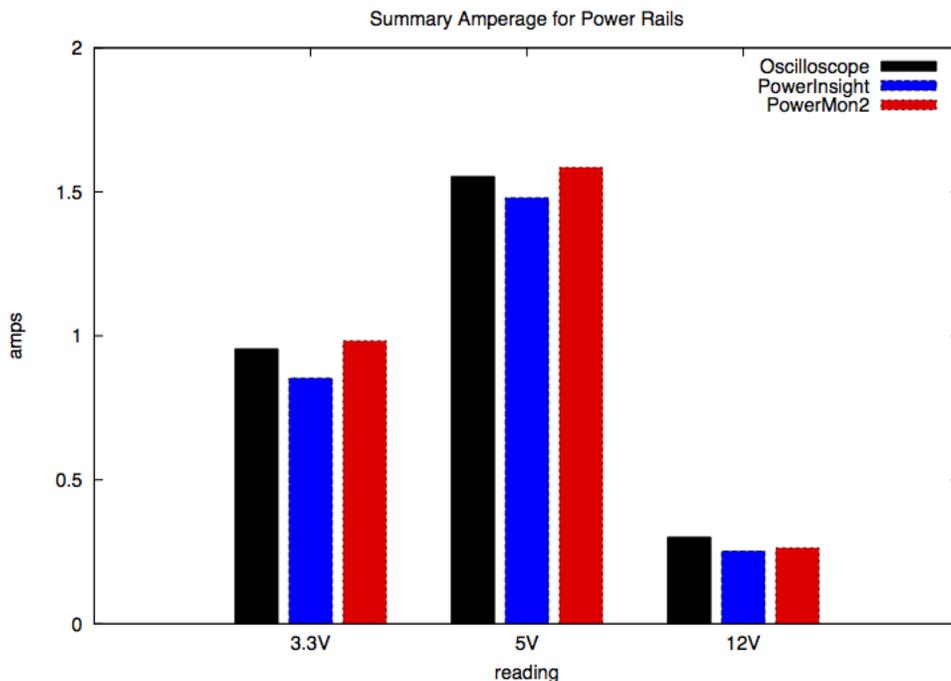


Figure 1.4: Comparison Plot for Rail Amperage at Idle

Table 1.1 through 1.2 give a summary of the data gathered from the experiments discussed in the methods section for idle states of the testbed node using PowerInsight and PowerMon.

	Ground Truth		PowerInsight		Deviation		
	V	A	V	A	V	A	W
3.3V	3.403	0.9548	3.286	0.8533	-3%	-11%	-14%
5V	5.224	1.553	5.047	1.479	-3%	-5%	-8%
12V	12.50	0.3010	12.14	0.2525	-3%	-16%	-19%

Table 1.1: Comparison of PowerInsight at Idle over Power Rails

Table 1.3 through 1.5 gives a summary of the data gathered from the experiments discussed in the methods section for both idle and load states of the testbed node using PowerInsight and PowerMon2 to monitor memory and PowerInsight to monitor CPU.

