Highlights from Track 7 - Facilities, Clouds & Containers

24th International Conference on Computing in High Energy & Nuclear Physics

4-8 November 2019, Adelaide, Australia

Oksana Shadura (Nebraska), Sang Un Ahn (KISTI), Christoph Wissing (DESY)

Track Statistics

- 48 oral presentations
  - 8 sessions – 2 of them in parallel on Thursday afternoon
- ~22 posters
- Typically 30-40 persons in the room
  - Track rather well attended
- Good atmosphere
  - Often lively discussions – cut to stay (almost) on schedule

Pretty difficult to address all content in ~12min

Personal selection – it is biased!

Topical Sessions

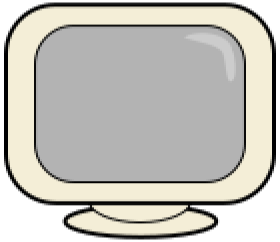

Monitoring &
Benchmarking

Cloud Computing

Trends &
new Approaches

Infrastructure

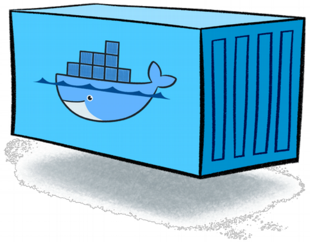

Containers

Non LHC Experiments

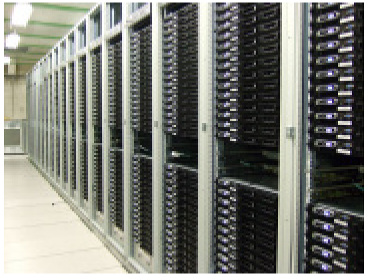

Opportunistic
Resources

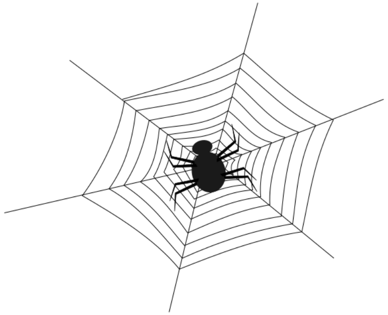

Network Technology

Monitoring & Benchmarking

- HEPSpec06 benchmark does not correlate with real HEP workloads of today
  - Real reference HEP tasks can serve as new benchmark
  - Recent container technology enables practical packaging and distribution
- Machine Learning techniques allow novel approaches for Monitoring
  - Anomaly detection via unsupervised networks
  - Inputs either metrics or logs
- SAND network monitoring project
  - Gather metrics from various sources, e.g. perfSONAR, FTS, network workflows
  - Employ modern analytics tools
- Common platform for monitoring: Elastic ecosystem
  - Used for MONiT at CERN, but also elsewhere

beats
logstash
elasticsearch
Grafana
kibana
ElastAlert

Cloud Computing

- Commercial and research Clouds remain important
- European Open Science Cloud (EOSC)
  - DODAS (Dynamic On Demand Analysis Service) using EGI FedCloud
  - CS3-APIs: Vendor independent layer for sync-and-share
  - Pre-commercial cloud resource procurement projects
- Projects with commercial partners pose new challenges
  - Many of them are not technical
- Commercial clouds: Data import, export and access
  - Providers appear to have very strong internal network fabric and also good WAN connectivity
  - Import is always for free
  - Hosting, accessing and exporting data comes with significant costs

New Approaches

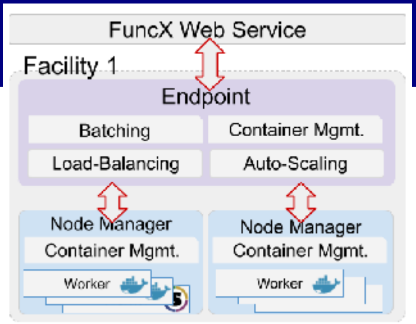

- Function as a service – funcX
  - Provide access to functionality/calculation via webservice
  - Prototype with HEP analysis example shows nice scaling
- Reduce efforts to run/provide site services
  - SIMPLE
    - Easy deployment of middleware components from minimal config files
    - Builds on Docker containers and Puppet
  - SLATE (Services Layer at the Edge)
    - Docker, Kubernetes, and Helm to package and deploy services
    - Central server to mediate requests being sent to participating edge Kubernetes clusters
- Noted (Network Optimized Transfer of Experimental Data)
  - On demand configuration of redundant network links
  - Employs Rucio, FTS and involved SDNs

