New developments in cost modeling for the LHC computing

Andrea Sciabà
on behalf of the HSF/WLCG System Performance and Cost Modelling WG

CHEP 2019
Adelaide, 4-8 November 2019

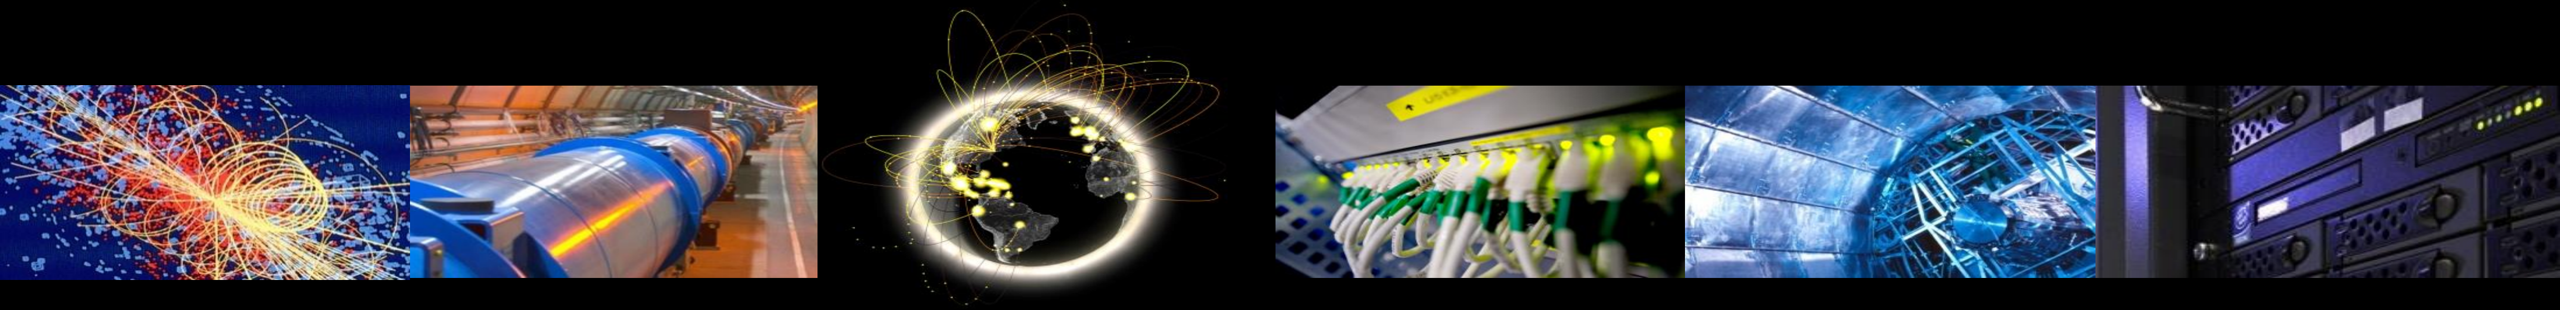

Worldwide LHC Computing Grid

The High Luminosity Challenge

- The effort to reduce the gap between the estimations of needed and available resources is continuing
  - Changes in the computing models
  - Improvements in software performance
  - GPU computing more accessible
  - Market movements driving down CPU prices
- Still, we are not there yet!

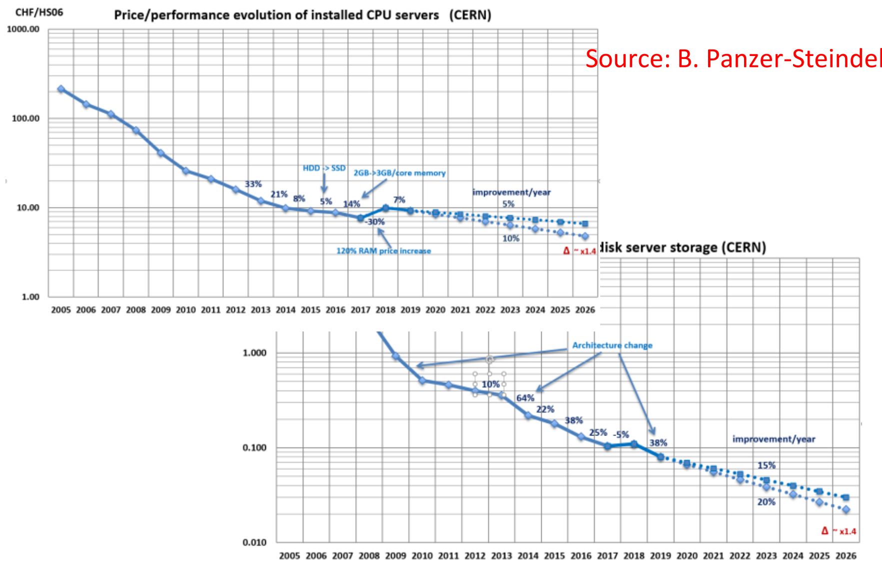

Source: B. Panzer-Steindel

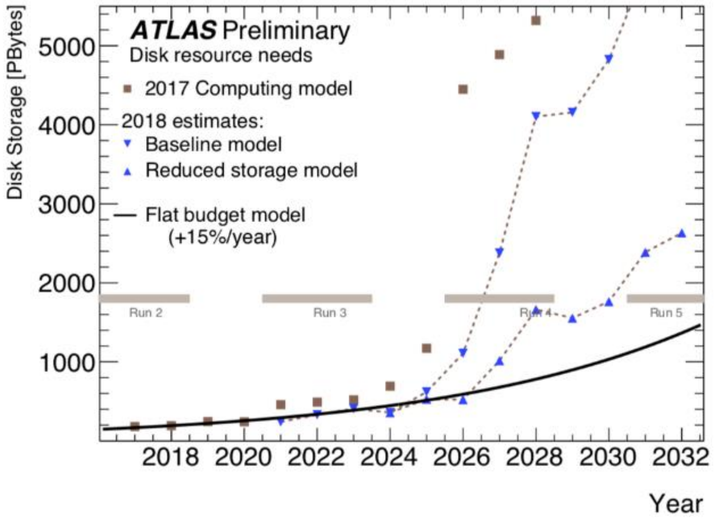

Source: D. Costanzo

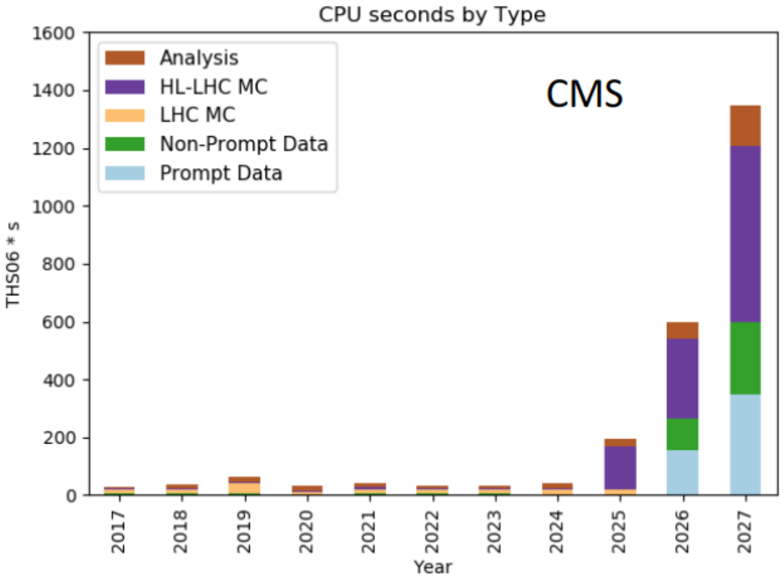

Source: D. Lange

WLCG
Worldwide LHC Computing Grid

The Working Group

- Created exactly two years ago, to help WLCG fitting into the available resources for Run3 and Run4
  - Context
    - In summer 2017, the best estimates gave a 10x factor between required and available resources at HL-LHC
    - The WG was born as the place where to focus on this issue from a cross-experiment, cross-site perspective
  - Mandate
    - Develop a deep understanding of current workloads, resource utilisation, and site costs
    - Explore future scenarios, estimate possible improvements in efficiency
    - Develop tools and methods for the above
  - Organisation
    - Conveners: J. Flix, M. Schulz, A. Sciabà
    - ~35 active members from LHC experiments, sites and including software experts
    - By-weekly meetings (https://indico.cern.ch/category/9733/)

WLCG Worldwide LHC Computing Grid

Initial and current areas of work

- Several activities were started to cover the various aspects that together determine the cost of computing
  - Identify representative experiment workloads
  - Define which metrics best characterise such workloads
  - Understand how to estimate resource needs
  - Define a process to evaluate the TCO of an infrastructure
- More activities followed
  - Measure the impact of new storage configurations
  - Identify potential areas for savings
  - Study the effect of resource constraints on experiment workloads
  - Detailed studies of hardware resource utilisation
  - Define a viable mechanism to pledge HPC resources in WLCG

