Data centers are a software development challenge

Luca Bortot
ENI IT Architectures Project Leader
luca.bortot@eni.com

Walter Nardelli
ENI Technical Services Manager
walter.nardelli@eni.com

Peter Seto
EE HPC WG
jordanruthseto@aol.com

Case study of a data center designed by computer scientists, with a focus on data-driven improvement, after 6 years of activity

Context, *aka* the tools

Green Data Center (GDC) plate data

- Up to 30MW IT power
- 5000 sqm in 6 data rooms
- TIER-IV compliant
- **Up to 50KW/rack**
- **2 HPC systems (22.4 Pflops peak)**
- > 10000 servers
- PUE < 1.2 (southern clover 2018: 1.17)

Datacenter automation

- **Designed and implemented in-house from scratch**
- Evolved and extended with the increased understanding of the GDC functioning
- Coordinates the subsystems, applying an energy strategy that depends on current, past and forecasted environment conditions

Monitoring system

- Designed to **be the facility "debugging" tool**
- Collects any available data, even those that seem unnecessary (may prove useful in the future)
- Was up & running when the datacenter was being built
- > 200k samples every 10s
- > 8 years worth of data online

Improving availability

Fault detection

- A fault-tolerant automation system may mask some faults or sub-optimal behaviors
- Using detailed ODA data, masked non-trivial faults can be detected and anticipated

Stress-testing

- Maintenance errors or wear can reduce a device effectiveness
- In a redundant facility this may be masked because no device is used at its limit, except in emergency
- The automation performs routine on-line stress tests of each component

Increasing efficiency

Dynamic setpoint adjustment

- Running a system at a fixed-point, worst-case scenario means wasting energy
- Most devices can exceed their data-plate values
- Adapt every subsystem setpoint to the current load and environment

Dynamic hardware selection

- Different component configurations can provide a given function, eg cooling vs humidity control
- Let the automation decide the best configuration and strategy to save energy and reduce hardware wear

Conclusions

Detecting masked faults

A masked defective turbine. Setpoint is attained, but energy spike is detected

Increased efficiency

- Without hardware changes, the PUE has dropped from 1.4 (2013) to 1.17 (2018)

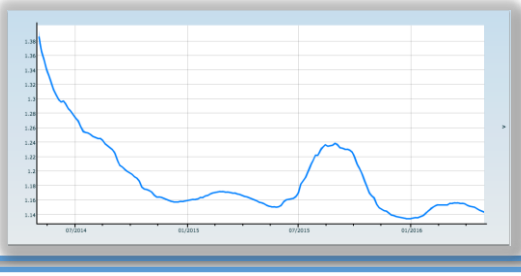

Knowledge base

- The documented work done at the GDC can be applied to future facility and software design