

Procedures and Policies to Respond to A/C Failures

Approved: February 27, 2008

1 Purpose

The purpose of this document is to describe how the clusters in room 74-2060 shall be configured in order to prevent equipment damage in the case of an A/C failure and to establish guidelines for responding to A/C failures during off-hours.

2 Overview

Computer equipment (CPUs, ram chips, hard drives, north and south bridges, etc.) are particularly susceptible to damage from excess heat, which can significantly reduce the lifetime of these critical components. Room 74-2060 is shared by multiple clusters and servers, and excess heat from any one cluster can damage both that cluster and the other equipment present in the room. In order to mitigate the effects of excess heat, we need to ensure that the room temperature remains below 76F (more specifically the air temperature entering the clusters). All clusters shall monitor the room temperature and idle jobs / kill jobs if it exceeds 75F. Clusters shall completely shutdown if the room temperature exceeds 90F. Software to monitor the room temperature will be provided, and the cluster shall be set up to automatically idle / kill jobs and shutdown at the appropriate temperatures. Cluster administrators may choose to set these thresholds lower.

3 Guidelines for cluster configuration

Clusters shall monitor the room temperature via the “bluebox” (Uptime Devices, bluebox, SH2, v4.3) and automatically respond as follows:

- If $T > 75F$ then idle, or optionally kill, all jobs
- If $T > 89F$ then completely power down
- Cluster administrators may choose to set these threshold temperatures lower, but not higher.

In addition, it is recommended that each cluster monitor the individual node temperatures and shutdown nodes as needed.

A computer outside of 74-2060 shall also receive temperature information from the bluebox and send alerts via e-mail and text messages if the bluebox had stopped sending current room temperatures.

4 Guidelines for handling A/C failure

4.1 Sensaphone alerts and First Responder Actions

In the event of an A/C failure the room temperature will rise and trip an alarm on the Sensaphone. The temperature threshold for this alarm will be set sufficiently low to catch the A/C failures early, but high enough to avoid false alarms. The Sensaphone will call a list of “First Responders” (FRs). The Sensaphone will cycle through the list calling each number one-by-one until one of the FRs indicates that he/she will attend to the emergency. A FR indicates this by entering the code 555. After entering this code the FR shall do the following:

- Immediately call campus safety 475-2853 (475-3333 is the emergency number) and request that the on-call A/C technician attend to the A/C failure
- Go to RIT to monitor the temperature in the cluster room and perform a manual shutdown of the clusters if it exceeds 90F

4.2 Shutdown procedure

The FR **must** shutdown the clusters if the room temperature exceeds 90F. The shutdown procedure is as follows:

- Log onto the clusters as user “heat” with a password provided to all FRs
- Wait 8 minutes and then turn off any remaining nodes, switches, etc, via the power buttons on the individual pieces of equipment, and then turn off power to all UPSs

The clusters shall be set up such that a login as user “heat” will trigger a shutdown, which can take up to 5 minutes. If the shutdown fails (likely because the cluster should have automatically shutdown in response to the temperature alert sent by the bluebox), or if the head node has crashed, or the cluster does not allow logins, then the cluster shall be shutdown by turning off power to each piece of equipment (via the power buttons, if present) and then turning off the UPSs.

4.3 Extreme Emergency

In the event that the room temperature exceeds 100F, the FR shall immediately shutdown all cluster, servers, etc. by shutting off the UPSs.

4.4 Post-Shutdown procedure

The FR can leave RIT once the clusters have been shut off.

4.5 Prolonged A/C failure and replacement FRs

If the temperature does not exceed 90F, but the A/C remains non-functional, then the FR shall monitor the room temperature and A/C status for 30 minutes. If after 30 minutes the A/C has not been fixed, then the FR shall attempt to contact another FR. Once another FR has agreed to go to RIT, the original FR shall monitor the room temperature for an additional 15 minutes. If after 15 minutes the room temperature has not reached the critical 90F, the original FR can leave. If no replacement FR is available, or if the room temperature reaches 90F at any time, then the

original FR shall shutdown the clusters via the procedure in subsection 4.2. Once the clusters are shutdown, the FR(s) can leave.

The A/C shall be deemed functional once cold air begins circulating. If A/C has been repaired, then the FR shall monitor the room temperature and can leave when the temperature falls below 80F.

4.6 FR report

The FR(s) shall write a brief e-mail describing the situation and any actions that he/she/they took.

5 Testing the Automated Shutdown System

Under full load, the room temperature in 74-2060 rises quickly. It is therefore vital that the automated idle / shutdown system works correctly. The cluster administrators shall arrange a time to test the automated idle / shutdown procedure and the manual “heat” login shutdown procedure. This test will consist of:

- Raising the temperature reported by the bluebox to 80F and confirming that both clusters idle / kill jobs.
- Raising the temperature reported by the bluebox to 90F and confirming that both cluster shutdown successfully within 8 minutes.
- Loading both cluster with jobs, logging in as “heat”, and confirming that both clusters shutdown successfully within 8 minutes.

In addition, the Sensaphone shall be tested weekly. The test will consist of raising the Sensaphone temperature to 80F and confirming that all numbers are called.

The temperature of the Sensaphone / bluebox will be safely raised to 80F / 90F for these tests using a dry heat source (with temperature $T < 100F$) applied directly to the thermocouple / thermistor.