Workload characterisation

- Identification of (today's) representative workloads
  - CPU intensive (MC simulation)
  - CPU + I/O (reconstruction)
  - I/O (digitisation with pileup, merge, skimming, …)

- Basic metrics
  - CPU (i.e. how much do I use?)
  - Memory (i.e. do I swap a lot?)
  - Disk I/O (i.e. how much do I read/write?)
  - Network (i.e. what are the access patterns?)
- Available tools
  - PrMon ([GitHub]), using ProcFS
  - Trident ([GitHub]), using hardware counters
- Testbeds
  - CERN and PIC

Job	Events	Processes/ threads	Wallclock time (sec)	CPU efficiency	Time per event (sec)	Memory per core (GB)	Read rate per core (MB/sec)	Write rate per core (MB/sec)
ALICE sim	1000	1	10901	100%	10.90	0.96	0.0788	0.1660
ATLAS sim G4	1000	8	33627	100%	269.02	0.44	0.0152	0.0090
ATLAS digireco	2000	8	13981	87%	55.92	1.12	0.3174	0.2412
ATLAS deriv	95741	8	8401	98%	0.70	1.20	0.6849	0.0705
CMS gensim	1000	8	2651	99%	21.21	0.19	0.0473	0.0377
CMS digi	1000	8	737	78%	5.90	0.65	0.2854	0.3004
CMS reco	1000	8	1221	83%	9.77	0.45	0.3073	0.2153
LHCb gensim	10	1	1782	100%	178.20	0.89	0.3115	0.0117

WLCG
Worldwide LHC Computing Grid

Example: CMS workloads

GEN-SIM

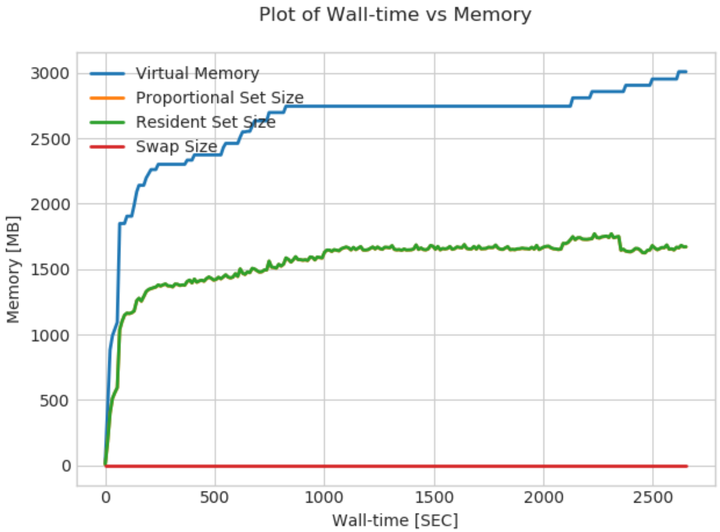

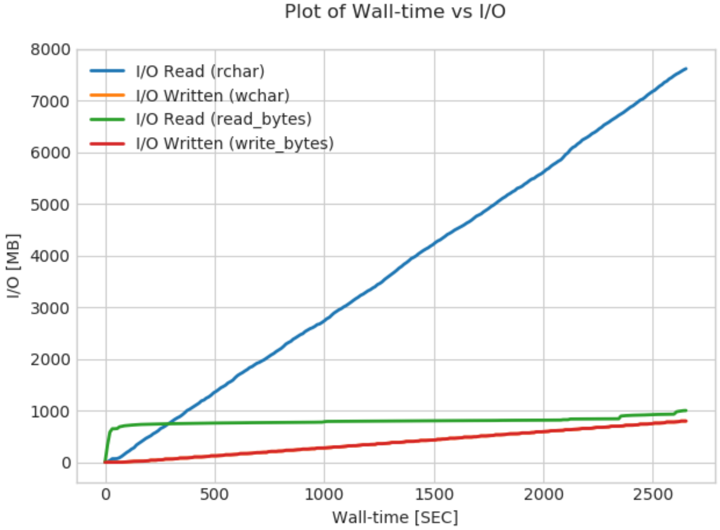

DIGI

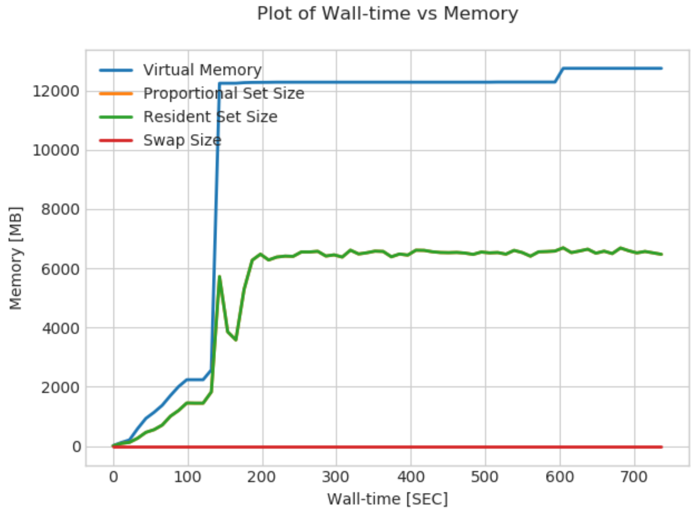

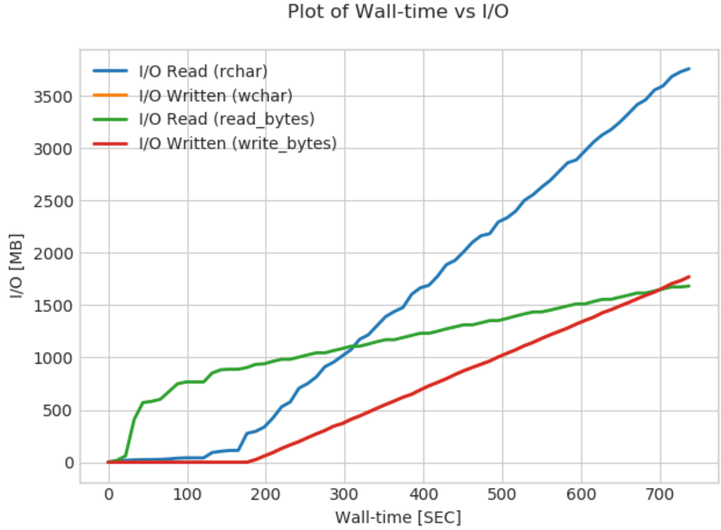

RECO

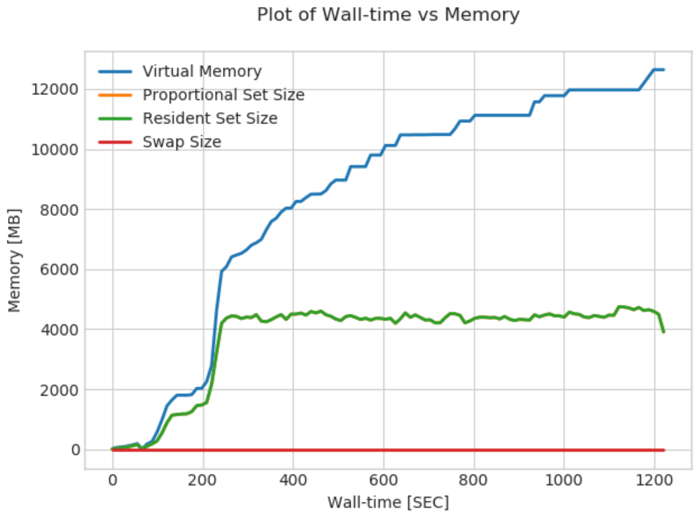

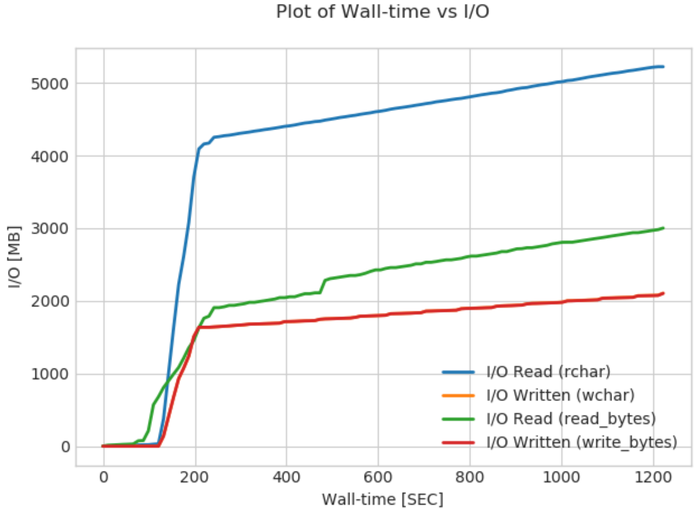

WLCG Worldwide LHC Computing Grid

Workload modelisation under resource constraints

- Apply restrictions to memory, network bandwidth and latency
- Measure the effect with PrMon
- Parametrise the time series
  - Using the CPOP (Continuous-piecewise-linear Pruned Optimal Partitioning)
- Next step: parametrise also the impact of different restrictions

