

Measuring and Managing a Data Center Cooling Shutdown

Columbia University Information Technology - August 24, 2012

Table of Contents

[Overview](#)

[Timeline](#)

[Cricket Graphs](#)

[What's being graphed](#)

[Mechanical](#)

[Electrical](#)

[Ambient Temperature](#)

[24 hours as of 9am 8/17/12](#)

[24 hours as of 8:46pm 8/18/12](#)

[36 hours as of 10:00 8/19/12](#)

[Final tenant system - EE cluster - restored 8/20/12](#)

Overview

A full mechanical shutdown of the cooling plant for the University Data Center was performed Saturday August 18, 2012. This work was required by construction for the Core Research Computing Facility (NIH grant 1G20RR030893-01) which is upgrading the Data Center's electrical service. The work consisted of demolition of unused cooling pipe and electrical work needed to prepare for further construction activities. The chilled water cooling system and electrical power to the computer room air conditioners (CRAC) were shut down. The shutdown was also used as an opportunity to perform extensive preventative maintenance on the cooling and electrical systems in this 50-year old facility.

As the Data Center is a critical facility that can't afford to be shut down for a full day, a 50-ton temporary cooling system was employed, providing enough capacity to operate at a reduced load by shutting down non-critical services for the day. The normal operating load is about 120 tons, so approximately $\frac{2}{3}$ of the approximately 700 servers in the data center were planned to be shut down.

Thanks to support from the New York State Energy Research and Development Authority (NYSERDA contract ST11145-1) extensive instrumentation of the cooling and electrical distribution systems had been recently installed. This monitoring capability helped us plan for and manage the shutdown. Following are representative graphs of the installed monitoring capability that helped us successfully plan for and manage this work.

Timeline

Planned and actual times for the work were developed and tracked for this event. Following is a summary of major milestones and incidents of note:

Milestone/Incident	Planned Time	Actual Time
First 3 hour temporary cooling test completed.	Thursday 8/16/2012 17:00	Thursday 8/16/2012 18:45
Second 3 hour temporary cooling test completed. 40 ton reduction in central cooling load measured.	Friday 8/17/2012 17:00	Friday 8/17/2012 16:45
Commence server shutdowns (confirm overnight batch complete)	Saturday 8/18/2012 06:00	Saturday 8/18/2012 06:00
Electrical monitoring data collector failed for 30 minutes (7:30-8)		Saturday 8/18/2012 07:30
Commence electrical and mechanical work	Saturday 8/18/2012 08:00	Saturday 8/18/2012 08:18
Cricket data collected failed for 1 hour (11-12)		Saturday 8/18/2012 11:00
Mechanical and electrical construction work complete. CRAC PMs ongoing. Some CRACs and air handlers (AC1&2) operating.	Saturday 8/18/2012 19:00	Saturday 8/18/2012 11:36
Chilled water operating through backup dry coolers (12-2:30)		Saturday 8/18/2012 12:00
Chilled water switched over to normal campus CHW operation		Saturday 8/18/2012 14:30
Recovery of IT systems commences	Saturday 8/18/2012 20:00	Saturday 8/18/2012 15:00
Majority of central IT systems recovered -- some tenants elected to wait until later to restart their systems.	Saturday 8/18/2012 23:59	Saturday 8/18/2012 21:26
Last major tenant IT system -- 800-core HPC cluster -- restored starting 2pm. Full load reached at 5:30pm.		Monday 8/20/2012 17:30

Cricket Graphs

What's being graphed

The following mechanical and electrical infrastructure elements are instrumented using the SNMP monitoring protocol (with translation from ModBus).

Mechanical

Mechanical systems monitoring consists of Flexim temperature and flow rate metering installed on three subsystems which measures temperature, flow in gallons per minute, and calculates heat transferred in tons.

ac1-2 chilled water loop serving the overhead air handlers (ac1 and ac2).

liebert1-6 chilled water through a plate-frame heat exchanger that serves the computer room air conditioners (CRAC)

drycoolers water loop is the backup system, switched to via butterfly valve to reroute the "liebert1-6" water through rooftop dry coolers. Notably, the dry coolers operate at ambient temperature and typically see water temperatures of over 90 degrees as the CRAC compress transfer heat to the water loop. This is to be contrasted with the normal operation chilled water loop with temperatures in the 42-50 degree range.

Electrical

Electrical systems monitoring is installed at several levels throughout the facility: Wattnode meters monitor the main 3-phase distribution panels (and several subpanels) for mechanical and IT loads; Raritan instrumented power distribution units are mounted in key server racks. Further monitoring is available via several of the rack-mounted UPSes as well as several of the servers.

ats2 & 3 are the automatic transfer switches feeding MCC1 and MCC2, respectively. These MCCs supply power to the CRACs, AC1&2 air handlers, chilled water pumps, and various lighting and office electrical loads throughout the Computer Center building.

ats4 is the automatic transfer switch feeding the majority of the IT load in the data center.

pp26 & 27 are two power panels that carry IT load but are not fed from the main IT panel, ats4. These were installed within the past 5 years and are connected to a totally different source from the rest of the Computer Center building.

unixrackN-pduN from Raritan are mounted in several equipment racks as mentioned above.

