



# Monitoring & Metrics

or

how to keep track of analysis and analysis infrastructure evolution

WLCG-HSF Pre-CHEP Workshop,  
Norfolk VA

6 May 2023

Dirk Duellmann, CERN

Profiting from discussions with B. Panzer, M. Schulz, A. Sciaba, A. Naumann, J. Blomer, A. Peters, J. Moscicki, G. McCance



# Which Metrics? Why Analysis? Why me?

- Why me?
  - did physics, infrastructure and access pattern analysis and an data analysis course (incl. deep learning), some of the above ages ago!
  - was not present at the meeting...
- We, also CERN IT, need an answer to
  - Q: does it pay-off to invest in non-standard analysis infrastructures?
    - if yes, which would be best to complement existing services?
- I assume, “pay-off” means
  - repartition funding for standard resources in a way that **some agreed target metric increases** - *not sure that metric has been established yet.*

# Why Metrics?

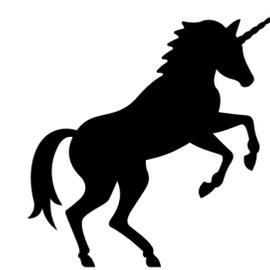
- **To manage services:** operational (facility metrics in doc)
  - Q: “all going as expected?”
    - yes -> boring
    - no -> what needs fixing?
  - usually: no effort-like quantities involved: no \$, Wh, t CO<sub>2</sub>
- Large majority of metrics in eg IT monitoring system
  - often not useful/designed for trending/optimisation
    - short term metrics, typically not even comparable across release boundaries
    - collected at rates/volume/retention period which make trending hard
- For AF deployment
  - Not yet needed at this point, as “working as expected” is still ill defined

# Using existing metrics: an example

- **CPU efficiency**  $:= t_{\text{cpu}} / t_{\text{wall}}$   
btw: a misnomer and should be called CPU utilisation
  - operational - measures *free CPU capacity*  
in units of the *current* average CPU speed
- If low, we usually conclude:  
bad for sites: who invested in buying it  
bad for experiments: who invested in obtaining it
- But, we could state equally:  
good the CO2 and energy efficiency (if workload still completes in time)  
confirms impact of new faster CPUs deployed
- Using operational metrics for trending comes with risks

# Why Metrics?

- To steer change towards a set goal or direction: evolution/optimisation
  - Q: “which (budget) combination for disk, network, memory and CPU investment maximizes throughput for the *currently* expected workload?”
    - Conf1 -> x evts/\$; Conf2 y evts/\$
  - Q: “does adding SSD increase throughput?”
    - Conf2 -> z evts/\$ -> better? worse? not significant within *error bars*?
- Our analysis of data center metrics is often far below the quality standards we apply elsewhere.
  - Definitely at CERN and I suspect also in some other places
  - Neither effort nor expertise is usually available in service teams
    - task is taken over by experiment and WLCG contributors



# So do we agree at least: Which Target Metric?

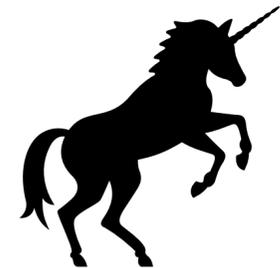
Maximize *science potential*

- for available budget
- what would I buy now? for the work we need.

# So do we agree at least: Which Target Metric?

Maximize *absolute throughput* (eg benchmark evt/s):

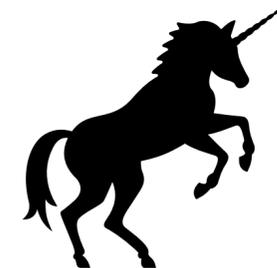
- for budget
  - what would I buy now? for the work we need.
- for existing kit
  - what would I run now? for the work that is also useful.



# So do we agree at least: Which Target Metric?

## Maximize *throughput* :

- for budget
  - what would I buy now? for the work we need.
- for existing kit
  - what would I run now? for the work that is also useful.
- for energy (and financial) budget.
  - slow-down to meet energy constraint?
- for CO2 (and financial) budget.
  - retain h/w for total CO2 rather than energy savings?



# Why consider (specialised) infrastructures for Analysis?

- Analysis
  - := anything that *is not* organised production (incl. sim) or pure SW development
  - too wide as a definition to reach any quantitative comparison
- Agreed benchmark examples are good, to get practical experience
  - but also do risk to *focus the discussion too quickly* on technicalities
    - eg the results may vary depending on who implements them
  - and may exclude some of the key strategy decisions
    - eg if main outcome is improved analyst efficiency (not CPU), how do we confirm/measure that with a benchmark?

# Why consider (specialised) infrastructures for Analysis?

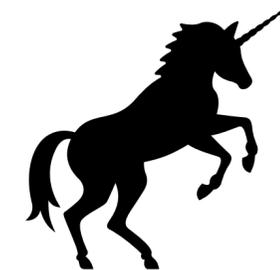
- **Analysis** := anything that *is not* organised production (incl. sim) or pure SW development
  - Let's stick for a moment with this broad definition
  - and be “success-oriented”: what could an AF exploit?
- Some (exploitable?) change wrt production:
  - **more (code) diversity**: bad news  
less likelihood and incentive to optimise
  - **“less” data diversity**: good news  
input *caching* or **train models** try to exploit this, also keeping **in-memory state**
    - metric to confirm: run time reduction
      - effort: cost of **cache** size
      - effort: additional latency of **train** concept
      - effort: loss of CPU utilization of “state **in-memory**”
  - **shorter time to next human decision**
    - target metric is maintained human attention/context (rather than CPU utilisation)

# Separate or On-the fly?

- Separate Infrastructure
  - additional cost -> which fraction of funding is adequate?
- Virtual infrastructure (eg created on the fly)
  - latency of creation/expansion/shrinking/destruction
    - creates additional CPU inefficiency
  - Is the minimum runtime to amortize and minimum time to create these on-the-fly AFs
    - compatible with observed “interactive” patterns?
- Unless we trade computing for brain efficiency?

# Is Single User Success already Success?

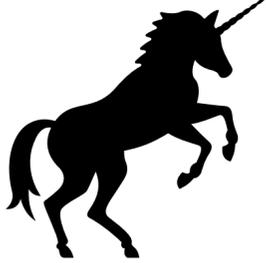
- is a necessary condition, but not a sufficient one for analysis community by-in
- Some parallel and scalable systems for analysis existed since decades (PIAF, PROOF)
  - but despite demonstrated improvements wrt to latency and throughput goals did not attract the majority of workload
  - only few user see benefit while others wait, in some cases without ETA
  - this has often not been not enough benefit for experiment resource managers to allocate dedicated analysis resources
- During the “honeymoon” period in the beginning most new services are attractive to early adopters - just because they are still empty



# Metrics for Resource Review interferes with (technical) optimisation

- Metrics used in higher level steering try to shield strategy discussion from (too much) technical detail
  - But what is too much?
- We see already adverse impact on attempts to improve efficiency
  - Eg storage & CPU has to remain full
  - Accounting one level of storage quality
- GPU and energy efficiency and CO2 goals will complicate the discussion further
- Some guidance concerning target metrics will be needed from the main stakeholders to disambiguate a quantitative discussion - *if that is desired.*

# Summary

- Sorry, I did not add anything new, but annotated the points with
-  in places where I believe missing data, evaluation work and eventually agreement is in the way of a quantitative approach to make **resource decisions**
- Even a more focussed evaluation could profit from an upfront target metric (eg which goals first)
  - how about focus just on “interactive analysis”
  - quantify current use, determine user impact with/without a specialised analysis infrastructure