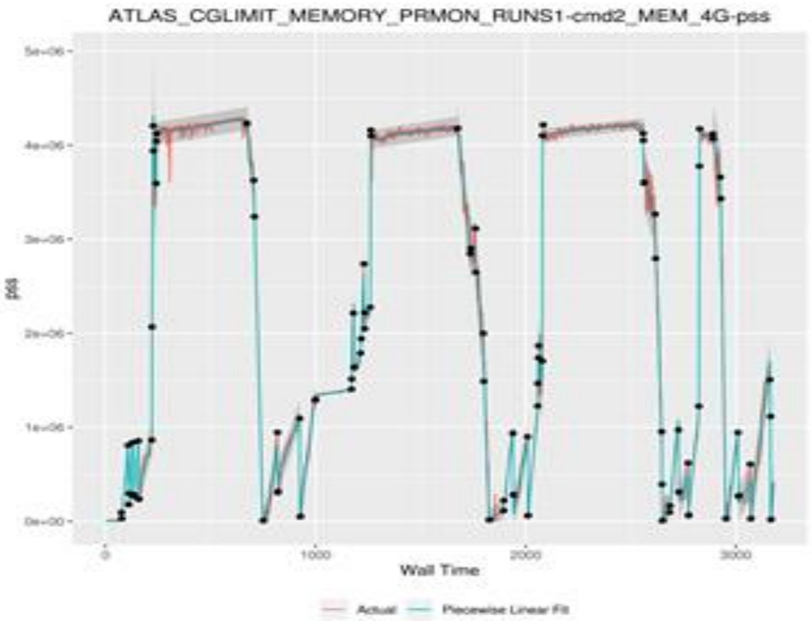

16 threads/processes

ATLAS G4 simulation memory and swap when memory is limited

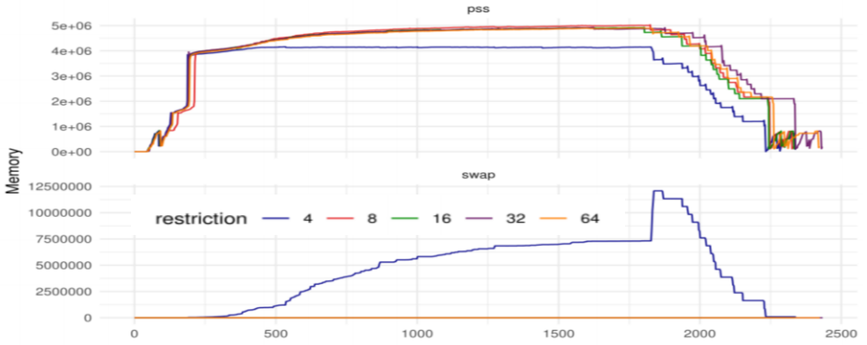

ATLAS digi-reco bytes read from disk when bandwidth is limited

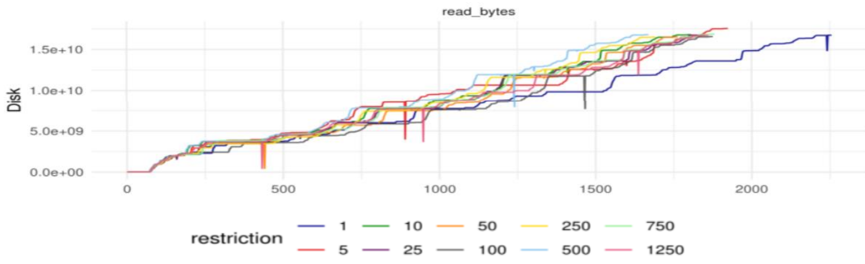

CMS digi bytes read from network when latency is limited

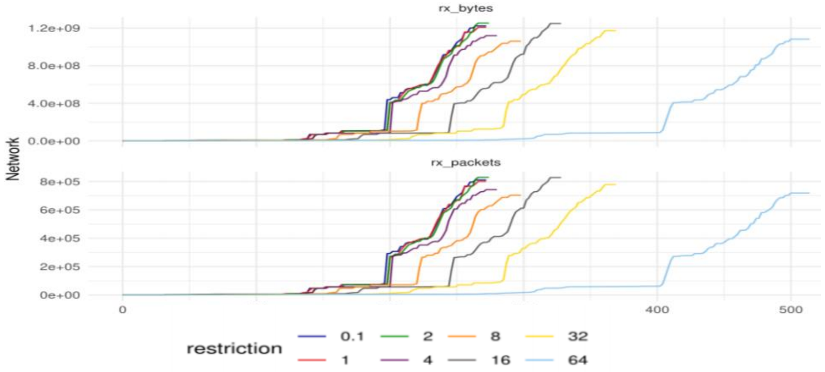

[1] Fearnhead, P., Maidstone, R., & Letchford, A. (2019). Detecting Changes in Slope With an L0 Penalty. *Journal of Computational and Graphical Statistics*, *28*(2), 265-275.

Effect of compiler optimisations on G4

- Running standalone Geant4 simulations, using either dynamic or static libraries
- Two versions of gcc compiler were used, and four gcc optimization levels (Os, O1, O2 and O3)
- G4 simulation with a CMS-based geometry has been used as a reference

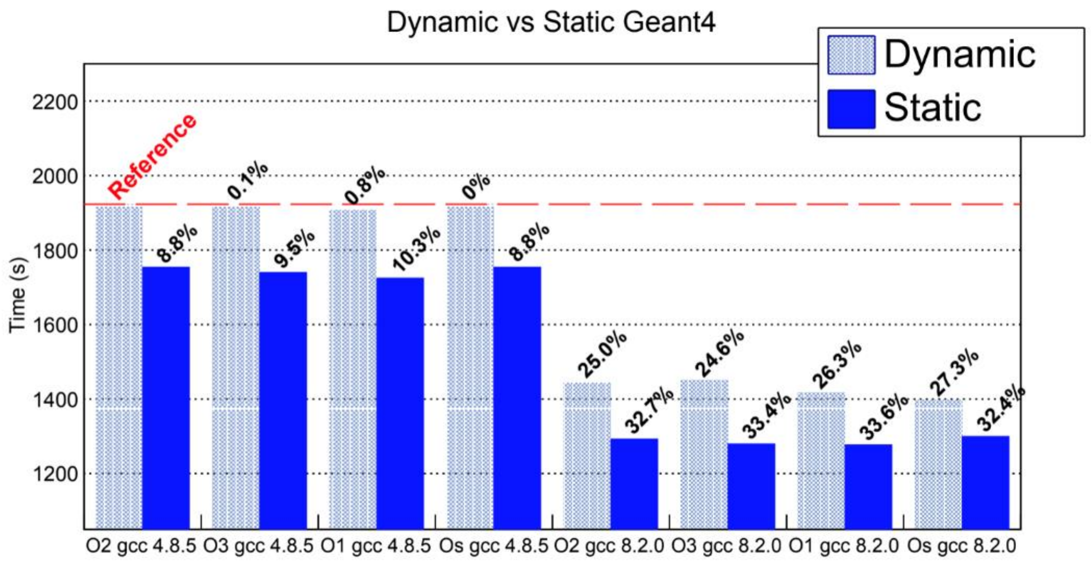

- The static approach, for both gcc versions, generally reduces the execution time by more than 10%
- Switching from gcc 4.8.5 to 8.2.0 results in an average of 30% improvement in the execution time
- See Caterina Marcon’s talk, “Impact of different compilers and build types on Geant4 simulation execution time”, on Thursday

