


# Energy Efficiency Investigation for the Magellan Super Computer at the National Energy Research Scientific Computing Center



Final Project Report

Prepared for: Brent Draney, NERSC

Prepared by: Steve Greenberg, Henry Coles

Lawrence Berkeley National Laboratory  
One Cyclotron Road  
University of California  
Berkeley, California 94720

March 1, 2012

LBL# - xxxxx

## **DISCLAIMER**

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

## **Acknowledgements**

The authors would like to thank

Jim Crow, Thomas Davis, Brent Draney, Jeff Grounds, John Hutchings and Howard Walter for their assistance in data acquisition and facility access. In addition the authors appreciate the use of equipment from the Pacific Energy Center Tool Lending Library.

## Contents

Introduction .....	4
System Description.....	4
Methods .....	4
As-Found Analysis .....	5
As-Found Results .....	6
Generic Performance Analysis .....	6
Generic Performance Results .....	7
Conclusions .....	7
References .....	8
Appendix A: Magellan Data – Oct 19, 2011 2-4PM.....	9
Appendix B: PDU and 12kV Transformer Power .....	10
Appendix C: As-Found Analysis Diagram.....	11
Appendix D: Whole Data Center Analysis Diagram.....	12
Appendix E: pPUE Calculation Details.....	13
Appendix F: Calculation Detail Notes.....	14
Appendix G: Taylor Engineering Water-side Economizer Chilled Water Plant Model With and Without Secondary Pumping .....	16
Appendix H: Power Level Comparison.....	17
Appendix I: Reference Only – Simplified PUE .....	18

## Introduction

LBNL's Oakland Scientific Facility (Building 943), a high performance computing (HPC) center also known as NERSC (National Energy Research Scientific Computing Center) currently contains five actively used super computers including the Magellan system. Magellan uses a unique cooling design consisting of rear door coolers. This paper investigates the energy use efficiency of Magellan with a focus on the efficiency of the cooling design as it compares to the total site energy use efficiency. We look at both the "as-found" installation at OSF and a hypothetical one with a full center of IT equipment cooled by rear doors.

## System Description

The energy use rate (power) at NERSC - for all systems and supporting infrastructure including cooling - varies, but is relatively constant compared to systems found within other LBNL data centers. The low variability of the load is due to a high computing demand that creates a backlog of jobs. The typical energy use rate is 4.7MW for the total IT load; the corresponding PUE<sup>TM1</sup> (Power Usage Effectiveness, Green Grid 2011) varies but is in the range of 1.26 to 1.29. Therefore the total site energy-use rate for infrastructure and cooling averages approximately 1.3MW. The Magellan IT equipment consumes approximately 290kW of the 4.7MW IT site total or approximately 6 percent of the total site IT power.

The Magellan computer system, an IBM iDataPlex system with 400 compute nodes, has a cooling system unique compared to the other super computer systems located at NERSC. The unique feature includes passive (no added fans) heat exchangers, referred to as rear door coolers, attached directly to the rear of the racks holding the compute node equipment. As will be presented later the majority of the heat generated by the compute nodes is removed by the rear door coolers.

## Methods

Two attributes highlight the key energy-saving design differences used for Magellan:

- The hot air exhausted by the IT equipment located in the Magellan frames is immediately cooled by air-to-water heat exchangers (rear-door coolers) attached at the back of each rack. This compares to the other systems that rely on CRAH's (computer room air handlers) or low-pressure refrigerant-to-water heat exchange cooling systems. The close proximity of the heat exchanger to the hot IT equipment exhaust air exposes the heat exchanger to higher temperatures compared to air exposed to heat exchangers in a CRAH type cooling design. In addition, some fan power savings are found because the rear door coolers do not have fans as would be found in a CRAH cooling environment. The IT exhaust air is pushed through the rear door coolers using the fans internal to the IT equipment. In general for rear door coolers an increase in IT equipment fan power may offset this advantage somewhat. However, for Magellan the IT equipment fans ran at minimum speed so no additional fan power was observed during our investigation.
- The performance of the rear door coolers allows using cooling water that is much warmer (~54F) than the chilled water available at this facility which is currently set at ~45F. To take advantage of and investigate the use of warm water cooling the Magellan rear doors are fed by two CDU's (cooling distribution units) that use the warm return water from CRAH's used

---

<sup>1</sup> PUE = Total Facility Energy divided by the IT Equipment Energy

to cool an adjacent super computer. The advantage of using higher temperature cooling water is substantial and is investigated using a hypothetical example where all the computer systems at a given site use the rear door cooler design.

