

A model for forecasting data centre infrastructure costs

System Performance and Cost Modeling Working Group

- Large computing requests to come (HL-LHC, astroparticles ...)
- Need to provide IT service budget
 - How much does the infrastructure cost today ?
 - How much is it going to cost in the future ?

What changes to make IT infrastructure able to address the needs in storage and data processing at HL-LHC ?

- WLCG Working Group created ~ 1 year ago
 - Studies workflow performance
 - Estimates future resources needs
 - Studies infrastructure costs

- No (or little) information available
 - Site expenses
 - Budgets modeling
- CC-IN2P3 carried on a study in 2014-2015
- Presented at CHEP 2015 (Okinawa)
 - Talk : <https://indico.cern.ch/event/304944/contributions/1672219/>
 - Proceedings : https://www.researchgate.net/publication/288022117_A_model_to_forecast_data_centre_infrastructure_costs


Cost model objectives

- Infrastructure costs
 - Hardware costs per sector
 - CPU and disk servers with rack, switch... → 'equipped' cost
 - Power costs per sector
 - IT & cooling
- Trends
 - How prices have been evolving so far
- Try to model yearly resource investment, per sector
 - Today
 - Next years


Hardware

Cost evolution & flat budget

- Costs drop **exponentially**
 - xxx % per year
 - Observed in storage, compute, tape cartridges...
 - « density » evolution
- Situation in many sites : **flat budget**
 - Every year, same amount of money to invest = purchase of servers, routers, tapes, drives, licences
 - 1st part : capacity replacement
 - 2nd part : capacity increase
- Price drops
 - Capacity replacement budget < Total budget
 - Available capacity increases with time

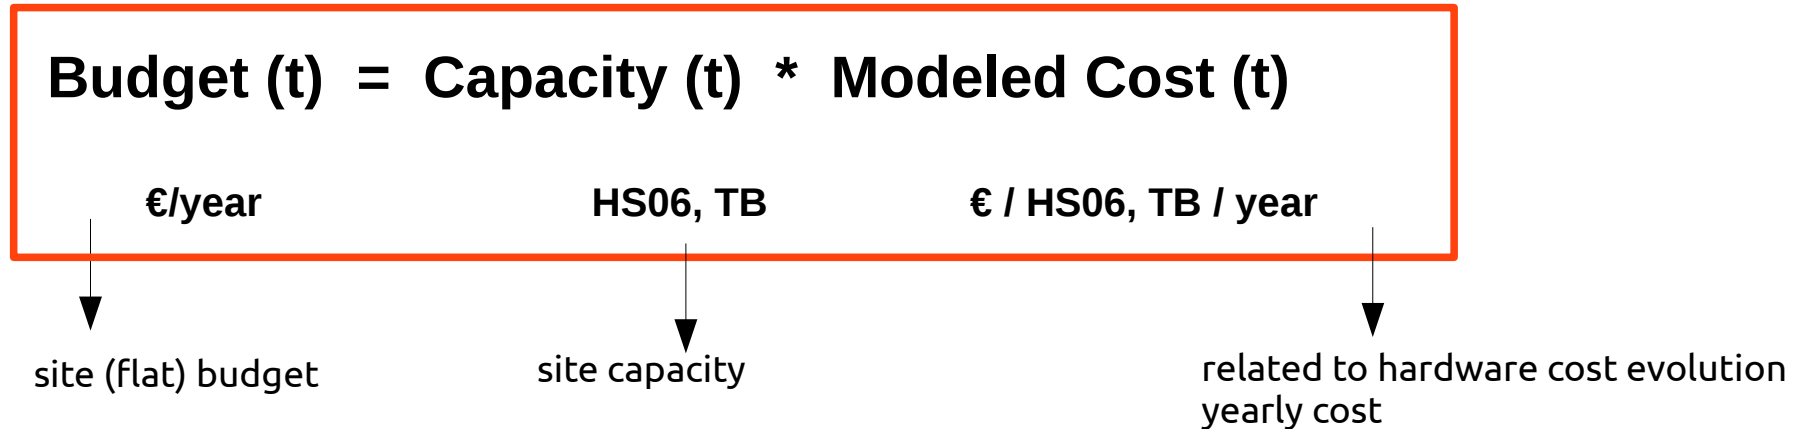


Capacity evolution at flat budget



What we can do with that

- Ideal hypotheses :
 - Flat budget
 - Exponentially-decreasing costs
 - Use only hardware under warranty
- → everything becomes exponential 😊
- Then site capacity relates to budget in an ~easy way




Modeled cost ???