WLCG Worldwide LHC Computing Grid

CMS GENSIM performance vs. CMSSW version

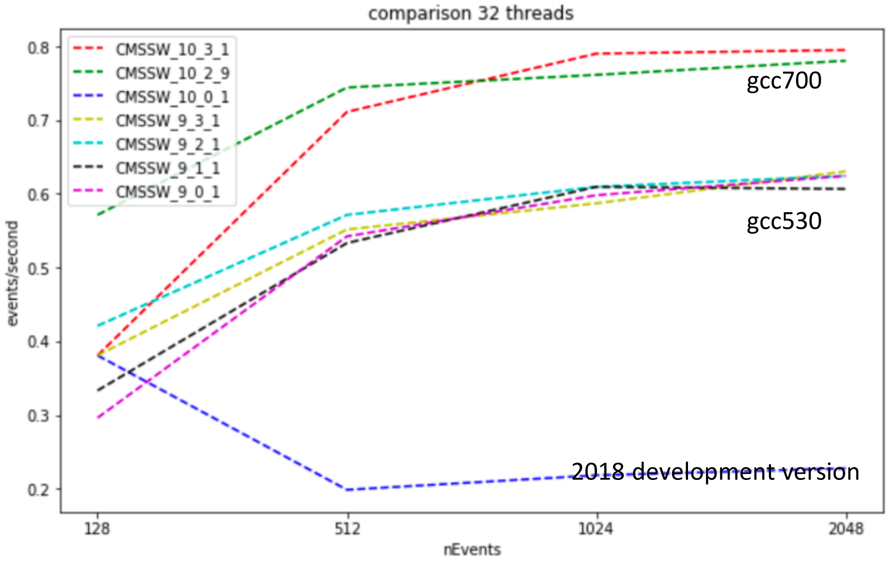

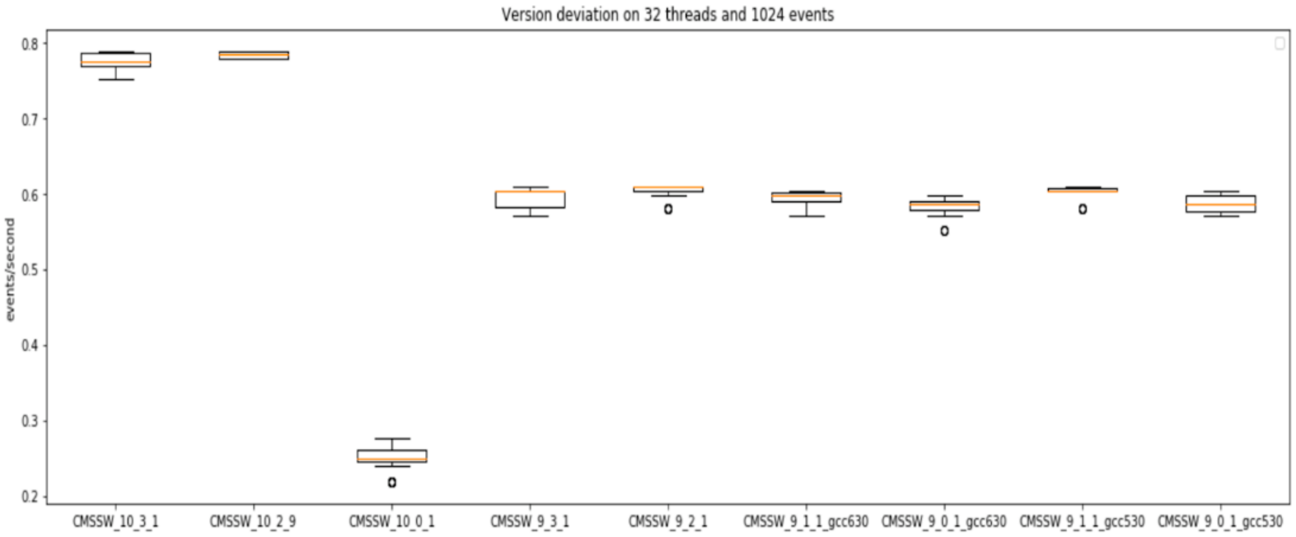

Source: A. Alonso

- Given that GEN-SIM is rather stable, differences are attributed to different version of gcc
  - +25% in performance consistent with Caterina's results

Resource estimation

- The initial goal was to define a common framework for modelling the computing requirements of the LHC experiments
  - Models as collection of parameters and generic calculations
  - Allow to play with different scenarios

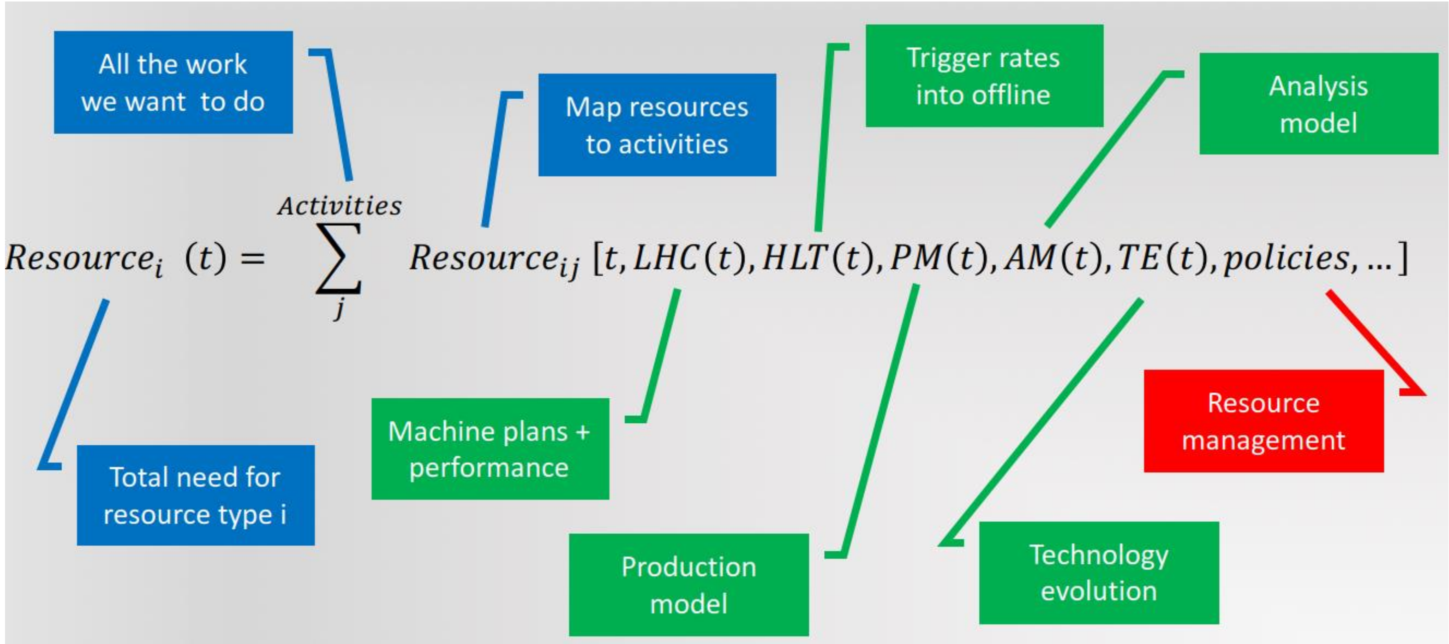

Source: D. Lange

WLCG Worldwide LHC Computing Grid

General observations from resource estimation

- Experiments have largely converged in terms of parameter sets and model complexity
  - But not on implementation, which is “nice to have” but not necessary
- Models address CPU, disk and tape
  - CMS started working on network capacity estimation
    - Starting from basic WAN needs (standard scheduled data transfers)
  - No GPUs or accelerators (still OK…)
  - Tape I/O will become a big limiting factor at HL-LHC
    - CMS estimates a x10 increase in the number of needed tape drives (source: D. Lange)
  - Need a good model for cost evolution
    - Already being developed in this WG!
- Effect of R&D or big changes in computing models difficult to anticipate
  - See D. Lange, “Modeling of the CMS HL-LHC computing system”, later today, track 9

Site cost estimation models

- Developed a method to calculate the TCO of the needed resources
  - “Atomic” or bottom-up approach implemented as a spreadsheet that can be copied and used by any site
  - “Holistic” or top-down approach starting from the data centre budget and deducing TCO per unit of resource capacity
- Model for cost evolution using as input prices and trends to calculate required budgets (spreadsheet)
  - Shown for the first time at CHEP2015 (R. Vernet, “A model to forecast data centre infrastructure costs”)
- Surveyed sites (confidentially) to understand the major costs
  - Current costs of CPU, disk, tape and power
  - Recent evolution
  - Variation across sites
- The work is not finished!

CC-IN2P3

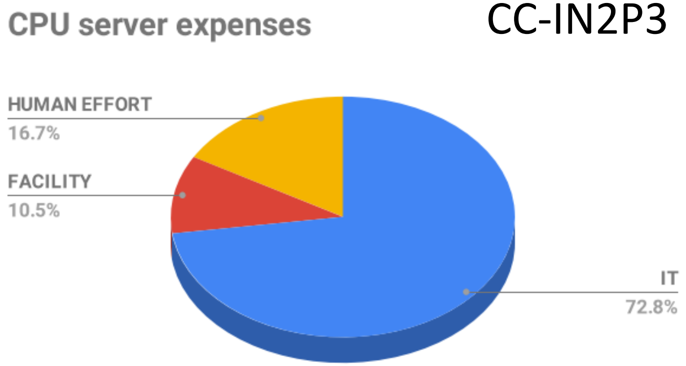

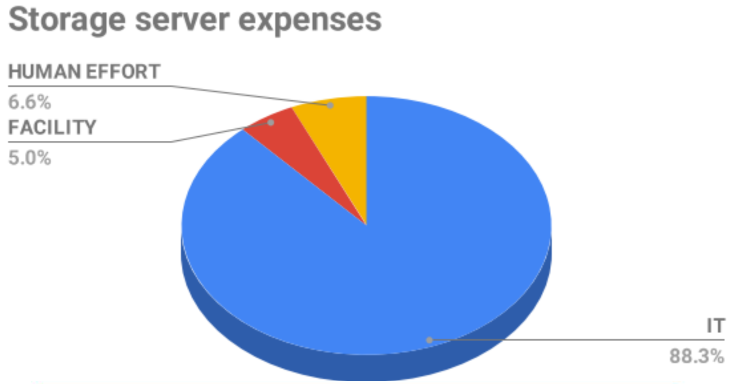

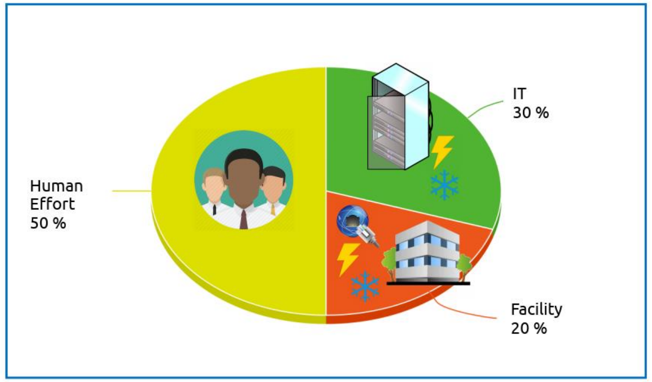

Source : R. Vernet

Cost evolution model

Example with dummy data

SITE DATA (input params)

