

Cost modelling - update -

*System performance & cost modelling
2019-03-20
Renaud Vernet*

Last WLCG workshop

- « Simple » exercice on resource costs
- Significant deviations found across sites
- Several reasons
 - Local contexts
 - Differences in hardware and infrastructure
 - Understanding of metrics

ref: <http://cern.ch/go/hDT9>

**A « simple » exercise found to be
« not so simple » at all**

A thought on what we could need

- Ideal scenario
 - Common metrics
 - Common measurement methods
 - Common framework to
 - aggregate site numbers
 - model technology evolutions
 - model budgets
- Plug in experiment resource requirements roadmap
 - → get the budget evolution in sites

**We are not there yet.
We may as well not do that .**

Resource cost model

Available tool

- Spreadsheet on googledocs

ref: <https://goo.gl/XGntJc>

- Infrastructure expenses only

- Can try to extend scope

- Site input params

- Hardware price (and evolution)

- Power price and efficiency

Example with dummy data

SITE DATA (input params)

Category	Metric	current value	yearly evolution type	yearly evolution rate
Local power situation	Datacenter PUE	1,7	none	none
	Power price	0,10 EUR/kWh	linear	0,030 EUR/kWh.year
	CPU	2,00 W/HS06	exponential	-18,0 %/year
Power consumption	Disk	10,00 W/TB	exponential	-17,0 %/year
	Tape	1,00 W/TB	exponential	0,0 %/year
	CPU price	10,00 EUR/HS06	exponential	-15,0 %/year
Procurement	Disk price	100,00 EUR/TB	exponential	-15,0 %/year
	Tape cartridge price	10,00 EUR/TB	exponential	-25,0 %/year
	Tape budget	cartridge budget over total tape budget	40%	
Hardware lifetime	CPU lifetime	5 years		
	Disk lifetime	6 years		
	Tape lifetime	7 years		


Site input

Capacity planning (input params)

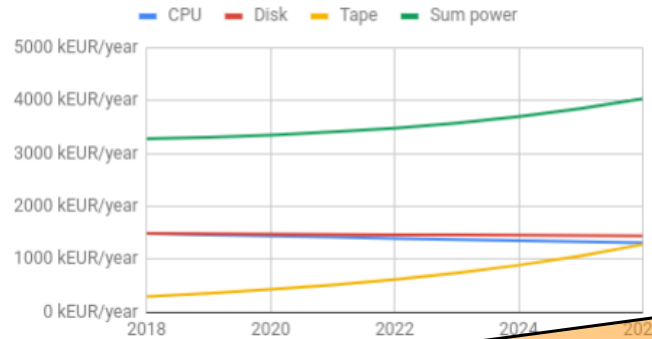
	current capacity	capacity growth rate
CPU	500 kHS06	20,0 %/year
Disk	100 PB	20,0 %/year
Tape	200 PB	20,0 %/year

Experiment input


Hardware costs



Power costs



Total cost (hardware + power)



Resulting budgets

But...

- What are site inputs ?
- Can we use a global tool ?
- Are sites going to evolve differently ?
 - Specialisation ? Diversification of technologies ?

→ need to address site expenses in detail

Survey to sites

- Survey sent to T1s in September

ref: <https://goo.gl/H5nx9W>

- All T1s contacted, few T2s
 - Fewer exchanges, provide tape, substantial capacities
- T2s very welcome to participate nonetheless

- Typical questions

How much do you pay for 1 HS06?
How much do you pay for 1 kWh?
What evolution over the last years?

- More general questions

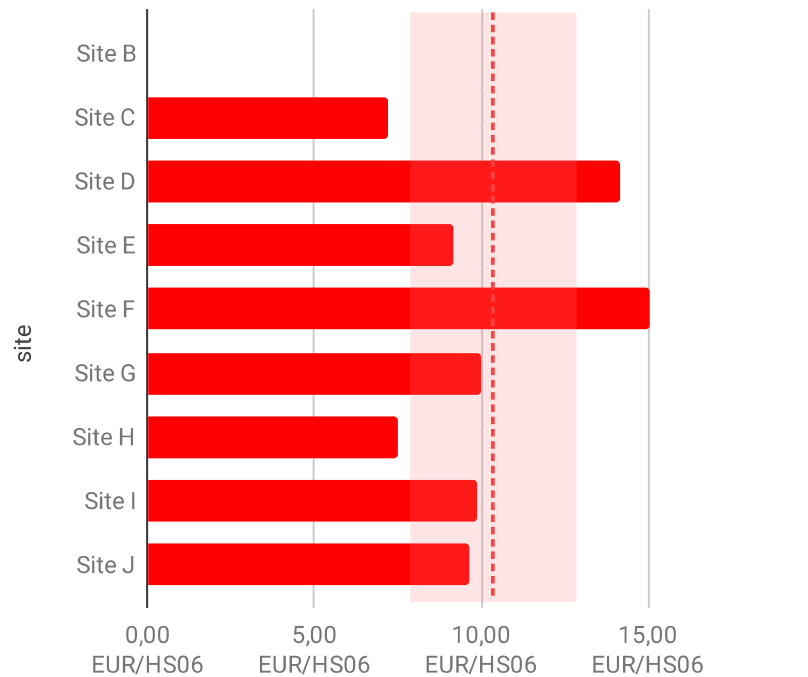
Library/tape technos
Connectivity
Type of cooling
...

