



Challenges in computing and software for our big data

JENAS May 2022, Madrid

G. Merino, merino@pic.es



Ciemat Centro de Investigaciones Energéticas, Medioambientales y Tecnológicas



R Institut de Física d'Altes Energies



Outlook

Computing (sw & data analysis) as a 1st class citizen in scientific collab. and institutions

- Incentive to contribute to computing, as compared to detector constr./ops.
- s/w dev. as an intellectual activity - Career path for Research Software Engineers

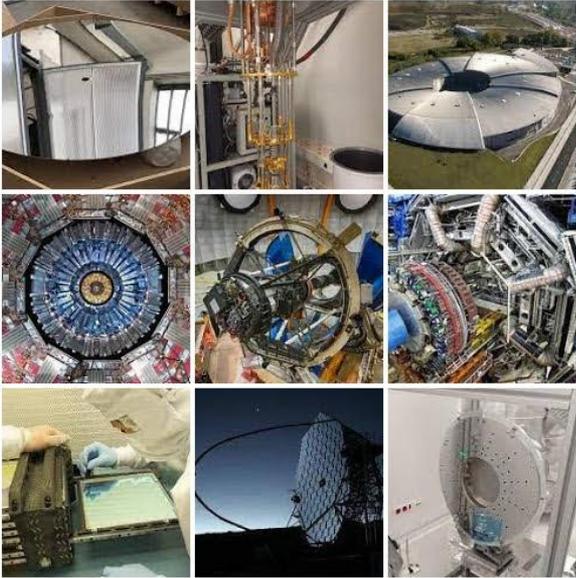
Extraordinarily dynamic environment

- Massive investments from industry - lots of innovation, short cycles
- Fierce competition for talent

Sustainability through Collaboration: ECFA - APPEC - NuPPEC, and beyond

Digital Infrastructure as integral part of Scientific Instrumentation

Scientific Instrumentation



digital infrastructure



```
patchUInt32Value, json: op
Op string json: path
Path string json: value
Value UInt32

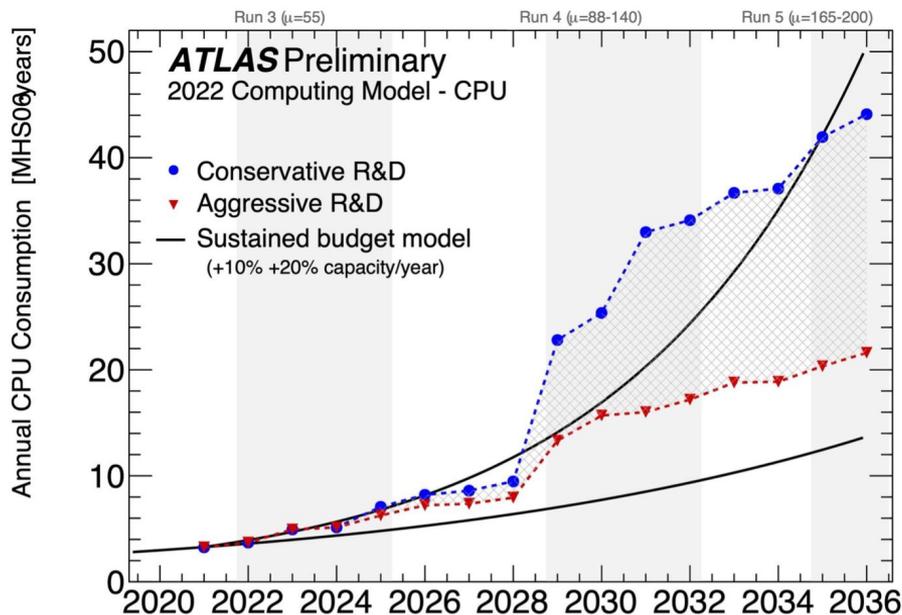
unc setVirtualServiceWeights(client dynamic.Interface, virtualServiceVR := schema.GroupVersionResource{
Group: "networking.k8s.io",
Version: "v1alpha3",
Resource: "virtualservices",
})
// weight the two routes = 50/50
patchPayload := make([]patchUInt32Value, 2)
patchPayload[0].Op = "replace"
patchPayload[0].Path = "/spec/http/0/route/0/weight"
patchPayload[0].Value = 50
patchPayload[1].Op = "replace"
patchPayload[1].Path = "/spec/http/0/route/1/weight"
patchPayload[1].Value = 50
// patchPayload

panel("default").Patch
```

analysis



HL-LHC future computing needs



[Source: CERN-LHCC-2022-005 report](#)

10x more data, 10x more complexity

Conservative scenario: current personpower and expertise maintained.