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

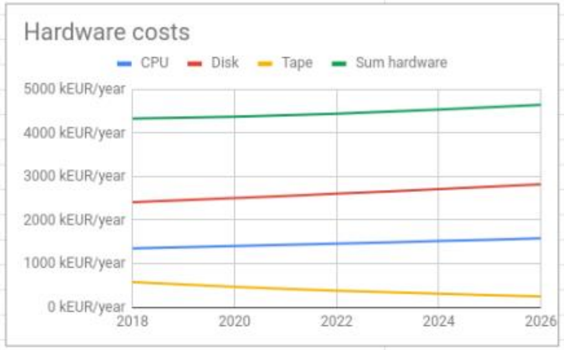

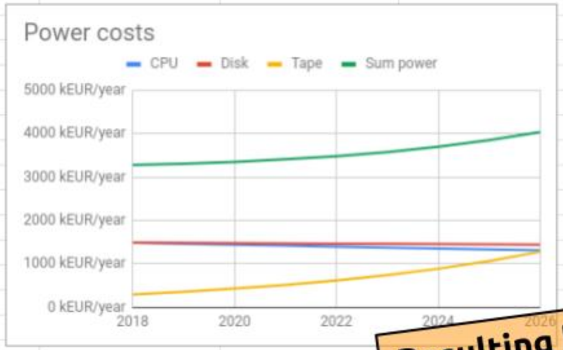

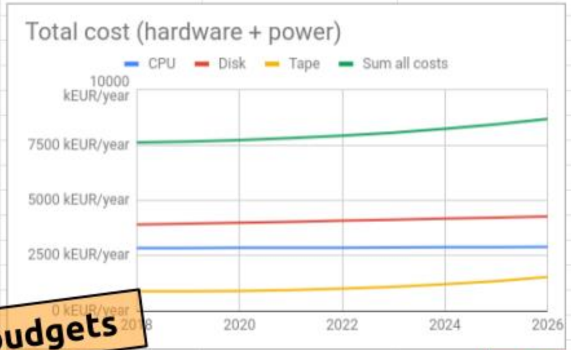

Resulting budgets

WLCG Worldwide LHC Computing Grid

Site cost Tier-1 survey

- Launched in September 2018 a survey among Tier-1 sites (and open to Tier-2s) to understand their costs for CPU, disk and tape
  - Questionnaire available [here]
  - Eight Tier-1s and one Tier-2 answered
- Average costs
  - CPU: €10.3/HS06, -12%/y
  - Disk: €126/TB, -15%/y
  - Tape: €22/TB, -14%/y
- 20-50% spread among sites

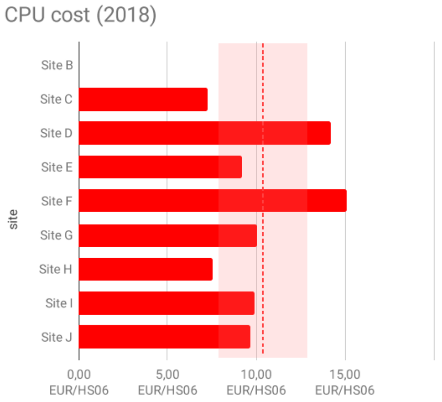

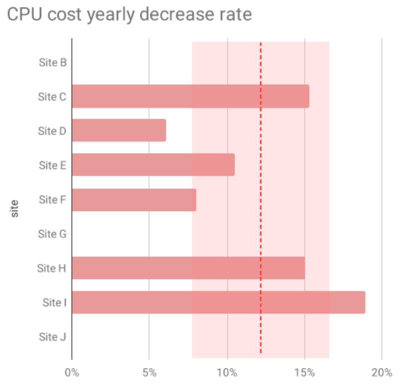

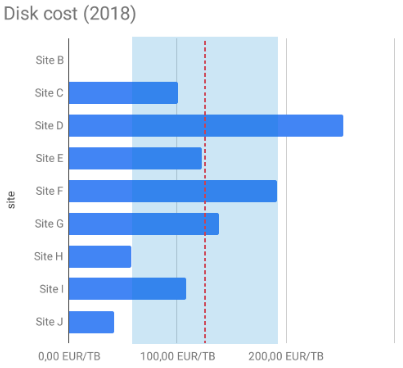

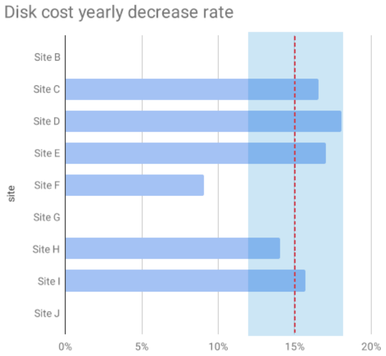

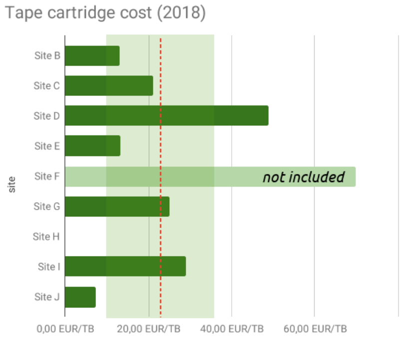

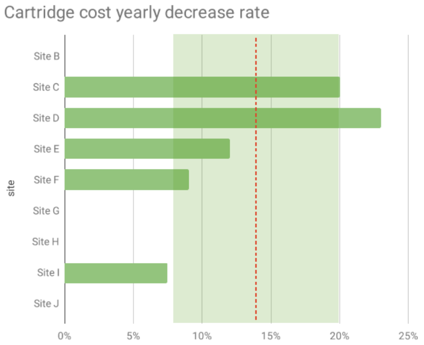

Source: R. Vernet

WLCG Worldwide LHC Computing Grid

Storage modeling and popularity studies

- Investigated the feasibility of storage caches using popularity data from ATLAS and CMS
  - Work is now converging into the DOMA (Data Organization Management Access) working group in WLCG
  - See Markus' talk, "Analysis and modeling of data access patterns in ATLAS and CMS", tomorrow in track 4
- Data access and popularity studies at PIC
  - Provides a more detailed picture than the general approach (including time of file creation and deletion) using the dCache billingDB
  - See Pepe's talk, "CMS data access and usage studies at PIC Tier-1 and CIEMAT Tier-2" tomorrow in track 4

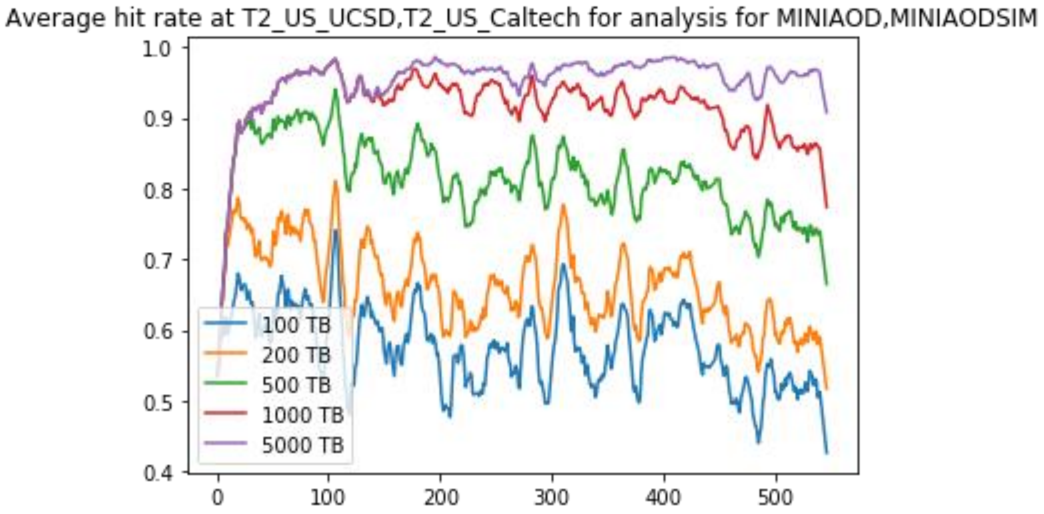

PIC Data Access and Popularity studies [CMS]