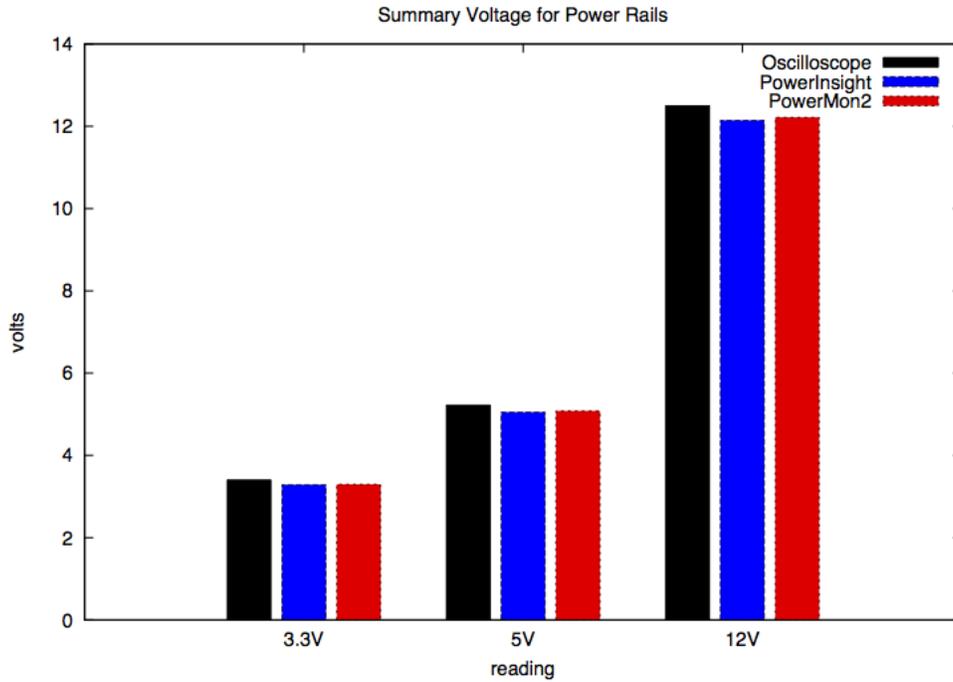


Figure 1.5: Comparison Plot for Rail Voltage at Idle

	Ground Truth		PowerMon2		Deviation		
	V	A	V	A	V	A	W
3.3V	3.403	0.9548	3.297	0.9823	-3%	3%	0%
5V	5.224	1.553	5.077	1.585	-3%	2%	-1%
12V	12.50	0.3010	12.21	0.2637	-2%	-12%	-14%

Table 1.2: Comparison of PowerMon2 at Idle over Power Rails

	Ground Truth		PowerInsight		Deviation		
	V	A	V	A	V	A	W
idle	12.44	1.308	12.13	1.248	-2%	-5%	-7%
load	12.43	3.605	12.12	3.534	-2%	-2%	-4%

Table 1.3: Comparison of PowerInsight measuring the CPU Power Rails (12V isolated)

	Ground Truth		PowerInsight		Deviation		
	V	A	V	A	V	A	W
idle	5.224	1.553	5.047	1.479	-3%	-5%	-8%
load	5.214	2.841	5.036	2.669	-3%	-6%	-9%

Table 1.4: Comparison of PowerInsight measuring the Memory Power Rails (5V motherboard)

	Ground Truth		PowerMon2		Deviation		
	V	A	V	A	V	A	W
idle	5.224	1.553	5.077	1.585	-3%	2%	-1%
load	5.214	2.841	5.052	2.866	-3%	1%	-2%

Table 1.5: Comparison of PowerMon2 measuring the Memory Power Rails (5V motherboard)

## Results & Discussion

The baseline measurements using only the oscilloscope were compared to the measurements with or without either the PowerInsight or PowerMon2 inline. Thus the remaining values gathered were with both PowerInsight and PowerMon2 inline since having either one did not perturb the values of the other. The readings were taken in parallel so that the shape of the waveform can be confirmed to align with each other and that any anomalous values would be seen on all test equipment equally (see Figure 1.6). The results are given as compared to the ground truth values in Table 1.1 and Table 1.2.