We investigate two questions:

Question -1: How does the partial PUE (pPUE<sup>2</sup> -The Green Grid Copyright 2011) of the Magellan system compare to the total PUE for the NERSC site?

*As-Found pPUE*

Method: This question is investigated by measuring as-found power and cooling levels associated with the Magellan system and calculating a pPUE. This value is then compared to the total site as-found PUE reported earlier (Greenberg and Coles 2011).

Question -2: What is the generic power usage effectiveness (PUE) of a site similar to NERSC if fitted only with the Magellan type IT and cooling design in a stated location?

*Generic pPUE*

Method: Calculate a pPUE for a hypothetical system where the rear door coolers provided 100% of the needed cooling for the IT power. In addition a model of a chilled water plant located in Santa Clara CA is used as the stated location in order to forecast the chilled water plant performance when producing only the higher temperature cooling water.

## As-Found Analysis

### As-Found partial Power Usage Effectiveness (pPUE) Method Description

The process used to estimate the as found pPUE (Question 1):

- measure the electrical power at the load side of the IT PDU's
- measure the CDU specific electrical power and cooling supplied to the rear doors
- estimate the performance (kW/ton) for the plant supplying the cooling water for the CDU's.
- calculate the pPUE, see Appendices C and E.

Measurements were taken on October 19, 2011 from 2 to 4PM to gather enough data to support the pPUE calculation. The CDU operating conditions are not recorded in a database so measurements were taken over a small time period of approximately 2 hours and used for the calculations.

The measurements are from the PDU's (power distribution units) feeding all the IT equipment assigned uniquely to the Magellan system, see Appendix A. This allows a calculation of the IT power, IT PDU loss, 12kV transformer loss, and utility supplied power, see Appendix B. Note there is approximately 240 watts of UPS loss included in the PDU loss. In addition thermal and electrical power measurements were collected for the two CDU units that supply cooling water to the Magellan rear door coolers.

---

<sup>2</sup> pPUE = Total Energy within a boundary divided by the IT Equipment Energy within that boundary

### Calculation Method for Magellan pPUE:

1. Sum Magellan IT power from the PDU measurements.
2. Determine the transformer losses using found and assumed efficiencies.
3. Estimate the heat removed by the rear-door coolers using calibrated temperature data measured at both rear door CDU's.
4. Calculate the remaining heat removed by the CRAH's – includes all power inputs except the chiller plant and the IT power not cooled by the CDU's.
5. Estimate the fan power needed by the CRAH's as a function of the total CRAH heat removed.
6. Determine the power needed by the chiller plant for the CDU's that doesn't include power required for secondary pumping.
7. Using subtraction, estimate the power needed by the chiller plant for the remaining cooling (provided by CRAH's). The chiller plant performance includes secondary pumping energy.
8. Calculate the total electrical power needed to support the Magellan system including the IT power and the total pro-rated chiller plant power.
9. Divide result from 8. by result from 1. to determine the pPUE ratio.

Our calculations show the rear door coolers are extracting approximately 77 percent of Magellan IT power. CDU#1 removes approximately 43% and CDU#2 removes 35% of the Magellan IT power generated heat. The pPUE is calculated at 1.25; see Appendices E and F for calculation details and notes. An investigation into the piping layout points to the basement plant receiving the heat from the Magellan CDU's. The electrical power needed to generate the cooling water was estimated at 0.66 kW per ton for the basement plant with secondary pumping included and 0.64kW per ton without secondary pumping included.

### As-Found Results

We calculated the pPUE for the Magellan super computer to be 1.25. As expected, this value is lower than the total site PUE of 1.28 as reported (Greenberg and Coles 2011). Appendix H shows IT power levels for Magellan are very similar for the two periods of data used for comparison. Appendix I shows a PUE near 1.24 during the measurement period with limitations, when 0.04 is added the PUE is 1.28 as stated above.