[talk]

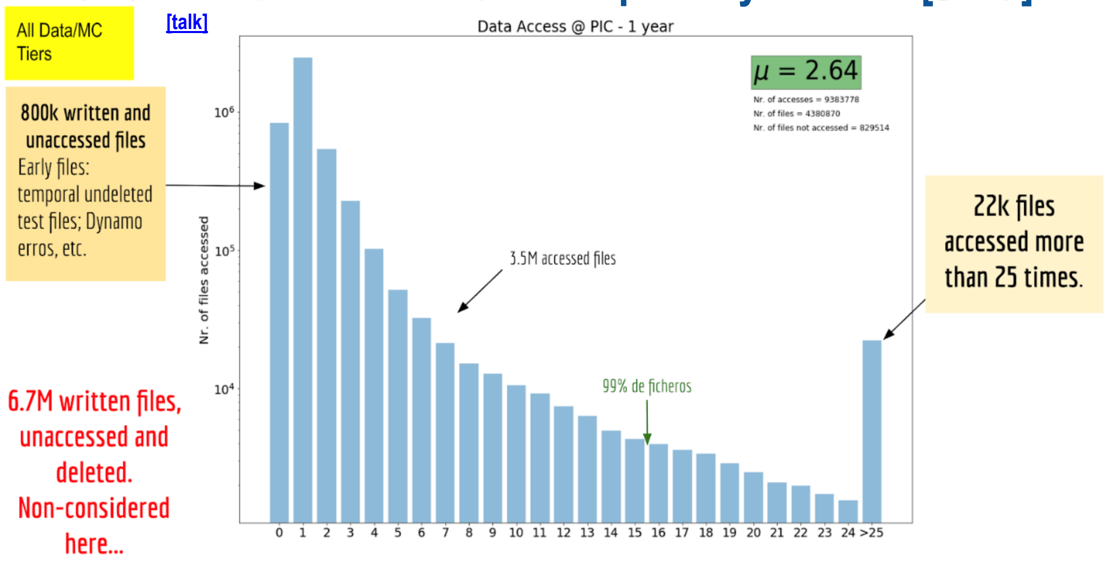

WLCG Worldwide LHC Computing Grid

Other areas of potential savings

- Many "small" improvements can stack to provide significant gains
  - Some of these estimates got more precise as a result of recent studies (data popularity, TCO estimates, effect of data loss, …)

Summary of how this evolved:

Change	Effort Sites	Effort Users	Gain
managed storage on 15 sites + caches elsewhere	Some on large sites/gain on small sites	little	40% decrease in operations effort for storage
Caches at most sites (dataLake strawman)	Some everywhere	Frameworks some	15% of storage
Reduced data redundancy	Some large sites	Frameworks some	30-50% disk costs
Reduced data replication and cold data	little	Frameworks some	30% disk costs
Compact data formats for analysis	none	Some	>15% disk costs
Scheduling and site inefficiencies	Some	Some	10-20% gain CPU
Reduced job failure rates	Little	Some-Massive	5-10% CPU
Compiler and build improvements	None	Little-Some	15-20% CPU
Improved memory access/management	None	Realistic	10%-15% CPU
Exploiting modern CPU architectures	None	Massive	100% CPU
Paradigm shift algorithms	Some	Massive-Infinite	Factor 2-100 CPU
Paradigm shift online/offline data	Little	Massive-Infinite	2-10 CPU 10-20 Storage

Source: M. Schulz

WLCG Worldwide LHC Computing Grid

Pledging HPC resources

- First attempt at formalising a way of quantifying the computing capacity of HPC resources with respect to the LHC computing workloads in the WLCG context
- The problem
  - HPC systems are (very) different from WLCG commodity resources
  - Usability can be a serious issue (network and storage access, scheduling, authorization…)
  - Diversity of architectures (x86, POWER, ARM, GPUs, FPGA…)
- How to measure the value of an HPC resource?
  - At the end of the day, the value is the achievable throughput
  - For each eligible workload, measure throughput 1) on the HPC system and 2) on a conventional system with a known HS06 (or Hepscore) value
  - Express 1) as an "equivalent HS06" (potentially different for different workloads)
- How to measure how well the HPC resource is used?
  - For example by measuring the average FLOPS consumed by the workload and comparing it with the LINPACK R_{max} FLOPS
- A final version of a document with all the details is in the works
  - Latest draft available

WLCG
Worldwide LHC Computing Grid

Where are we now?

- After two years of activity it's time to re-examine the roadmap and the goals
  - Many activities now gravitate in other working groups
    - Data access and storage cost optimisation → DOMA (access/QoS)
    - Detailed studies of workload behaviour on CPUs
      - Trident, code analysis → Various activities within experiments and HSF
    - Build and compiler studies (AutoFDO, static builds) → experiments and tools (G4, …)
  - Other have become standard activities
    - Site cost calculations
    - T1 anonymised site cost assessment
    - Compute Resource Estimation (ATLAS and CMS have now Python frameworks)
    - Pilot for storage and compute integration (BEER) has been in production
  - Focus on efficiency and performance is now common
    - Computing Schools, Working Groups within the experiments, HSF, GPU, ML…
    - Generator workshops, several activities around simulation, faster reconstruction, compact analysis formats…

Where to go

- Relocate and re-scope
  - Move activities to the better fitting working groups
    - Workload, storage and access to DOMA and Benchmarking
    - Profiling to benchmarking or HSF
  - Identify topics unique to the cost and performance modeling WG
    - Site cost calculation with multiple sites
    - Confidential statistical analysis of cost differences
    - Refinements of the Resource Estimates
  - Run topical micro-workshops with experiments and sites
    - Every 3 months could be linked to GDB, HSF-WLCG Workshops

WLCG Worldwide LHC Computing Grid

Conclusions

- This working group was established to improve our understanding of the performance and the cost of computing for LHC (and HEP) and its evolution
- The WG is active in many areas, is already achieving important results and these activities are much more mature now
- Right time to re-scope the mandate

WLCG
Worldwide LHC Computing Grid

Membership

- Alaettin Serhan Mete, Alessandra Forti, Alessandro Di Girolamo, Andrea Sartirana, Andrea Sciabà, Andrea Valassi, Andrew Sansum, Andrey Kirianov, Antonio Pérez-Calero, Bernd Panzer-Steindel, Carlos Perez Dengra, Catherine Biscarat, Chris Hollowell, Concezio Bozzi, Costin Grigoras, Daniele Bonacorsi, David Lange, David Smith, Davide Costanzo, Dirk Duellmann, Domenico Giordano, Duncan Rand, Eric Fede, Erik Mattias Wadenstein, Frank Wuerthwein, Gareth Roy, Graeme A Stewart, Helge Meinhard, Jan Iven, Johannes Elmsheuser, José Flix, Markus Schulz, Martin Gasthuber, Michel Jouvin, Michele Michelotto, Oxana Smirnova, Paul Millar, Raul Cardoso Lopes, Renaud Vernet, Servesh Muralidharan, Shigeki Misawa, Tommaso Boccali, Torre Wenaus, Xavier Espinal Curull, Xiaomei Zhang, Yves Kemp,