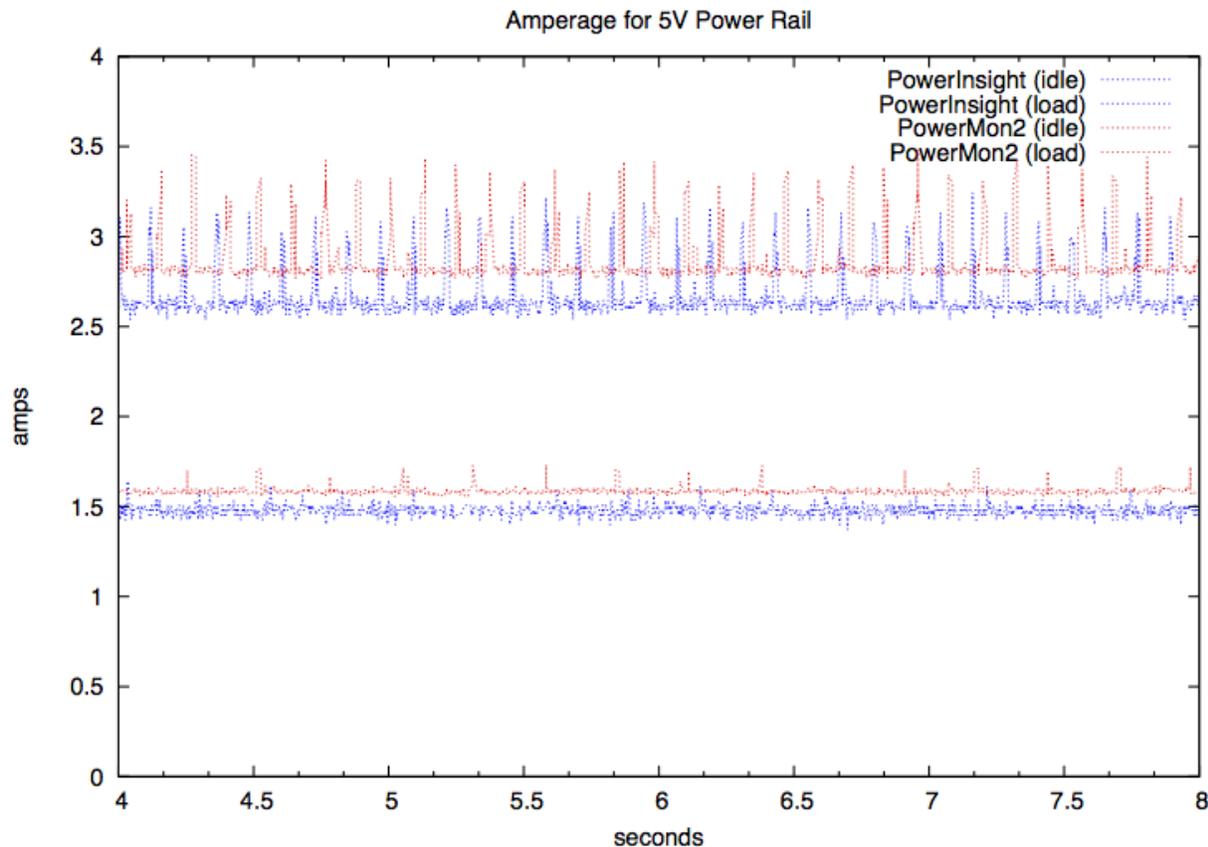


Figure 1.6: Comparison Plot for Memory Rail Reading at Idle and Under Load

Comparing readings between PowerInsight and the oscilloscope, we see that a deviation between 5% to 16% over the main power rails to the motherboard. This inaccuracy is likely due to the design of PowerInsight, which according to the manufacturer was not designed to measure currents below 1.5A. Looking at the readings from PowerMon2 we see a similar inaccuracy on the 12V rail operating at a low current, though the impact is barely noticeable on the 3.3V and 5V rails.

After discussing the accuracy with Penguin Computing, the device is intended to be effective within the range of 1.5 amps to 18 amps. The primary component on the PowerInsight

which accounts for its accuracy is a combination of ADC and operational amplifier which supplies the reference voltage for sampling. Key to accuracy is the precision of the op amp to divide VCC to 1/10th its value and the way in which this value is represented in the drivers conversion equations. On further inspection of the Allegro op amp device, Penguin Computing sampled a variation between chips such that the internal trim resistors of the chip had a manufacturing tolerance that would allow the reference voltage to be off by up to 6%. Future plans are to calibrate each sensor prior to delivery so that a custom configuration file could be used to improve the accuracy of the conversion equations.

The difference between idle and load measurements is quite apparent in our experiment. The results are consistent over multiple runs and repeatable. We see that the deviation is kept to below 6% over both sets of runs though the accuracy of the PowerMon2 device is superior.

## Conclusion

As anticipated, measurements using the PowerInsight on rails exhibiting low current draw deviated substantially (greater than 5%). This was true of the PowerMon2 measurements as well due to the susceptibility to noise and the values being so low that deviations are exaggerated when looking at pure percentage of measurement.

When analyzing the CPU and memory, which have higher current draws, we indeed observe improved accuracies for both PowerInsight and PowerMon2. Under load, both devices produce measurements that are within 6% of the oscilloscope measurements, which is close to our 5% accuracy target.

# References

- [1] D. Bedard, Min Yeol Lim, R. Fowler, and A. Porterfield. Powermon: Fine-grained and integrated power monitoring for commodity computer systems. In *IEEE SoutheastCon 2010 (SoutheastCon), Proceedings of the*, pages 479–484, 2010.
- [2] Sandia National Laboratories. Advance systems technology test beds, October 2013.
- [3] James H. Laros, David DeBonis, and Phi Pokorny. *PowerInsight - A Commodity Power Measurement Capability*. Apr 2013.

## DISTRIBUTION:

- 1 MS 0899 Technical Library, 9536 (electronic copy)