### Generic Performance Analysis

To answer Question 2, the energy usage effectiveness (PUE) of a data center was estimated as if all the IT equipment was cooled using the same cooling design currently used for Magellan.

Key assumptions:

- The IT power is assumed to be the same as the "As Found" analysis.
- Rear door heat exchangers provide 100% of the IT equipment cooling.
- Chilled water plant supplies 55.5°F instead of the ~45°F supplied in as-found analysis.
- The stated location is Santa Clara CA, the climate data from this location is used in the plant model supplied by Taylor Engineering (Coles 2010). It should be noted that different climates will produce different results.

- IT PDU and 12kV transformer efficiencies are 99% and 98% respectively. This is consistent with (Greenberg and Coles 2011).

The calculation follows the same method as that used to arrive at the results for question 1, see Appendix E and F.

### **Generic Performance Results**

The pPUE for the generic case is calculated as 1.18. It is expected that the pPUE will be lower because the supplied cooling water is higher 53.5°F compared to ~45°F. The use of rear door coolers is more energy efficient compared to CRAH cooling. This follows the published results on the Chilloff-2 (Coles 2010) project. It is important to note that this generic case is using Santa Clara CA climate data and therefore should not be compared directly to the "As-Found" case using Oakland weather.

### **Conclusions**

The partial PUE of the as-found condition is only marginally lower than the whole-center PUE. Even though the CDU/rear door combination provides cooling more efficiently than the CRAH arrangement, the overall effect isn't large.

The PUE of an all-rear-door equipped data center would be significantly better, with the main contribution being improved chiller plant efficiency (roughly 1/3). Overall results will vary depending largely on electrical distribution losses. This improvement is due both to an increase in chilled water supply temperature and the use of a water-side economizer when conditions allow.

## References

Azevedo, D., J. Cooley, M. Patterson and M. Blackburn, 2011. "Data Center Efficiency Metrics: mPUE™, Partial PUE, ERE, DCcE, The Green Grid, Technical Forum 2011, [http://www.thegreengrid.org/~media/TechForumPresentations2011/Data\\_Center\\_Efficiency\\_Metrics\\_2011.pdf?lang=en](http://www.thegreengrid.org/~media/TechForumPresentations2011/Data_Center_Efficiency_Metrics_2011.pdf?lang=en)

Coles, H.C. 2010. "Demonstration of Rack-Mounted Computer Equipment Cooling Solutions", Report to the California Energy Commission Public Interest Energy Research (PIER) Program.

Greenberg, S., H. Coles, 2011. "PUE Measurement Verification", internal LBNL memo Nov 8, 2011