WLCG Worldwide LHC Computing Grid

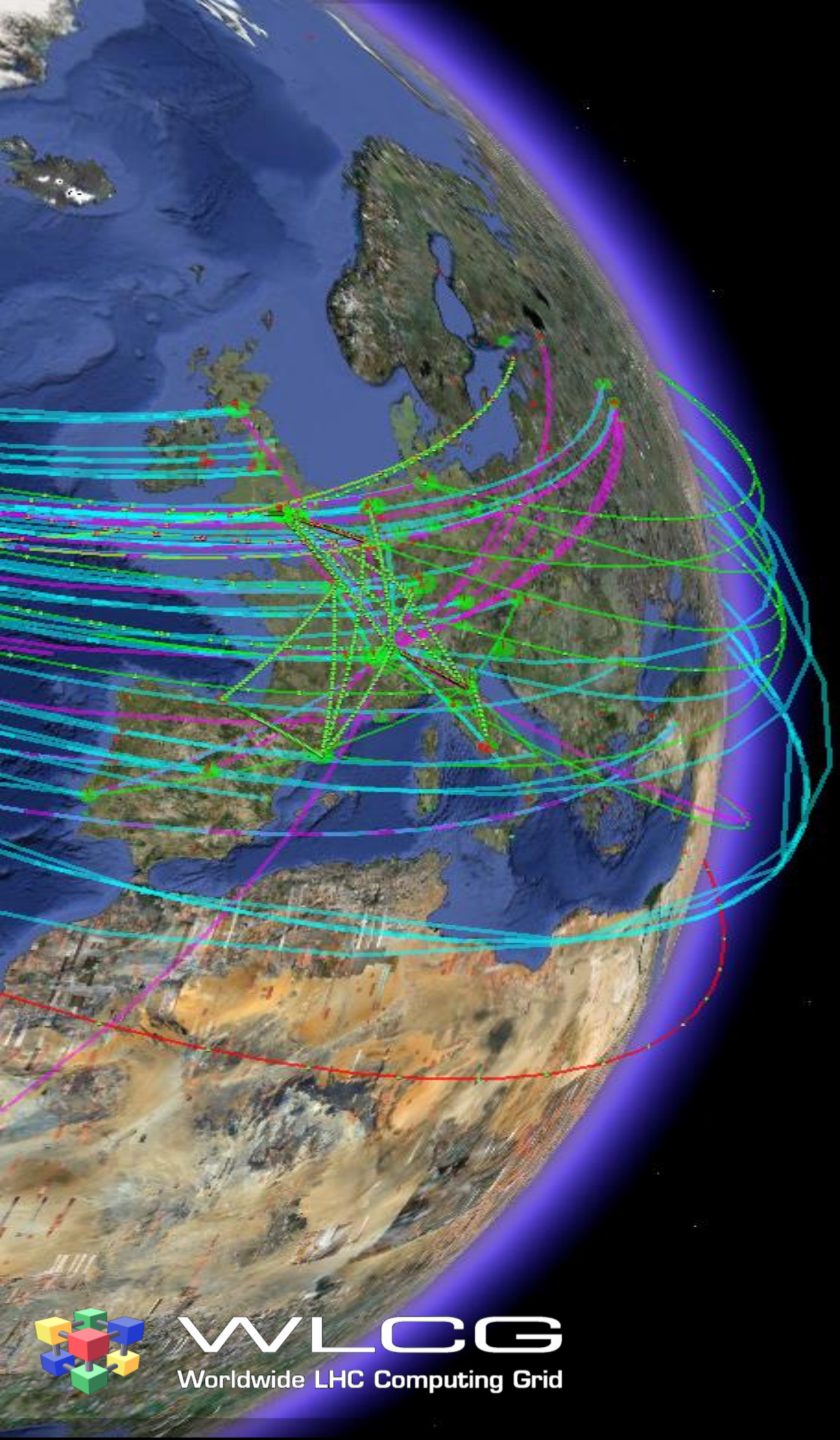

Backup slides

Measuring performance with Trident

- Measures CPU, IO and memory utilization based on hardware counters, memory and IO information, e.g. to identify bottlenecks
- Several metrics calculated
  - CPU: IPC, total cycles, top-down analysis (front-end bound, back-end bound, retiring, bad speculation)
  - Core backend utilization: compute (ports 0,1,5,6) vs memory (ports 2,3,4,7)
  - Memory: bandwidth usage, transaction classification (page-hit, page-empty, page-miss)
- Can be used to see how workloads differ (or resemble) the benchmarks we use (e.g. HS06)
- CPU counters are a powerful (but complex) tool and Trident makes them accessible