Survey to sites : results

- 9 sites answered the survey (thanks!)
 - 8 T1s ; 1 T2
- Most sites could not answer all the questions
 - so we do with what we have...
- Global picture
 - All sites pay for pledged resources
 - A few sites pay for network connectivity
 - A few sites have free power
 - PUE is not necessarily known


- Current cost and yearly decrease

CPU cost (2018)




10.3 €/HS06

CPU cost yearly decrease rate

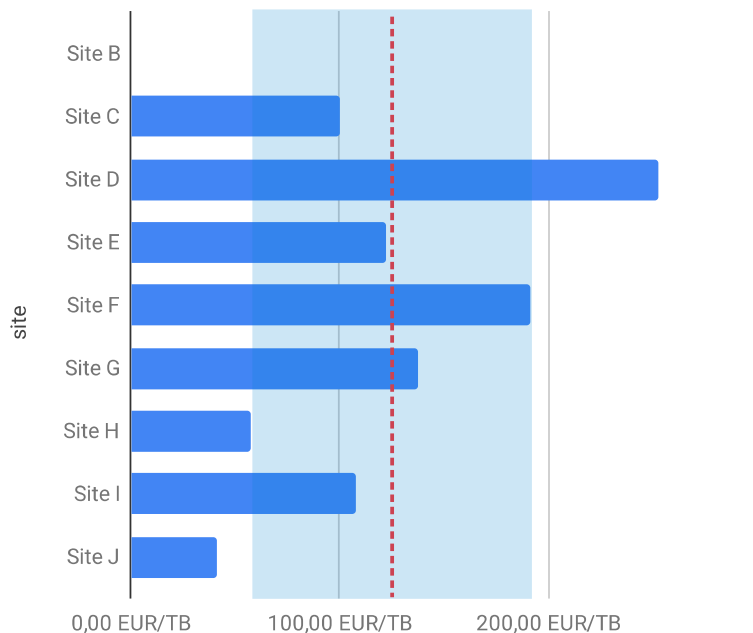


-12 %/year



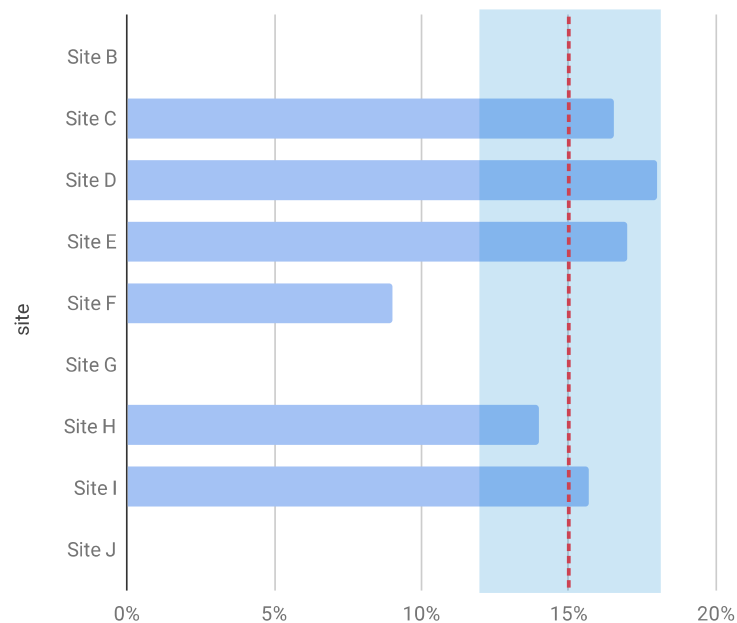
- Current cost and yearly decrease

Disk cost (2018)



126 €/TB

Disk cost yearly decrease rate



-15 %/year

Tape


- Current cost and yearly decrease

Tape cartridge cost (2018)