Taylor Engineering 2010. Personal communication with Mark Hydeman 3/17/2010



# Appendix A: Magellan Data – Oct 19, 2011 2-4PM

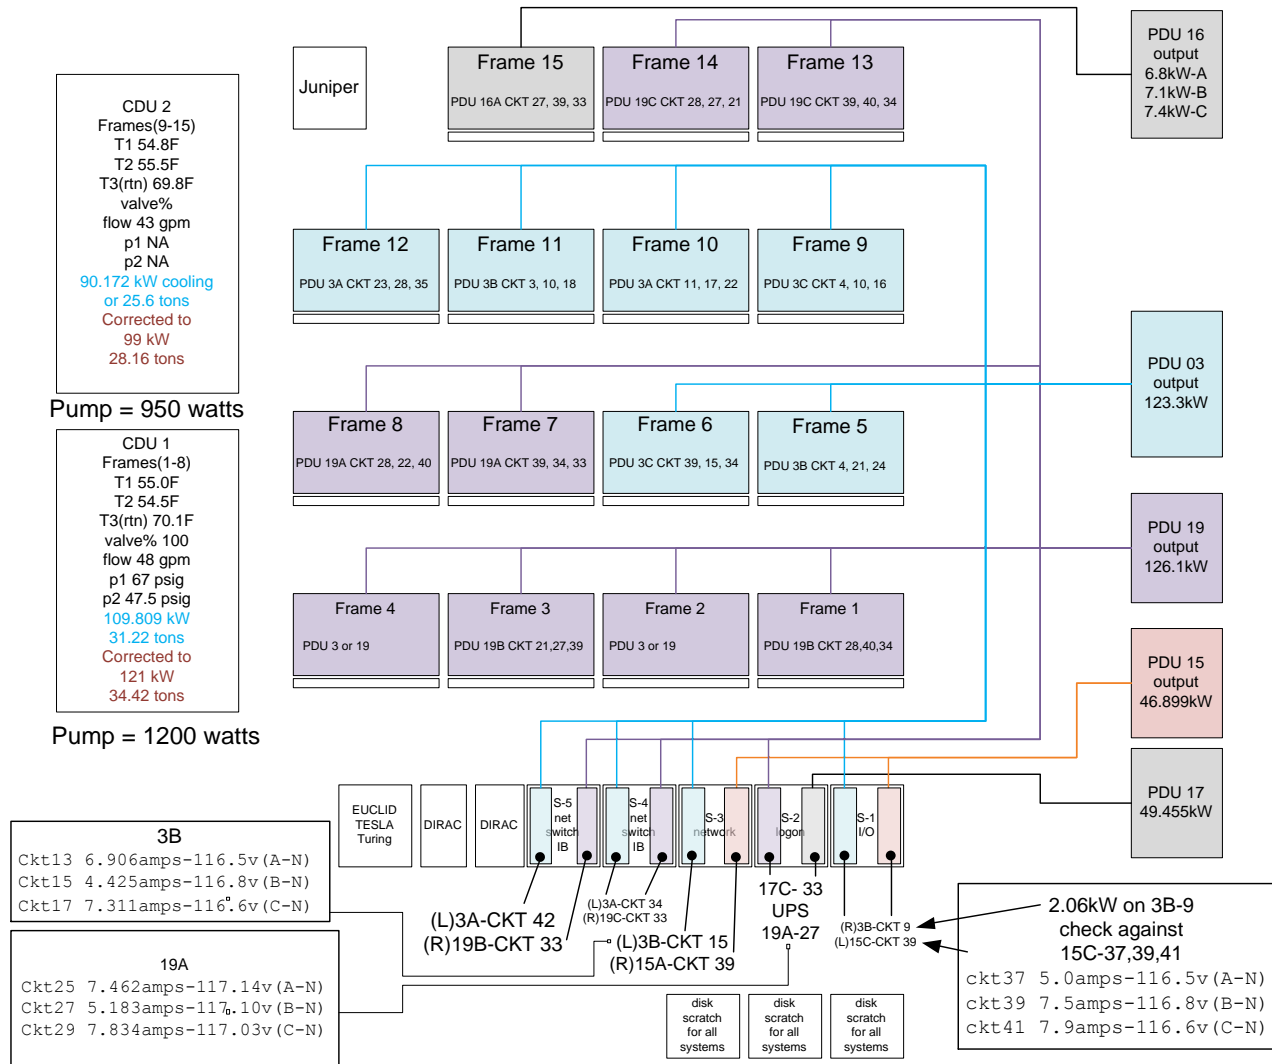


Figure A-1: Power and Thermal Data Recorded for Magellan – Oct 19, 2011 2-4PM

## Appendix B: PDU and 12kV Transformer Power

### IT Room NERSC Magellan Power Measurements

Recorded Oct 19, 2011 2-4PM

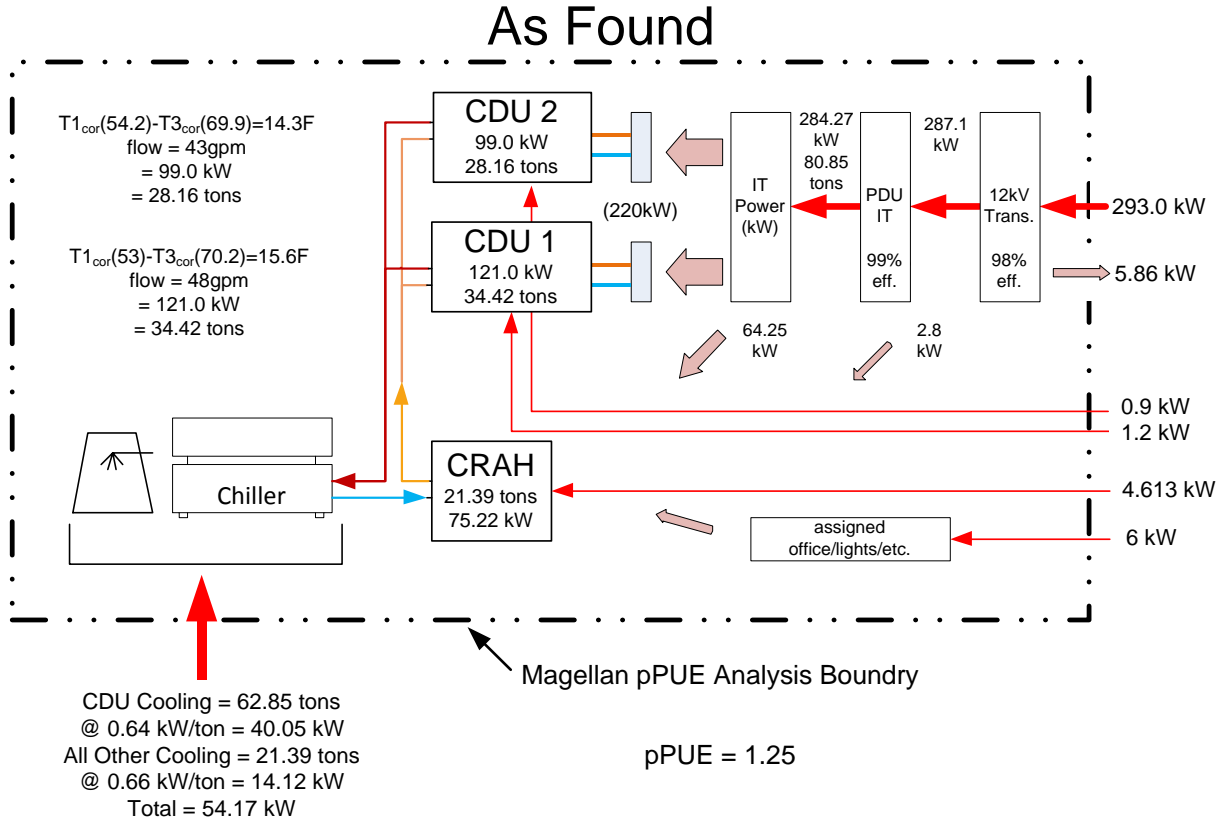
PDU Output Location	IT Power (kW)
Output of PDU 19 to Magellan Frames 1-4, 7, 8, 13, 14, S-2, S-4, S-5 Measured at PDU Panel	126.1
Output of PDU 3 to Magellan Frames 5, 6, 9-12, S-1, S-3, S-4, S-5	123.3
Output of PDU 16 to Magellan Frame 15 PDU 16A CKT 27 6.8kW PDU 16A CKT 39 7.1kW PDU 16A CKT 33 7.4kW	21.3
Output of PDU 15 to S-1 2.38 kW 2.06 kW	4.44
Output of PDU 15 to S-3 2.174 kW x 2	4.35
Output of PDU 17 to S-2 2.389 kW x 2	4.78
<b>Total</b>	<b>284.3</b>

Table B-1: As Found IT Power Summation at PDU Outputs

PDU and 12kV Transformer Data Calculated		
Description	value	unit
IT Power	284.3	kW
PDU Efficiency	98.0	%
PDU Input Power	290.1	kW
PDU Loss	5.8	kW
12kV Output Power	290.1	kW
12kV Transformer Efficiency	99.0	%
12kV Input Power	293.0	kW
12kV Transformer Loss	2.93	kW