Ambient Temperature



Ambient temperature monitoring is available from several systems:

all sensors graph shows ambient temperature data from several zones throughout the room using Neteye monitors which provide video as well as environmental conditions monitoring.


unixrackN-pduN from Raritan are also set up with ambient input air temperature probes that are mounted on the front grilles of the equipment racks.

XIV enterprise storage system has input temperature monitoring for each of 15 storage shelf modules.

24 hours as of 9am 8/17/12




Liebert1-6 heat-flow, interval is 1 day





Last updated Fri Aug 17 09:01:39 2012


ac1-2 temp, interval is 1 day



Last updated Fri Aug 17 09:01:38 2012




ac1-2 volume-flow, interval is 1 day



volume-flow 138



Last updated Fri Aug 17 09:01:38 2012



drycoolers volume-flow, interval is 1 day





volume-flow 3



Last updated Fri Aug 17 09:01:39 2012






24 hours as of 8:46pm 8/18/12









pp27 PowerSum, interval is 1 day







Last updated Sat Aug 18 20:46:57 2012



36 hours as of 10:00 8/19/12

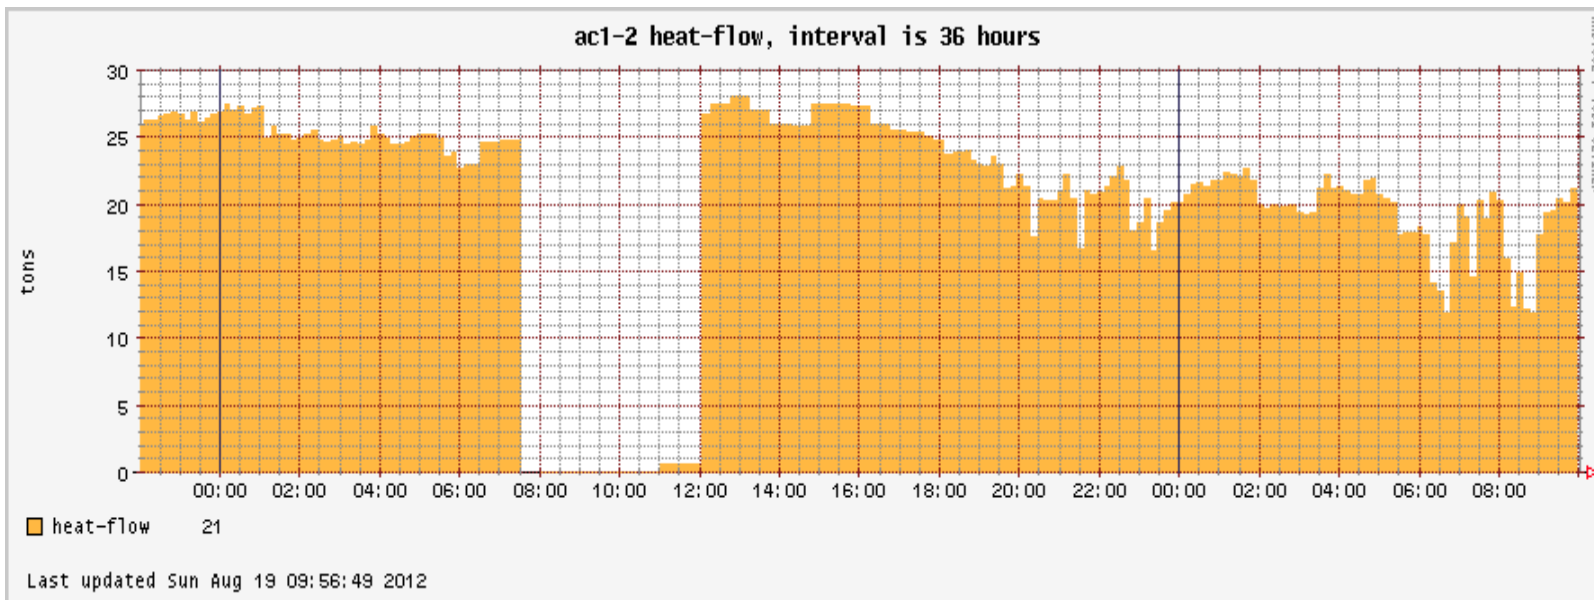



















all sensors temperature, interval is 36 hours





Last updated Sun Aug 19 09:56:49 2012


all sensors humidity, interval is 36 hours



Last updated Sun Aug 19 09:56:49 2012




unixrack34-pdu1 Temperature, interval is 36 hours



Temperature 24.50

Last updated Sun Aug 19 10:57:48 2012


unixrack34-pdu2 Temperature, interval is 36 hours



Temperature 22.50

Last updated Sun Aug 19 10:57:48 2012

unixrack4-pdu1 Temperature, interval is 36 hours



Last updated Sun Aug 19 10:57:48 2012

Final tenant system - EE cluster - restored 8/20/12

