

Matt Crawford

Amitoj Singh

Adam Walters

Jason Allen

Jun 15, 2011; revised Jul 7, 2011; Aug 11, 2011

ACTION PLAN FOR GCC CRB COOLING

This report explains the observations and reasoning behind the GCC load-shed action plans. The associated ITIL Problem Investigation is PBI00000000256. Load shedding is managed through the critical incident process.

The nature of the problem

GCC Computer Room B's total power is, and has always been, significantly below its designed maximum, and below the rating of the CRACs (Computer Room Air Conditioners). When at least 9 out of the 10 CRACs are operating, they should be able to remove 945 kW of heat. The room is intended to operate up to 900 kW of computing, and has been at 700 kW in recent months. It was around 600 kW in the summer of 2010.

Without short-cutting the root cause analysis, we can say that the difficulties appear to center on the condensers and the circumstances they operate under. In the worst weather conditions, a condenser may not be getting enough air to fully condense the coolant sent by the CRAC, or the air may be too hot. There are notes and photos attached to the PBI in Remedy to explain this more fully, and CD DocDB Document 4359 gives a background on our experience with cooling CRB.

Once the condensers start returning a liquid/gas mixture, effectiveness of the cooling system falls, the compressors in the CRACs run harder, but the room temperature rises only slightly until one or more CRACs fail outright, or give the "HIGH HEAD" alarm (signaling high pressure at the coolant return) and their compressors stop. Then the room temperature rises rapidly and CRAC failures accelerate.

Shedding heat load

The first critical incident involving GCC cooling in 2011 was from June 6–9. Soon afterward, several service and stakeholder representatives prepared the first draft of a staged power reduction plan to be used if needed to prevent another cooling failure if technical experts considered such an event imminent. This plan was revised with input from more stakeholders in early July. The first stage involves reducing power by about 30%, shutting down mostly older-vintage systems in racks 3000-3027. During the planning process we had considered a possible second stage of power reduction by another 20 to 30%. However, we strongly doubted that a second stage would prevent a cooling failure if the first reduction did not. But proving that "No plan survives first contact with the

enemy,” some intermediate shutdowns were performed during the next invocation of the load shed plan. This was done in order to preserve as much computing capacity as possible for preparation of results for summer conference presentations. Those intermediate shutdowns were denoted as “Stage 1.5a” and “Stage 1.5b” during that incident but will be referred to below as stages 2 and 3.

The final stage shedding heat load is, of course, shutting off all computing equipment.

Racks to power off in stage one, listed by the half-row.

3003	3007–3012
3014–3020	3021–3027
(none)	(none)
3054–3055	3042
(none)	(none)
(none)	3075

Racks to power off in stage two, listed by the half-row.

Incremental shutdowns beyond stage one are underlined.

<u>3000–3002, 3003, 3004–3006</u>	3007–3012, <u>3013</u>
3014–3020	3021–3027
(none)	(none)
3054–3055	3042
(none)	(none)
(none)	3075

Racks to power off in stage three, listed by the half-row.

Incremental shutdowns beyond stage *one* are underlined.

<u>3000–3002</u> , 3003, <u>3004–3006</u>	3007–3012, <u>3013</u>
3014–3020	3021–3027
<u>3035–3037</u>	<u>3031–3033</u>
3054–3055	3042, <u>3043</u> , <u>3045–3047</u>
(none)	(none)
<u>3077–3078</u>	3075

Affected stakeholders

The racks listed above contain computing equipment used by the following primary stakeholders. It should be borne in mind that a count of racks shut down is not a complete measure of the impact on stakeholders. Racks contain computers of different power, and stakeholders have different amounts of computing power in other locations. This table shows the racks to be shut down in the first stage and the total racks in GCC CRB used by each stakeholder.

stage	CDF	CMS	D0	FermiCloud	GPCF	GPFarm	GridServices
1	5 / 17	10 / 36	6 / 15	1 / 2	0 / 2	3 / 10	0 / 1
2	12 / 17	10 / 36	6 / 15	1 / 2	0 / 2	3 / 10	0 / 1
3	15½ / 17	18 / 36	6 / 15	1 / 2	0 / 2	3½ / 10	0 / 1

Invoking the load-shed plan

Comparison of last summer's cooling incidents to Fermilab weather records suggests that the following three factors combine to cause cooling failures:

- Outside temperature above 88°F, especially for consecutive days,
- High solar radiation to the ground,
- Low wind.

We believe that waiting until the hottest part of the day to reduce the heat load is risky. Once multiple CRAC compressors are running at 75 to 100% of capacity, the room may already be in imminent danger of cooling failure. Reducing the load mid-morning when adverse conditions are expected is safer, but of course has a greater impact on computing throughput and a risk of a false preventive action. Invoking the plan in the morning allows

orderly communication through the Service desk. System managers can then shut down their designated machines remotely. If load reduction happens in the afternoon, implementation may have to be rushed, even to the point of turning off power at the electrical panels.

Facilities will notify the Service Manager as early as possible when the former judges that conditions may require shedding heat load in GCC. Stakeholders and service providers will be notified through the Service Desk. If the load shed is needed, service providers will perform the shutdowns and a critical incident team will be assembled. Restoration of service will be decided by the critical incident team.