Table B-2: As Found PDU and 12kV Power Levels and Loss

## Appendix C: As-Found Analysis Diagram





## Appendix E: pPUE Calculation Details

	<i>red italics</i> = measured	<i>blue Cambria</i> = assumed/estimated from other data		
	<i>green</i> = calculated			
		As Found	Whole DC	
Item#		Value	Value	UOM
1	IT Power	<i>284.27</i>	<i>284.27</i>	kW
2	PDU Efficiency	<i>99</i>	<i>99</i>	%
3	Power in to IT PDU	<i>287.14</i>	<i>287.14</i>	kW
4	Power Loss in IT PDU's and UPS	<i>2.871</i>	<i>2.871</i>	kW
5	12kV Transformer Eff.	<i>98</i>	<i>98</i>	%
6	Power in 12kV Transformer	<i>293.00</i>	<i>293.00</i>	kW
7	Loss in 12kV Transformer	<i>5.86</i>	<i>5.86</i>	kW
<b>Measured CDU Heat Transfer</b>				
8	CDU #1 Heat Transfer	<i>121.0</i>	<i>142.14</i>	kW
9	CDU #2 Heat Transfer	<i>99.0</i>	<i>142.14</i>	kW
10	Total CDU Heat Transfer	<i>220.0</i>	<i>284.3</i>	kW
11	CDU #1 Heat Transfer	<i>34.42</i>	<i>40.43</i>	tons
12	CDU #2 Heat Transfer	<i>28.16</i>	<i>40.43</i>	tons
13	IT Power Cooled by CDU's	<i>0.774</i>	<i>1.000</i>	ratio
14	IT Power Cooled by Rm. CRAH's	<i>64.25</i>	<i>0.000</i>	kW
<b>Measured CDU Pump Power</b>				
15	CDU#1 Pump Power	<i>1.2</i>	<i>1.49</i>	kW
16	CDU#2 Pump Power	<i>0.9</i>	<i>1.49</i>	kW
17	Misc. Power Lights/etc.	<i>6</i>	<i>6</i>	kW
19	Total Heat Cooled by CRAH's	<i>75.22</i>	<i>11.86</i>	kW
20	CRAH Fan Power	<i>4.613</i>	<i>0.727</i>	kW
<b>Calculate Total Heat Cooled and Power Used by Chiller Plant</b>				
21	Total Elec. Power to D.C. Room	<i>301.10</i>	<i>301.99</i>	kW
22	Chiller Plant Cooling other than CDU	<i>75.22</i>	<i>11.86</i>	kW
23	Item 22 Converted to tons	<i>21.39</i>	<i>3.37</i>	tons
25	Chiller Plant Efficiency w/pumping	<i>0.66</i>	<i>0.438</i>	kW/ton
26	Chiller Plant Power with Pumping	<i>14.12</i>	<i>1.48</i>	kW
27	Chiller Plant CDU Cooling	<i>220.0</i>	<i>284.3</i>	kW
28	Item 27 Converted to tons	<i>62.58</i>	<i>80.85</i>	tons
29	Chiller Plant Efficiency wo/pumping	<i>0.64</i>	<i>0.392</i>	kW/ton
30	Chiller Plant Power without Pumping for CDU's	<i>40.05</i>	<i>31.69</i>	kW
31	Total Chiller Plant Elec. Power Use	<i>54.17</i>	<i>33.17</i>	kW
<b>Calculate pPUE</b>				
32	Total pPUE Elect. Pwr. Required	<i>355.27</i>	<i>335.16</i>	kW
33	IT Power	<i>284.27</i>	<i>284.27</i>	kW
34	<b>pPUE</b>	<i>1.250</i>	<i>1.179</i>	ratio

## Appendix F: Calculation Detail Notes

Note #	As Found	Generic
1	from PDU measurements, see Appendices A and B	assumed same as "As Found"
2	PDU losses based on input and output measurements	assumed same as "As Found"
3	calculation	calculation
4	calculation – assumes UPS loss of ~10% or 240 watts is included	calculation
5	direct measurements not available, assumed consistent with Greenberg and Coles 2011 and Synapsense analysis	assumed same as "As Found"
6	calculation	calculation
7	calculation	calculation
8	measured, then corrected with temp. calibration instrument, water flow rate was not verified and could be a source of error	assume 50% of total IT power removed
9	measured, then corrected with temp. calibration instrument, water flow rate was not verified and could be a source of error	assume 50% of total IT power removed
10	calculation	calculation
11	units conversion	units conversion
12	units conversion	units conversion
13	calculation	assumed 100%
14	calculation	calculation
15	measured at feed panel	ratio calculation using "As Found" measurements = 0.0369 kW/ton
16	measured at feed panel	ratio calculation using "As Found" measurements = 0.0369 kW/ton
17	(40kW+40kW)*7% Magellan = ~5.6kW round up to 6kW; 40kW from Synapsense	assumed same as "As Found"
18	calculation	calculation
19	calculation: 32tons cooling = 6.9kW fan power per Liebert specs.	calculation: 32tons cooling = 6.9kW fan power per Liebert specs.