Aggressive scenario: increase personpower and expertise.

Black lines: +10-20% yearly capacity increase, assumed in flat-budget scenario.

R&D for reducing resource needs

Substantial part of current R&D efforts towards porting code to specialized hw (eg GPUs)

- Generators - up to 20% of CPU budget
- Simulation - fast simulations using ML techniques
- Reconstruction (80% of cpu budget for CMS) - offloading parts of the reconstruction to GPUs

Data management

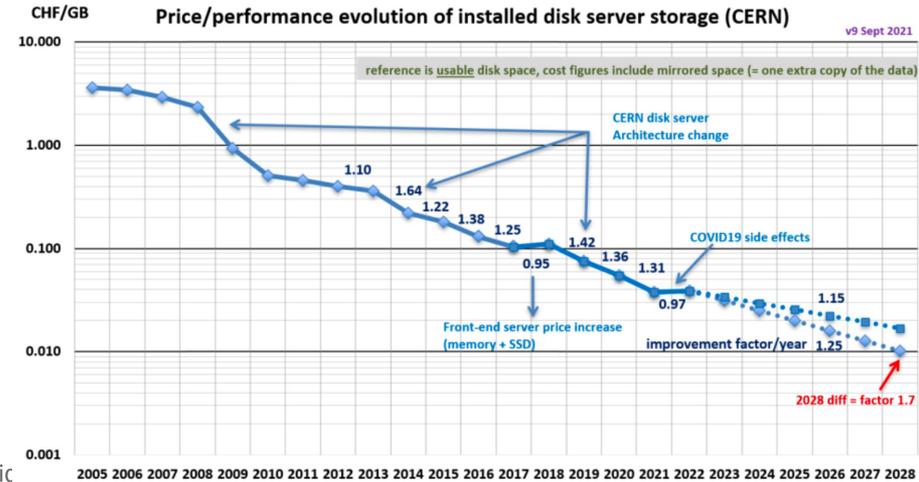
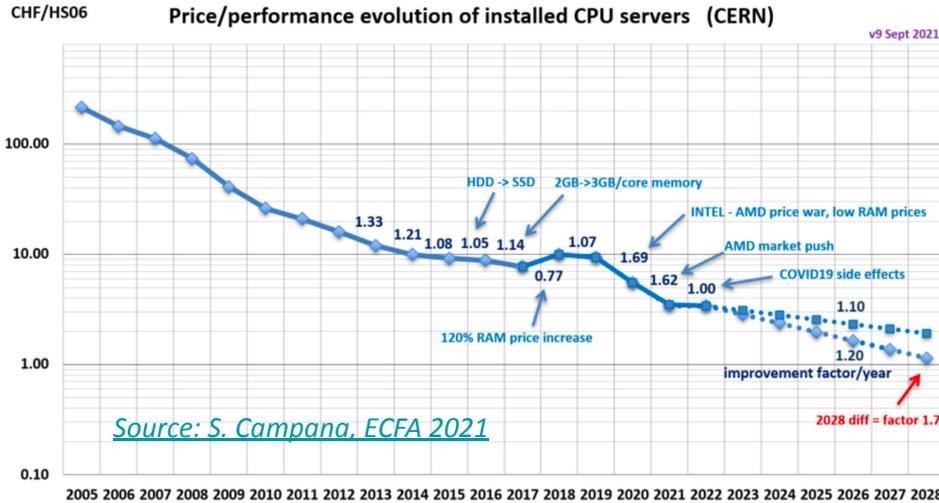
- Contain disk needs - reduced size analysis formats, data placement/caching
- New columnar data formats improving read throughput
- Minimize infrastructure cost by coordinating tiered storage
 - Main facilities getting ready to scale-up on tape capacity