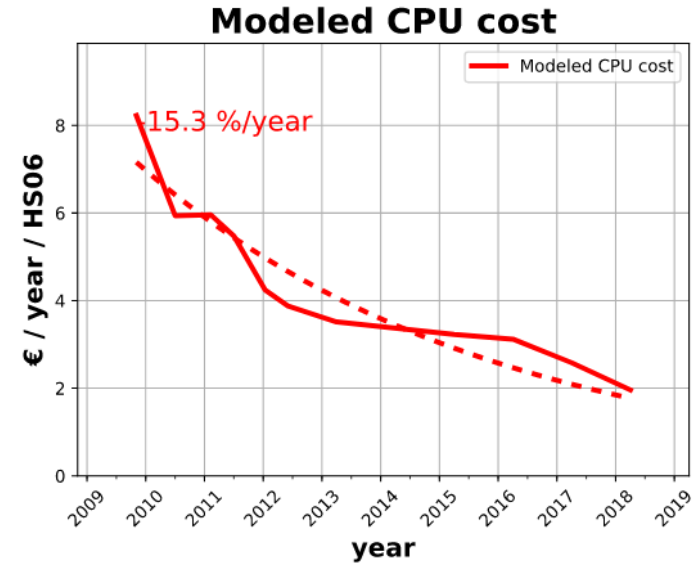

- Needs previous hypotheses

$$c^*(t) = c(t) \times \frac{r}{1 - (1 - r)^\tau}$$

- c^* : the modeled cost
 c : the procurement cost (= the real cost)
 τ : hardware warranty time (years)
 r : cost decrease rate ($0 < r < 1$)





Examples of modeled cost at CCIN2P3



- Nota bene :
 - Yearly costs
 - Integrates hardware replacement
 - Aware of mix of hardwares (some newer than others)
 - Similar in most sites (?)

Examples of modeled cost at CCIN2P3



- Nota bene :
 - Yearly costs
 - Integrates hardware replacement
 - Aware of mix of hardwares (some newer than others)
 - Similar in most sites (?)

Check point

$$\text{Budget (t)} = \text{Capacity (t)} * \text{Modeled Cost (t)}$$

€/year

HS06, TB

€ / HS06, TB / year

↓
site (flat) budget



↓
site capacity

↓
related to hardware cost evolution
yearly cost

- Budget ✓
- Modeled cost ✓
- → We can predict capacity increase with time

Tape...

- Tape modeling: more complex
 - Cartridges : cost decrease exponentially
 - Library, drives, maintenance, licence etc.
 - IN2P3 : cartridges cost ~ 1/3 total cost
- Make sure total cost decrease = cartridge cost decrease



Cross check

- Last verification was for CHEP (3 years ago)
- New verification this year

$$\text{Investment (t)} = \text{Capacity (t)} * \text{Modeled Cost (t)}$$

- Cross check with
 - WLCG budget @ ccin2p3 (~flat)
 - Current capacities deployed @ ccin2p3 for WLCG
 - Current values of modeled cost
 - completely independant from budget & capacities above
- Result : 20 % difference
 - Most of the uncertainty comes from tape modelization


Predictions on hardware costs

deployed capacities at CCIN2P3


modeled cost (c*)

Modeled budget


LHC: CPU resource profile



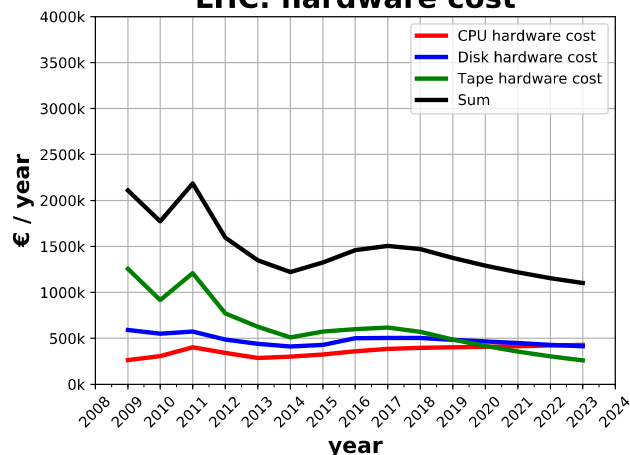
LHC: Disk resource profile



LHC: Tape resource profile



LHC: hardware cost




To bear in mind

- Life is always a bit different
 - No true exponential costs (especially Oracle tapes !)
 - Budget can vary
 - 20 % deviations around flat budget is OK
 - Media lifetime can change over time
- Model may not be exact for a given year
 - Unless the starting hypotheses hold true (rare)
 - See it as an indication over a longer period of time (a few years)
- We could try to calculate the cost of WLCG