20	calculation	calculation
21	calculation includes 100% CDU motor power to ambient air, this conservative assumption raises the PUE a small amount	calculation includes 100% CDU motor power to ambient air, this conservative assumption raises the PUE a small amount
22	units conversion	units conversion
23	calculation	calculation
24	calculation	calculation
25	Explanation below*	Taylor Engineering model using 53.5°F cooling water <u>without</u> pumping see Appendix E
26	calculation	calculation
27	calculation	calculation
28	calculation	calculation
29	Explanation below*	Taylor Engineering model using 53.5°F cooling water <u>without</u> pumping see Appendix E
30	calculation	calculation
31	data copied from item#1 above	data copied from item#1 above
32	calculation	calculation

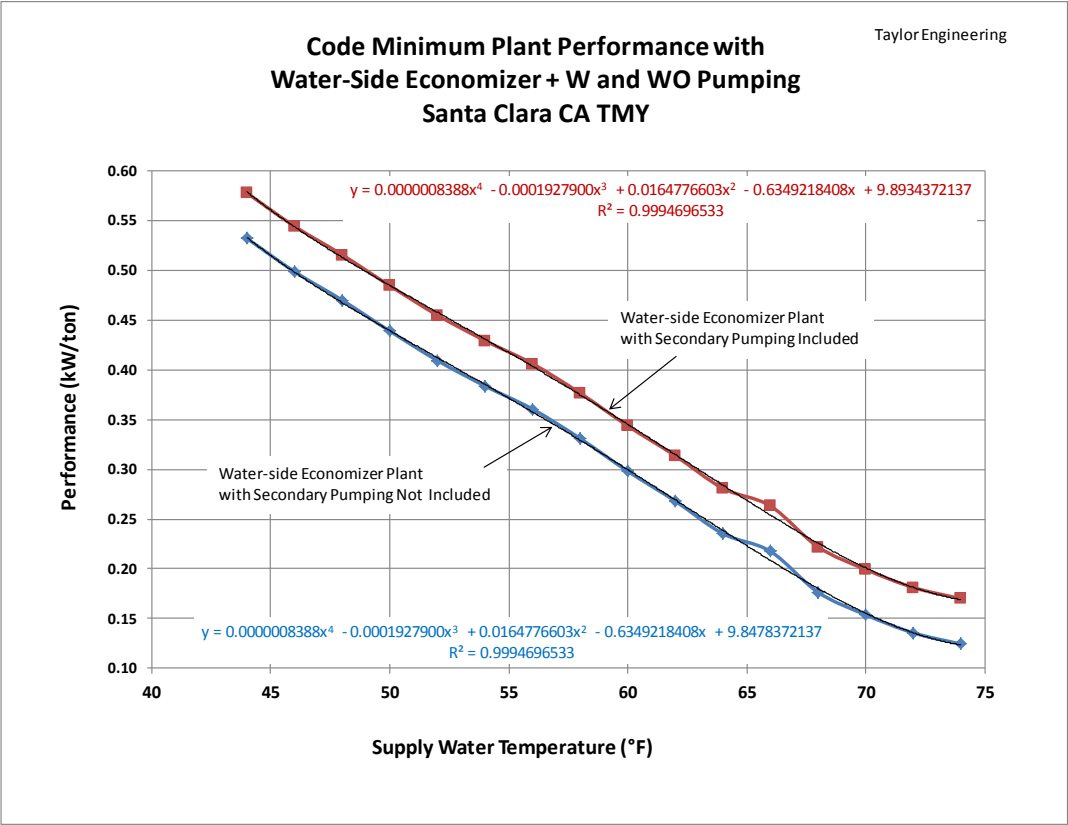
\* The cooling for Magellan is provided by the basement chiller plant, which rejects its heat from cooling towers on the main building roof. In order to determine the electrical power required at the plant to provide cooling water to Magellan, the plant kW/ton needed to be established with two variants: one with all plant components included (for chilled water delivered to the CRAHs providing auxiliary cooling), and the other with the secondary chilled water pumps excluded (for the CDUs, whose pumps remove water from the chilled water return piping and return it back to the return, with no net pumping requirement from the plant's secondary pumps.