Hardware costs

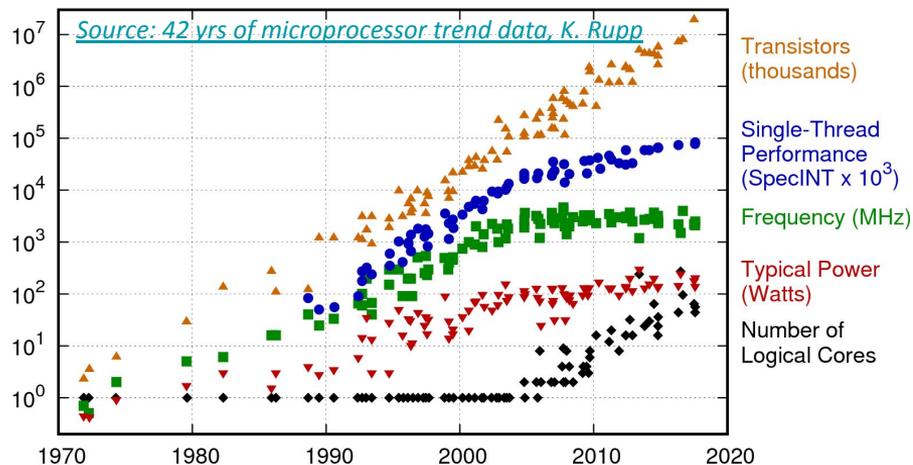
Subject to large fluctuations, even on the timescale of months.

For many components it is not technology that drives the prices, but market forces, lack of availability, and uncertainty over the future of certain technologies.

Future cost of computing is almost impossible to predict.



Semiconductor evolution



Moore's Law and Dennard scaling

- Exponential performance improvement for decades.

End of D.S, slowdown of M.L.

- Rising costs of semiconductor advances
- Ever-faster computers, more expensive

In an exponential world, CPU clock is now flat.

- Growth will happen by parallelization, distributed resources.
- The challenge is automation. Orchestration to manage complexity.

Federated Infrastructure

Technology and needs evolve exponentially, but available people's effort don't.

- The key is to find the way to SCALE.

The future is likely to be **distributed** - use resources wherever they are.

Multiple centers + integration with **HPC** and **Cloud**

- Some HPC will be usable, some won't
- Cloud
 - Need to understand TCO (experiences at large-scale already under way: [LSST](#), [ATLAS](#) ...)
 - Cloud advantage is **elasticity**
 - Keep control over the data - avoid vendor lock-in

High Performance Computing

Big investments in strategic agendas of several countries.

Most systems in top500 have **GPUs** - growing trend.

Need to find the best ways to use HPC for scientific experiments - stronger collaboration.

- Leverage big pool of GPU resources for ML analysis still, a small fraction of the overall need.
- Success in exploiting GPUs will play an important role.
- Also an opportunity for experiments to access opportunistic CPU and mitigate its shortage risk.



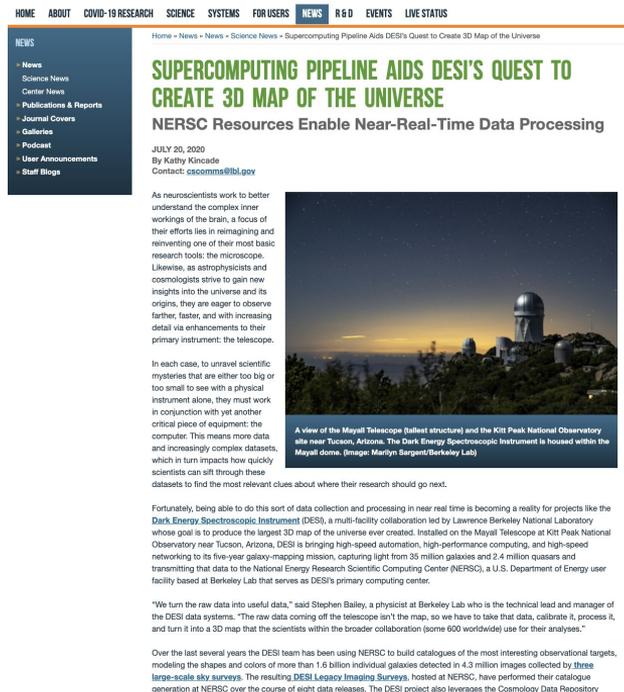
High Performance Computing

Barriers to use HPC for scientific experiments

- Access policies - quarterly calls for short projects
- WAN not designed for massive data in/out
- No outgoing connectivity from the worker nodes

Moreover: there will be need to access HPC in real-time

- DESI already doing this at NERSC, more to come ...