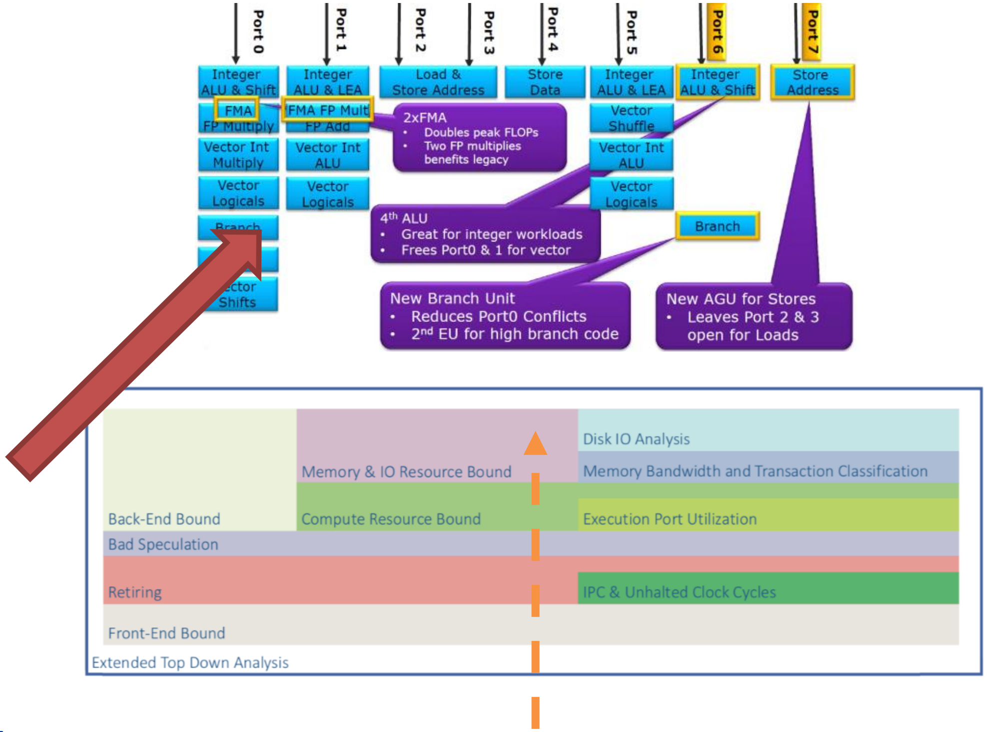

Full exploration of CPU utilization

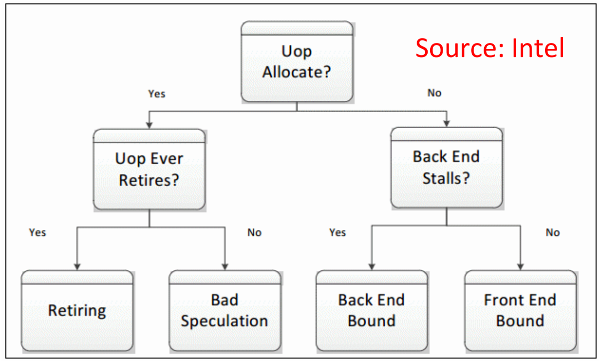

Trident plots: ATLAS Geant4

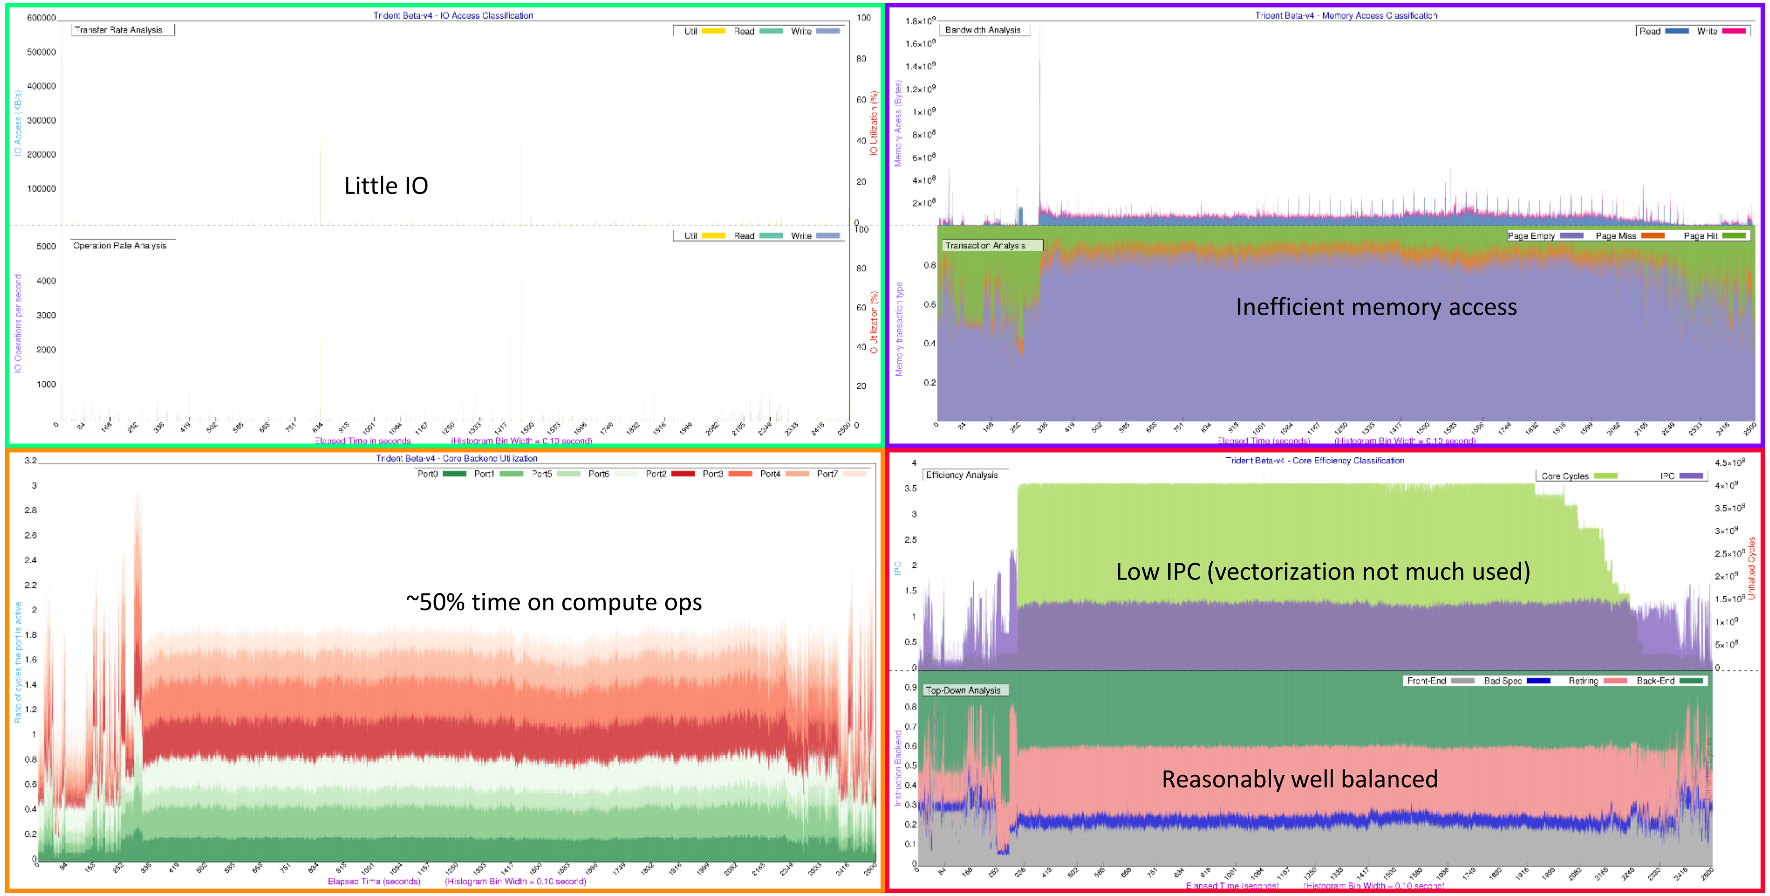

More on top-down analysis here

Source: Servesh Muralidharan

WLCG Worldwide LHC Computing Grid

Site infrastructure TCO calculation

INPUT HARDWARE COSTS	Memory (per GB)	30 EUR/GB
	Server local disk (per TB)	200 EUR
	Motherboard + CPU + housing	1600 EUR
	Rack high power (compute usage)	4000 EUR
	Rack low power (storage usage)	800 EUR
	PDU	200 EUR
	JBOD array	1400 EUR
	JBOD array disk	400 EUR
CPU SERVER CHARACTERISTICS	number of cores	40
	HS06 per core	10.0 HS06/core
	memory	128 GB
	local disk	2 TB
	max power	300 W
	NIC speed	10 Gbps
	lifetime	5 years
CPU SERVER COST	on full lifetime	5840 EUR
	per year	1168 EUR
STORAGE CHARACTERISTICS	JBOD arrays	2
	disks per array	24
	disk size	8 TB
	max power	750 W
	lifetime	4 years
STORAGE SERVER COST	disks	19200 EUR
	server + shelves	8640 EUR
	full storage, on full lifetime	27840 EUR
	full storage, per year	6960 EUR
RACK CHARACTERISTICS	lifetime	10 years
	low power	4000 W
	high power	12000 W
RACK COSTS	# CPU servers per rack	40
	cost per CPU server, on full lifetime	100 EUR
	cost per CPU server, per year	10 EUR/year
	# storage servers per rack	5.3
	cost per storage, on full lifetime	150 EUR
	cost per storage server, per year	15 EUR
PDU & power outlets & 1 Gb link	lifetime	10 years
	cost per year	20 EUR/year
	cost per CPU server	20 EUR/year
	cost per storage server	50 EUR/year
NETWORK	cost per CPU server per year	33 EUR/year
	cost per storage server	70 EUR/year
BUILDING	cost of new building	5 MEUR/MW
	lifetime	30 years
	cost per CPU server	50 EUR/year
	cost per storage server	125 EUR/year
POWER	power cost	120 EUR/MWh
	average power usage CPU	80%
	average power usage storage	60%
	datacenter PUE	1.65
	cost per CPU server	416 EUR/year
	cost per storage server	781 EUR/year
HUMAN OPERATIONS	low level operation time (per CPU server)	2.5 h/year
	high level operation time (per CPU server)	2.5 h/year
	network operations per server ([er CPU server)	0.6 h/year
	low level operation time (per storage server)	2.7 h/year
	high level operation time (per storage server)	5.5 h/year
	network operations per server (per storage server)	0.8 h/year
HUMAN OPERATION COSTS	low level (per CPU server)	131 EUR/year
	high level (per CPU server)	170 EUR/year
	network (per CPU server)	39 EUR/year
	low level (per storage server)	141 EUR/year
	high level (per storage server)	375 EUR/year
	network (per storage server)	52 EUR/year

SUMMARY

CPU			per server	per HS06	breakdown
IT	73 %	server	1168 EUR/year	2.92 EUR/HS06/year	57 %
		rack	10 EUR/year	0.03 EUR/HS06/year	0 %
		IT power	252 EUR/year	0.63 EUR/HS06/year	12 %
		network	33 EUR/year	0.08 EUR/HS06/year	2 %
		power outlets & PDU	20 EUR/year	0.05 EUR/HS06/year	1 %
FACILITY	10 %	building	50 EUR/year	0.13 EUR/HS06/year	2 %
		non-IT power	164 EUR/year	0.41 EUR/HS06/year	8 %
HUMAN EFFORT	17 %	human	341 EUR/year	0.85 EUR/HS06/year	17 %
TOTAL		total	2038 EUR/year	5.10 EUR/HS06/year	100 %

STORAGE			per server	per TB	breakdown
IT	88 %	server + shelves	2160 EUR/year	5.63 EUR/TB/year	25 %
		disks	4800 EUR/year	12.50 EUR/TB/year	56 %
		rack	15 EUR/year	0.04 EUR/TB/year	0 %
		IT power	473 EUR/year	1.23 EUR/TB/year	6 %
		network	70 EUR/year	0.18 EUR/TB/year	1 %
		power outlets & PDU	50 EUR/year	0.13 EUR/TB/year	1 %
FACILITY	5 %	building	125 EUR/year	0.33 EUR/TB/year	1 %
		non-IT power	307 EUR/year	0.80 EUR/TB/year	4 %
HUMAN EFFORT	7 %	human	568 EUR/year	1.48 EUR/TB/year	7 %
TOTAL	100 %	total	8569 EUR/year	22.32 EUR/TB/year	100 %

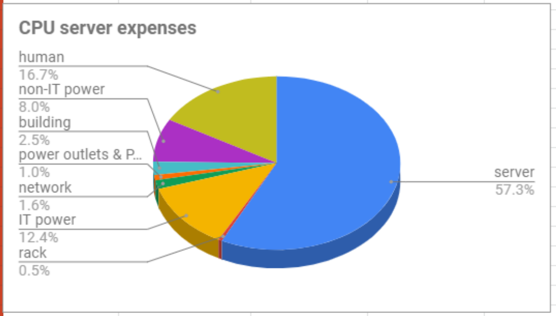

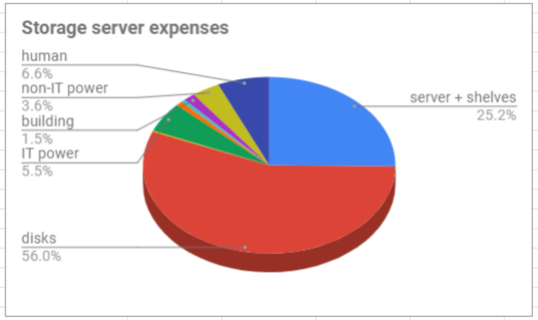

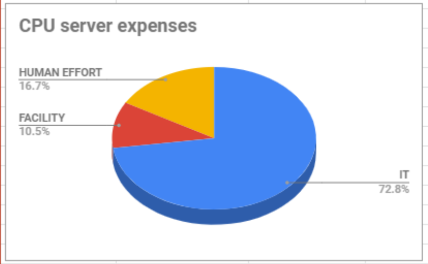

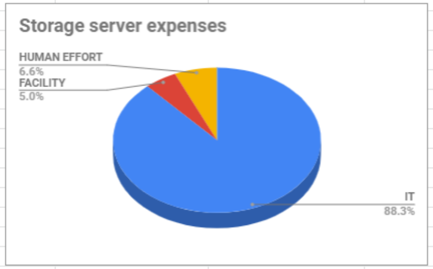

Differences in TCO

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