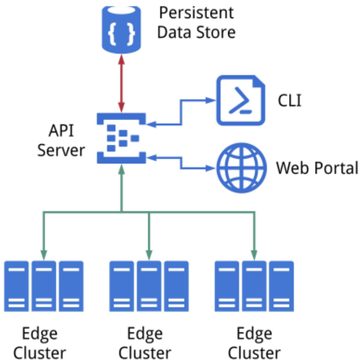

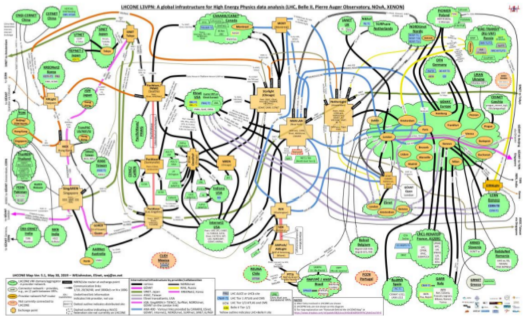

Fabiola

Trends

- Jupyter
  - Increasing interest on the user side
    - Interactive, fast learning curve
    - Easy development and sharing
  - Jupyterhub enables access to diverse HTP and HPC resources
- Overview by Hepix TechWatch working group
  - Hyperscales (Google, Amazon) drive the market
  - x86 market: AMD is back
  - Magnetic disk: Market is shrinking
  - Tape: Risks – essentially one company left for R&D
  - Ethernet evolving very fast
    - Pace of change exceeding IEEE standards process

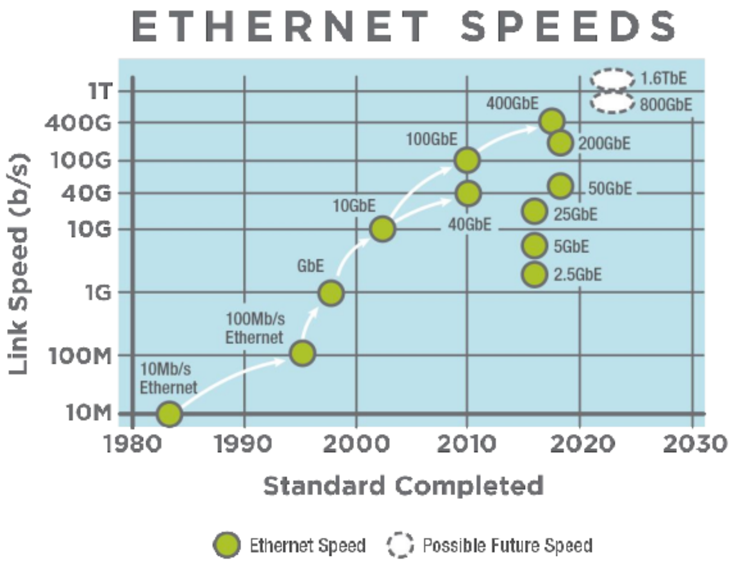

Infrastructure

- Software infrastructure
  - WPAD (Web Proxy Auto Discovery)
    - Mechanism to easily auto setup Squid configuration for non-static resources
  - Modernization of the ATLAS online web service environment
- Site and hardware infrastructure
  - Everlasting struggle with space and power
    - CERN: After stopping extension at Wigner leverages containers from LHCb
    - BNL: Setup of a new data center
  - Challenges for locations with several sciences on campus
    - Different demands from different communities

slurm workload manager

- HPC Cluster
- Users: Photon/Accelerator Scientists

HTCondor High Throughput Computing

- HTC Cluster
- Users: HEP Scientists

- Difficult to understand and use one system well
- Do you want to enforce learning a second system?

Containers

- Kubernetes, kubernetes, kubernetes, …
  - Intensive deployment of containers based on Kubernetes engine
  - Job scheduling without batch system, simply relying on Kubernetes
    -> effective and promising, aspect to simplifying site operations, no CEs...
  - Registry solutions to deploy container images is one of concerns
- ScienceBox
  - Complete solution for scientific set of services
    from highly-scalable storage solutions (EOS)
    to user-friendly application, Jupyter notebook
  - All nicely packaged in containers
- Container technology facilitates use of various resources
  - HPC, HTC, Grid resources, etc.
- Moving CERN batch from Openstack VMs to Kubernetes
  - First benchmarks indicate 5% performance gain
- All major experiments use containers in production

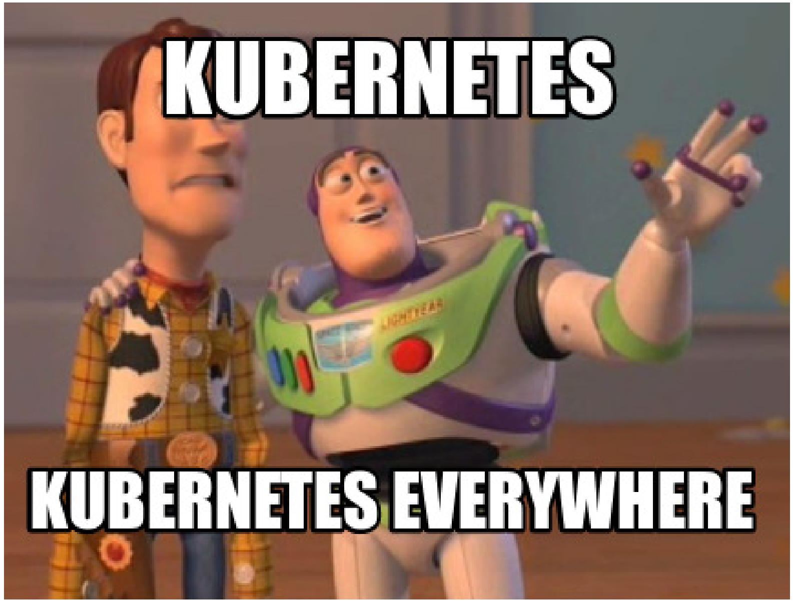

Non-LHC Experiments

- Leveraging resource access via Cloud interfaces
  - LHAASO distributed computing
    - Large High Altitude Air Shower Observatory
    - cloudSchedular for multi Cloud access
  - IceCube real-time processing in AWS