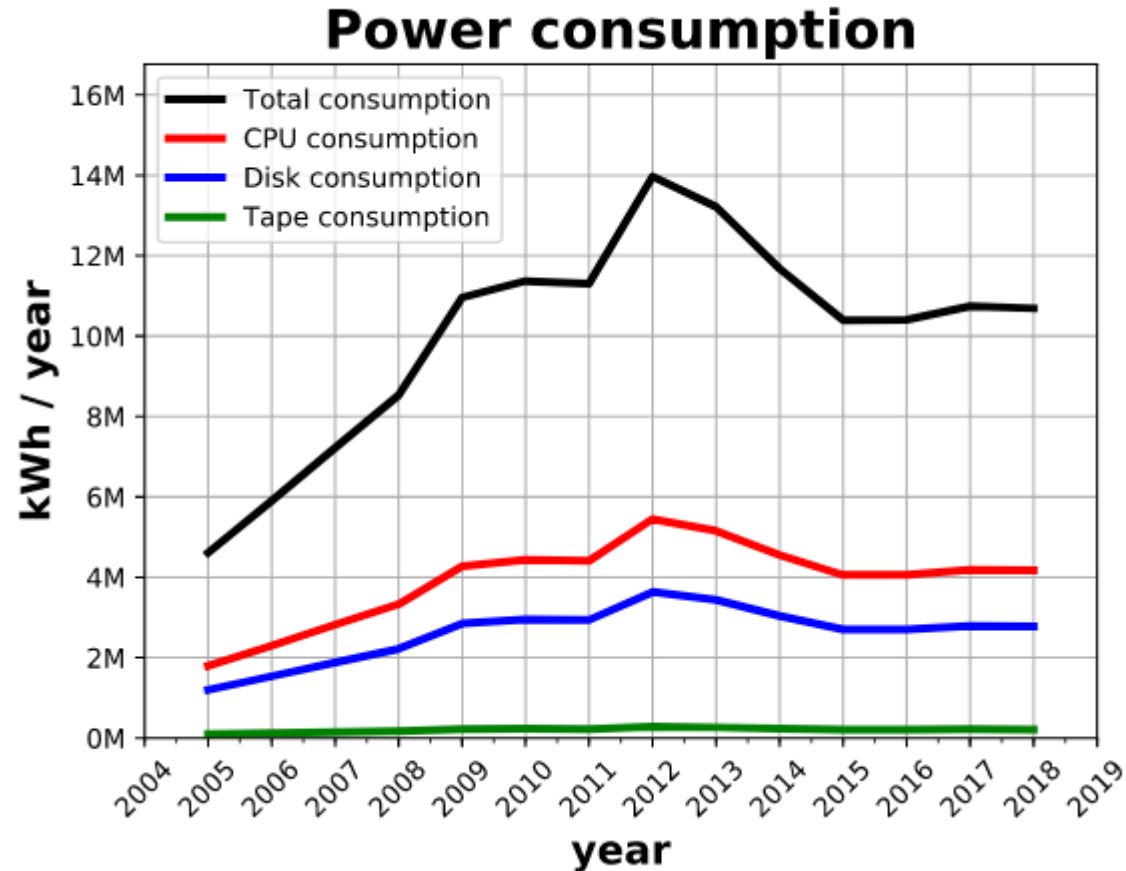
Power

Power price evolution






Power consumption breakdown

- Measurements on server PDUs
- Include PUE (cooling)
- Not easy to compute power consumption per sector
 - Work in progress
 - → currently : estimate
- Breakdown per sector
 - CPU : 39 %
 - Disk : 26 %
 - Tape : 2 %
 - Rest : 33 %



Consumption per unit of capacity



Total cost

Final Cost

hardware cost


+

power cost


=

total cost


LHC: hardware cost



LHC: power cost



LHC: total cost



Summary

- Framework to model data centres investments
 - Allows for hardware obsolescence and cost evolution
 - Computes costs per resource sector
 - Hardware + power consumption
 - Analyse cost evolutions
 - Make predictions : **translate flat budget into capacity evolution**
- Gives global picture
 - No year-to-year fluctuations
- Done for CCIN2P3
 - Principles should apply to (several) other sites
- Missing components
 - Network (difficult, ongoing)