The screenshot shows a news article from NERSC. The main headline is "SUPERCOMPUTING PIPELINE AIDS DESI'S QUEST TO CREATE 3D MAP OF THE UNIVERSE". Below the headline is a sub-headline: "NERSC Resources Enable Near-Real-Time Data Processing". The article is dated "JULY 20, 2020" and written by "Kathy Kincaid". A navigation menu on the left lists categories like News, Center News, Publications & Reports, Journal Covers, Galleries, Podcast, User Announcements, and Staff Blogs. The main text discusses how neuroscientists and astrophysicists use high-performance computing to process data from the Dark Energy Spectroscopic Instrument (DESI) at the Kitt Peak National Observatory. It mentions that DESI is producing the largest 3D map of the universe ever created and is using high-speed automation and high-performance computing. A photograph of the Kitt Peak National Observatory is included, with a caption: "A view of the Mayall Telescope (tallest structure) and the Kitt Peak National Observatory site near Tucson, Arizona. The Dark Energy Spectroscopic Instrument is housed within the Mayall dome. (Image: Marilyn Sargent/Berkeley Lab)".

Energy Efficiency

Data centers cause a sizeable contribution to GHG emissions across sectors (and growing)

Scientific computing centres have to become more energy efficient

- Increasingly important role in HPC - Green500 list ranks HPC by performance/watt

Initiatives ongoing to improve accounting of carbon footprint of scientific computing

- Facilities: maximize cooling efficiency, use electricity from renewable sources
- Users & sw providers: improve code/algorithm efficiency

Labs are already analyzing these and publishing annual environmental reports

Include carbon footprint in the planning decisions and tradeoffs

- Consider power/carbon envelope, besides cost envelope

Shared Infrastructure

Many of the large scientific computing facilities are supporting multiple experiments.

Economies of scale for operations - Support Nx experiments with delta effort.

Reduce operational support load by: open source, open protocols & federated resources.

- open source portable building blocks, used and developed for broader communities
- widely used standards (gridftp → http, x509 → OIDC ...)
- middleware to integrate heterogeneous resources - workloads and dataflow orchestrators.

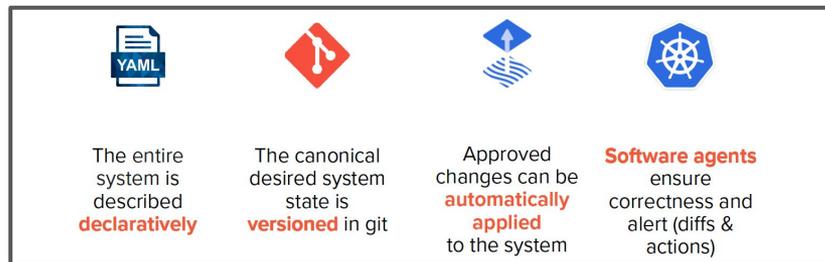
Cloud Native technologies

Methodology to deploy applications & infrastructure

- service-based architecture, declarative APIs, containers, immutable infrastructure ...

Emerging as de facto standard in industry

- Using sw development tools and processes to manage infrastructure
 - Containers + cont. orchestrators (k8s) - Declarative code for infrastructure
- opportunity to improve operational efficiency
 - sharing processes & code between facilities (more portability, less operational costs)



Principles of GitOps: <https://gitops-community.github.io/kit/>

Sustainability

Lots of science programs are spanning decades - **sustainability** of the computing infrastructure is as big a challenge as raw capacity.

Increased **collaboration** is one of the keys to sustainability.

2020 update of the European Strategy for Particle physics

- Highlights the importance of synergies between the neighbouring fields: APPEC - NuPECC
- Mutually benefit from scientific exchanges and technological cooperation in areas of common interest such as R&D for s/w and computing.



European Open Science Cloud - EOSC

An EC action to implement the Open Science vision and policies shared by member states

- A cloud for research data in Europe allowing seamless access to data and interoperable services
- Will federate existing e-infrastructures, national data centres and research infrastructures, allowing researchers and citizens to publish, access and re-use data



**EUROPEAN OPEN
SCIENCE CLOUD**

<https://eosc-portal.eu/>

ESCAPE

EU project in response to H2020 call to build the EOSC by connecting it to the **ESFRI RIs**.

ESCAPE: Astronomy and Particle Physics ESFRIs (2019 - 2022)

- Goal: Develop common “e-infrastructure” solutions for a wide range of particle physics & astronomy research facilities.
- Acknowledge that both communities will generate massive data and will push the state-of-the-art in data mgmt. and computing R&D.
- Common challenges for long-term data preservation, sustainability and open access to data.
- Exploit synergies between both communities expertise



