

# Service Incident Report for the Cooling problem at PIC on 21 June 2012

Incident start: 21-Jun-2012, 13:30 UTC

Incident end: 21-Jun-2012, 14:30 UTC

## Description

Due to a loss of cooling system performance the temperature of the datacenter room raised and the automatic control that acts to lower the IT load in case of temperature problems stopped part of the PIC Worker Nodes (WN).

## Impact

160 WNs were shutdown on June 21st due to this incident. After the cooling system recovered, 64 WNs could be powered on again so by the time of writing this report 96 WNs remain halted, to avoid repeating the overheat incident. The plan is to keep these 96 WNs stopped until a more detailed analysis of the cooling system performance depending on outdoor temperatures is performed together with the equipment vendor.

The total computing capacity at PIC which was of 39738 HS06 before the incident, has therefore been reduced of about 17% because of this incident, down to 33130 HS06.

The Tier1 computing service continues running with no problem, but with reduced capacity. It is worth to note that 33130 HS06 is already well above PIC pledged capacity for 2012 (26367 HS06).

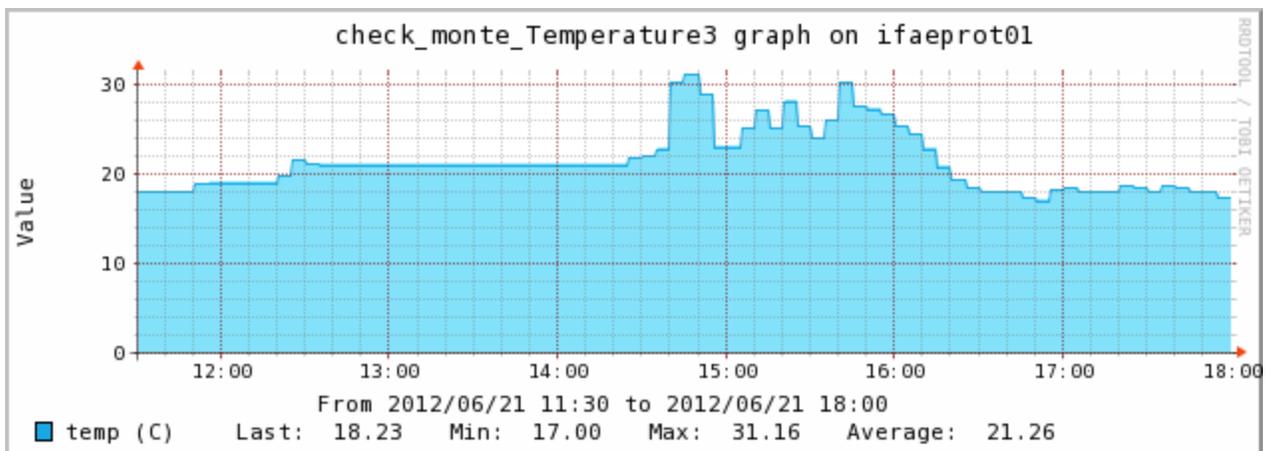
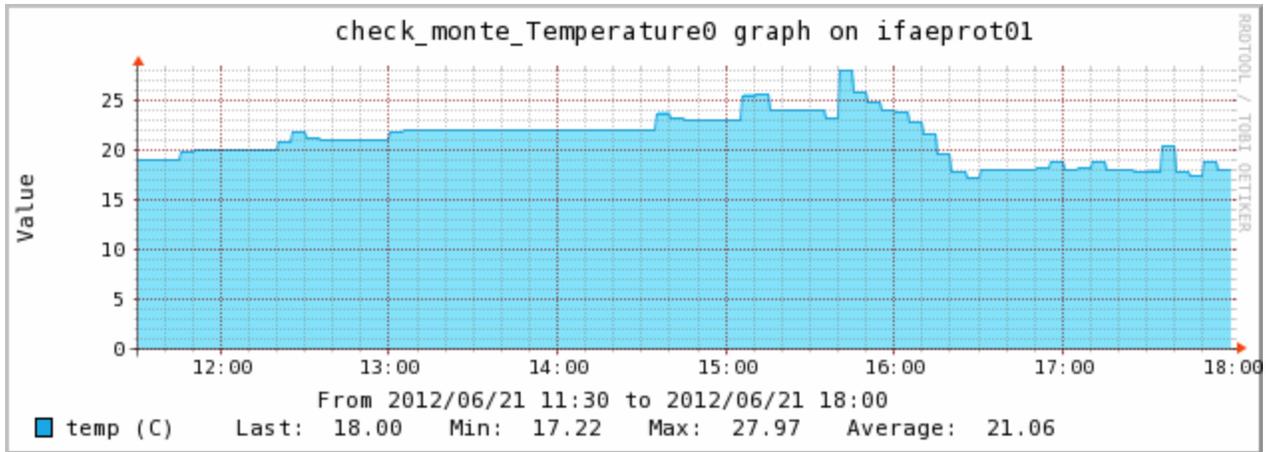
## Time line of the incident

(all times in UTC)

- 21-June
  - 08:00 - The temperature in the datacenter began to rise slightly on Thursday morning
  - Infrastructure team checks if there are any errors in the equipment but did not detect any hardware problem
  - 12:00 - Decide to reduce some IT power. An internal ticket is opened to the computing administrator to gracefully drain and shutdown some WNs
  - The temperature kept rising slowly until the cooling system lost performance about 13:30 UTC. Then the automatic temperature control detected that it had rose above the safety threshold, so 112 WNs were stopped automatically in order to prevent a major incident.
  - The IT power was reduced 40kW and the cooling system returned to be within the working thresholds.
  - 14:00 - Once the cooling system gained performance, decided to power on 48 WNs.

- 25-June, 10:00: 16 additional WN were powered on.

The two graphs below show the temperature value time evolution at two different locations inside the datacenter. The x-axis time values are in local time (CEST) units.



## Analysis

The outside temperature was exceptionally high on Thursday June 21st. There is no available sensor reporting this value on the specific equipment that caused the problems, but a similar equipment located on the same roof reported that the outside temperature was 40C. Although the max. allowed temperature is 46C according to the specifications, the performance of the system degraded such that the right temperature inside the IT room could not be kept.

## Followup actions

1. The annual maintenance check for this cooling equipment has been scheduled for July 4th. The possibility to monitor the temperature on the air condenser will be evaluated.
2. Collect more data about the performance degradation behaviour of the cooling system with rising outside temperatures. Until this is done, and to avoid unnecessary risks, 96 WNs (~17% of total capacity) will be kept powered off during the hottest part of the summer.

3. Study different options for gathering sensor data about the temperature and operation mode of the air condenser unit to be able to foresee similar problems with more time.