22 €/TB

Cartridge cost yearly decrease rate

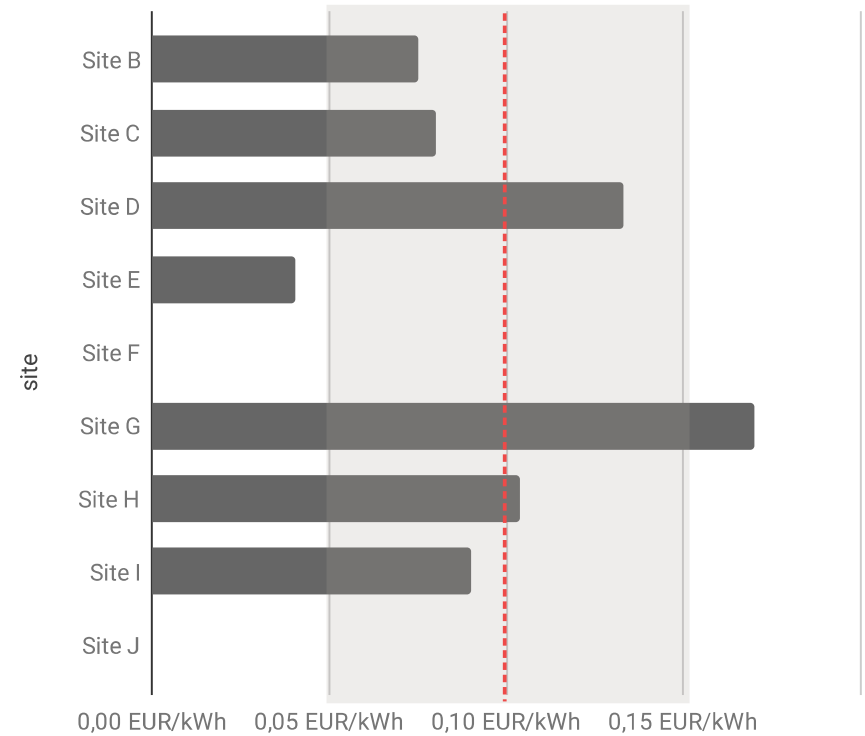


-14 %/year

Power

- Current power price

Power price (2018)



Results on resource costs at sites

- First bunch of results obtained
- Essentially **resource costs**
 - Little information on other topics
- Substantial dispersion in costs across sites
 - Local context & performance effect
 - but not « massive » : $\sigma \sim 20-50 \%$

But some metrics follow exponential behaviour with time

→ even 20 % error makes a real difference!

→ more stats for (hopefully) more precision

Addressing the TCO...

TCO ?

- TCO = Total Cost of Ownership
- Why do we need TCO ?
- How to calculate a data center resource TCO ?
- What should I include in the TCO ?
- 2 approaches
 - "atomic TCO" : all hardware & equipment + include all other relevant sources of expense
 - "holistic TCO" : include all expenses of data center


- Inspired from CERN procurement model
 - IT
 - Server, rack, PDU, disk arrays and IT power consumption
 - Network equipment (NIC, switch, router, uplink)
 - Includes losses and redundancies
 - Building (PUE)
- Manpower salary + taxes
 - Admin, operator, network expert
 - Holidays, overhead
 - Trainings etc.

Sources :


<https://goo.gl/97HXbD>
<https://goo.gl/ctiMkE>

Atomic TCO

CPU server expenses



Storage server expenses




→ "IT" is the main source of expense

- CCIN2P3 full budget (except tertiary)
- Full yearly budget sorted in categories
 - IT
 - Facility
 - Human effort
- Exercice done for 6 consecutive years [2012 – 2017]
 - Then averaged per year

Holistic TCO

- Average over period 2012 → 2017



→ "IT" is NOT the main source of expense

Holistic TCO

Atomic TCO

Facility

- Building & equipment
- PUE

IT

- Filled racks
- Network
- Power consumption

Human effort

- System admin
- Service operator
- Network expert

Facility

- Upgrades
- Maintenance

Other human effort

- Facility people
- Dev ops
- Administration
- Project managers
- User Support

Comments

- Need to converge on what to include in TCO
- ECAR TCO framework can help <https://library.educause.edu/resources/2015/4/tco-for-cloud-services-a-framework>
- IT resource price per unit of capacity/performance evolves quite fast (exponential)
- Human resources do not
 - constant budget dedicated to the « manpower and infrastructure » enveloppe
- Capacity & performance evolution over time
 - Driven by resource cost evolution, not manpower cost
 - Ratio « IT budget » / « TCO » does not influence much

Summary

- Survey on site costs → clearer global picture of main sites
- Interactive capacity evolution forecast tool available
- Resource TCO method still in discussion
- Atomic TCO tool can be played with
- Network expenses not modeled (yet)
 - Still a bit blurry, how to do that ?
- Interact more with computing model people ?