

TRIUMF incident report

compiled by R. Tafirout
April 25 2012

This report describes the time line and events that occurred on April 10 when TRIUMF suffered two major power failures within a few hours, each lasting a few hours and causing a significant downtime (~20 hours) affecting all Tier-1 services.

At present, the Tier-1 centre relies on utility power to be highly reliable and available effectively 24/7/365. Historically this has always been the case at the TRIUMF laboratory. In light of these events, we are revisiting our strategy for the longer term and are exploring a few options. TRIUMF core network services are backed up with diesel power that is also used for TRIUMF's core operations. The current diesel infrastructure is of limited capacity and so the Tier-1 data centre cannot tap into it at present. We would need a separate feed and a much larger capacity.

The root cause of the power failures is not fully understood. It was a ground fault causing a relay trip. TRIUMF load was put onto another transformer to try to understand the cause, but this caused also another ground fault few hours later. The load was reverted back to the initial transformer. Once the power has been restored, we found no serious damage that impacted production due to built-in redundancies.

The time line of events is the following: (time stamps are all in UTC and some are approximate)

April 10, 07:42 - utility power failure #1

This triggers an automatic shutdown sequence of large loads first (worker nodes), awaiting for personnel to arrive on site and assess the situation and complete an orderly shutdown of the remaining equipment. The UPS battery capacity eventually ran out and some of the remaining equipment was shut off uncleanly.

April 10, 13:10 - utility power restored

Because the UPS batteries have been fully discharged, we had to issue a service call (we do have 4 hours on-site support 24/7/365). The UPS power was ready only at around 18:15 . At that time we started powering on the data centre.

April 10, 20:25 – utility power failure #2

We were not back into production yet as the second power failure occurred. The shutdown procedure was initiated again and all the equipment powered off cleanly as personnel already on-site.

April 11, 00:15 – utility power restored

After assessing stability of power we started bringing back online systems at around 01:00 and Tier-1 services restored shortly after. After the usual functional tests ATLAS resumed its production in the following hours.