ESCAPE
ESFRI
projects,
landmarks
and a few
more RIs

Radio



SKA

**JIVE-
VLBI**

Visible light



ELT



ESO



EST

Gamma rays



CTA

**Accelerator-based
Particle Physics**

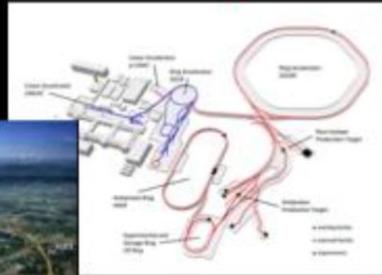


HL-LHC



CERN

**Accelerator-based
Nuclear Physics**



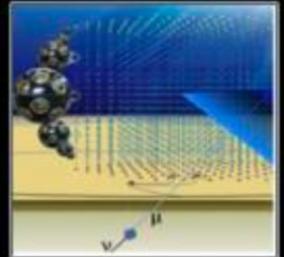
FAIR

**Gravitational
Waves**



EGO-VIRGO

**Cosmic-rays
Neutrinos**



KM3NeT

ESCAPE Work Programme

Prototype and deploy open data services that span the AP, HEP and NP domains.



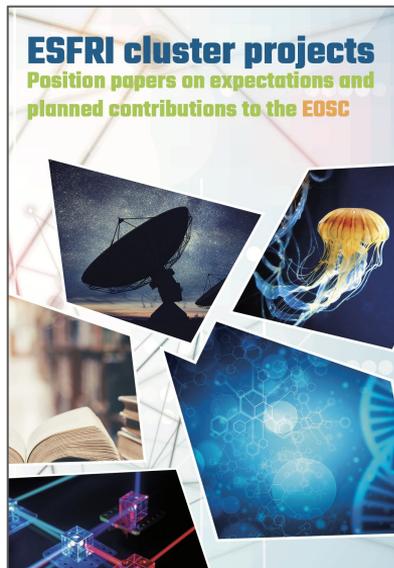
Integration demonstration through two **Test Science Projects** - Linked to two JENAA EoIs

- Dark Matter TSP - [iDMEu EoI](#)
- Extreme Universe TSP - [Gravitational Waves for fundamental physics EoI](#)

Broader synergies with other research clusters

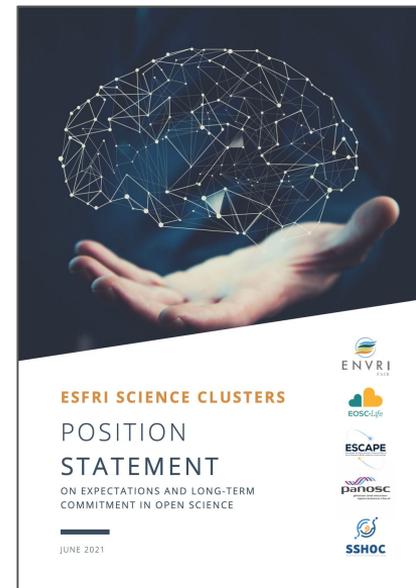


*Five thematic **Science Clusters**
funded under the same H2020
call (80% of ESFRI RIs)*



Dec 2019

[doi:10.5281/zenodo.3675081.svg](https://doi.org/10.5281/zenodo.3675081.svg)



June 2021

[doi:10.5281/zenodo.4892245.svg](https://doi.org/10.5281/zenodo.4892245.svg)

Summary - challenges

Extremely dynamic environment - almost impossible to make long-term predictions.

- Need a framework and processes to manage change.
- Metrics for continuous assessment - identify when changes are needed - iterative improvement.

Massive innovation from industry. Hard for research Labs to compete.

- Our big value is the capacity to generate high-quality high-value data.
- We need to engage with industry to integrate and exploit new technologies.

New boundary conditions for building experiment computing models. Several trade-offs:

- Peak performance, throughput, cost, ease of ops, ease of use, risks, carbon footprint
- Different optimal points for different communities - support hybrid environments

Summary - challenges

Engage with HPC to enable effective resource usage by scientific RIs

- Increase researchers ability to provision services at HPC centers
- Experiments to develop tools to provision and orchestrate infrastructure at HPCs
- Enable production use of novel architectures at HPC centres
- Technology to support real-time experimental workloads

Ability to scale-up is key - Federated distributed infrastructures.

- Effort needed to tackle the issues and study the tradeoffs
 - HPC: access policies, heterogeneous architectures
 - Commercial Clouds: cost predictability, data sovereignty

Summary - challenges

Long-term data preservation is challenging for our projects - multi-decade timescale

- Need organizations with sustainable funding model, trusted by the community
- With technical expertise on data management

Personpower to support scientific computing is not increasing.

- Need to increase automation to be more efficient to scale-up capabilities with constant effort.

Scientists' time is the most precious resource. We need to reduce the data management burden on them so they can focus on the science.

thank you