These two variants were estimated based on metered data and equipment ratings. All significant loads for this plant are equipped with electrical power meters, and the small remaining loads are estimated for completeness. The chiller gross cooling is estimated based on chiller input power and chiller rated kW/ton, since no reliable flow meter was available, and the piping configuration makes proper flow measurement impractical.

The estimated kW/ton for the period of the Magellan power readings is 0.66 with all equipment included, and 0.64 with the secondary chilled water pumps excluded.

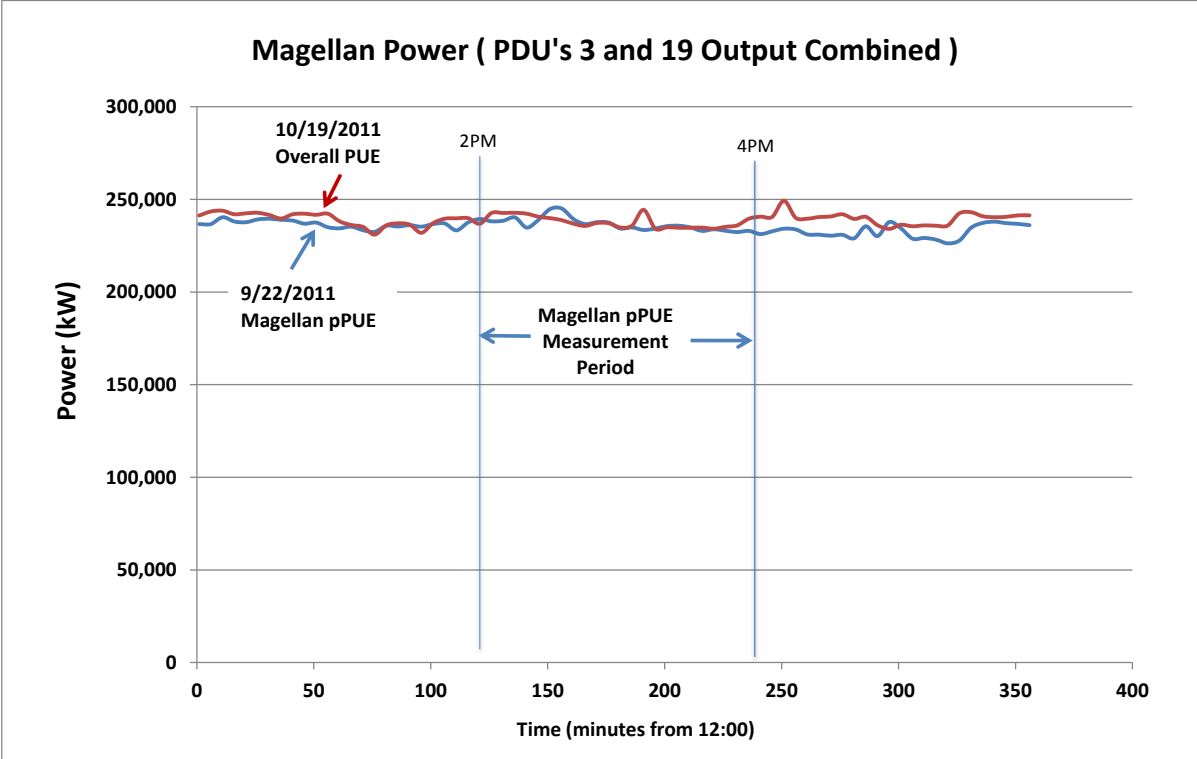
The plant equipment includes: 3 chillers, 3 each primary chilled water, secondary chilled water, and tower water pumps, 6 cooling tower fans, one tower water separator pump, and one chiller room exhaust fan.

# Appendix G: Taylor Engineering Water-side Economizer Chilled Water Plant Model With and Without Secondary Pumping





# Appendix H: Power Level Comparison



## Appendix I: Reference Only – Simplified PUE

Actual PUE approximately 0.04 above trend shown