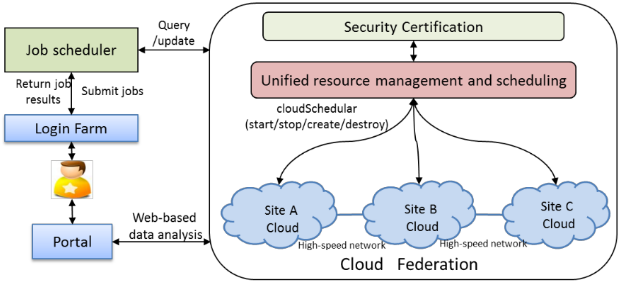

- Non-LHC Communities building on LHC tools (or at least consider them)
  - RUCIO, FTS, HTCondor, CVMFS – very common for LHC groups
  - VIRGO and Gravitational Waves computing in Europe
  - AENEAS: Designing a Federated Regional Centre for SKA Computing
    - Minimum storage: **750PB** over 10 years
    - Several Exabytes after 15 years
- Smaller experiments (e.g. Project 8 and ADMX) move to tools like CI & K8s
- Online computing for new generation photon science experiments
  - Diverse requirements from various groups
  - Modern Cloud stack allow to build virtual computing centers
    - Isolate one group from another

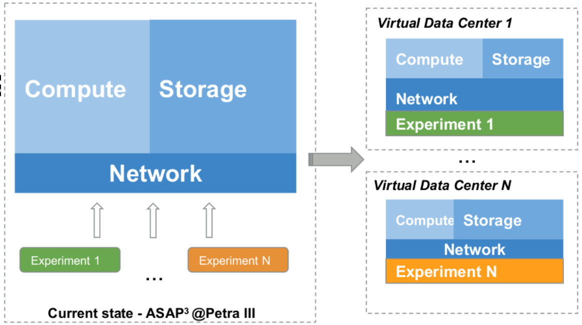

Opportunistic Resources

- Adding academic or commercial Clouds and HPCs very common
- Different tools presented (there are more!)
  - COBalD/TARDIS from KIT
  - Google Cloud Platform Condor Pool Manager (GCPM) from Tokyo
  - GlueX (JLAB) utilizes resources at NERSC, PSU and OSG
- PROMINCE tool used for the fusion community (ITER)
  - Very flexible tool to access a spectrum of resources
  - “Users don’t need to worry about provisioning infrastructure on clouds
    • Or even know what a cloud is”
- Using HLT farms for offline processing
  - Common approach for LHC experiments
  - Recent improvements by ATLAS
    - Shorten the startup of offline environment
    - Persistent CVMFS cache

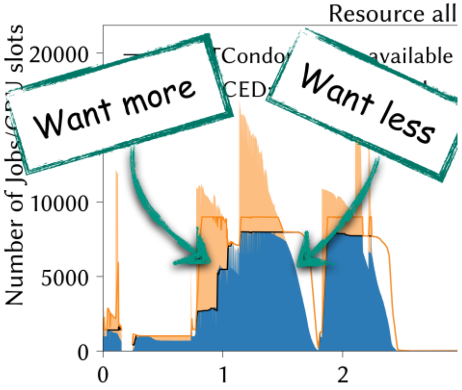

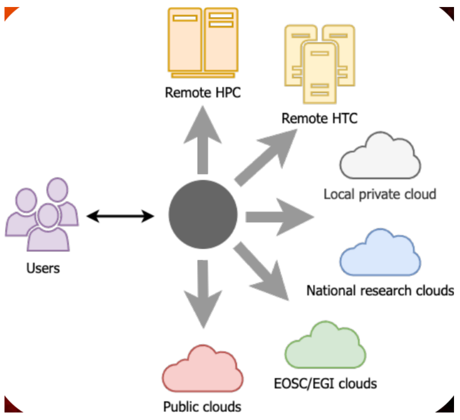

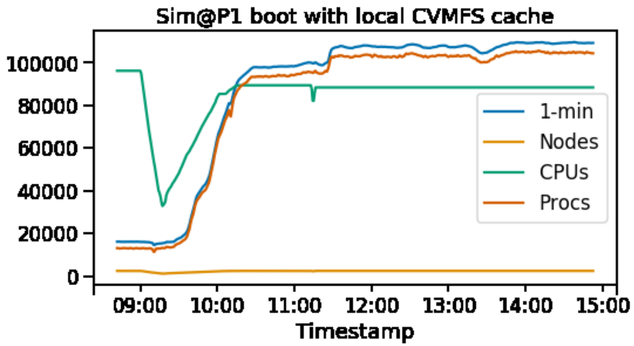

Networking

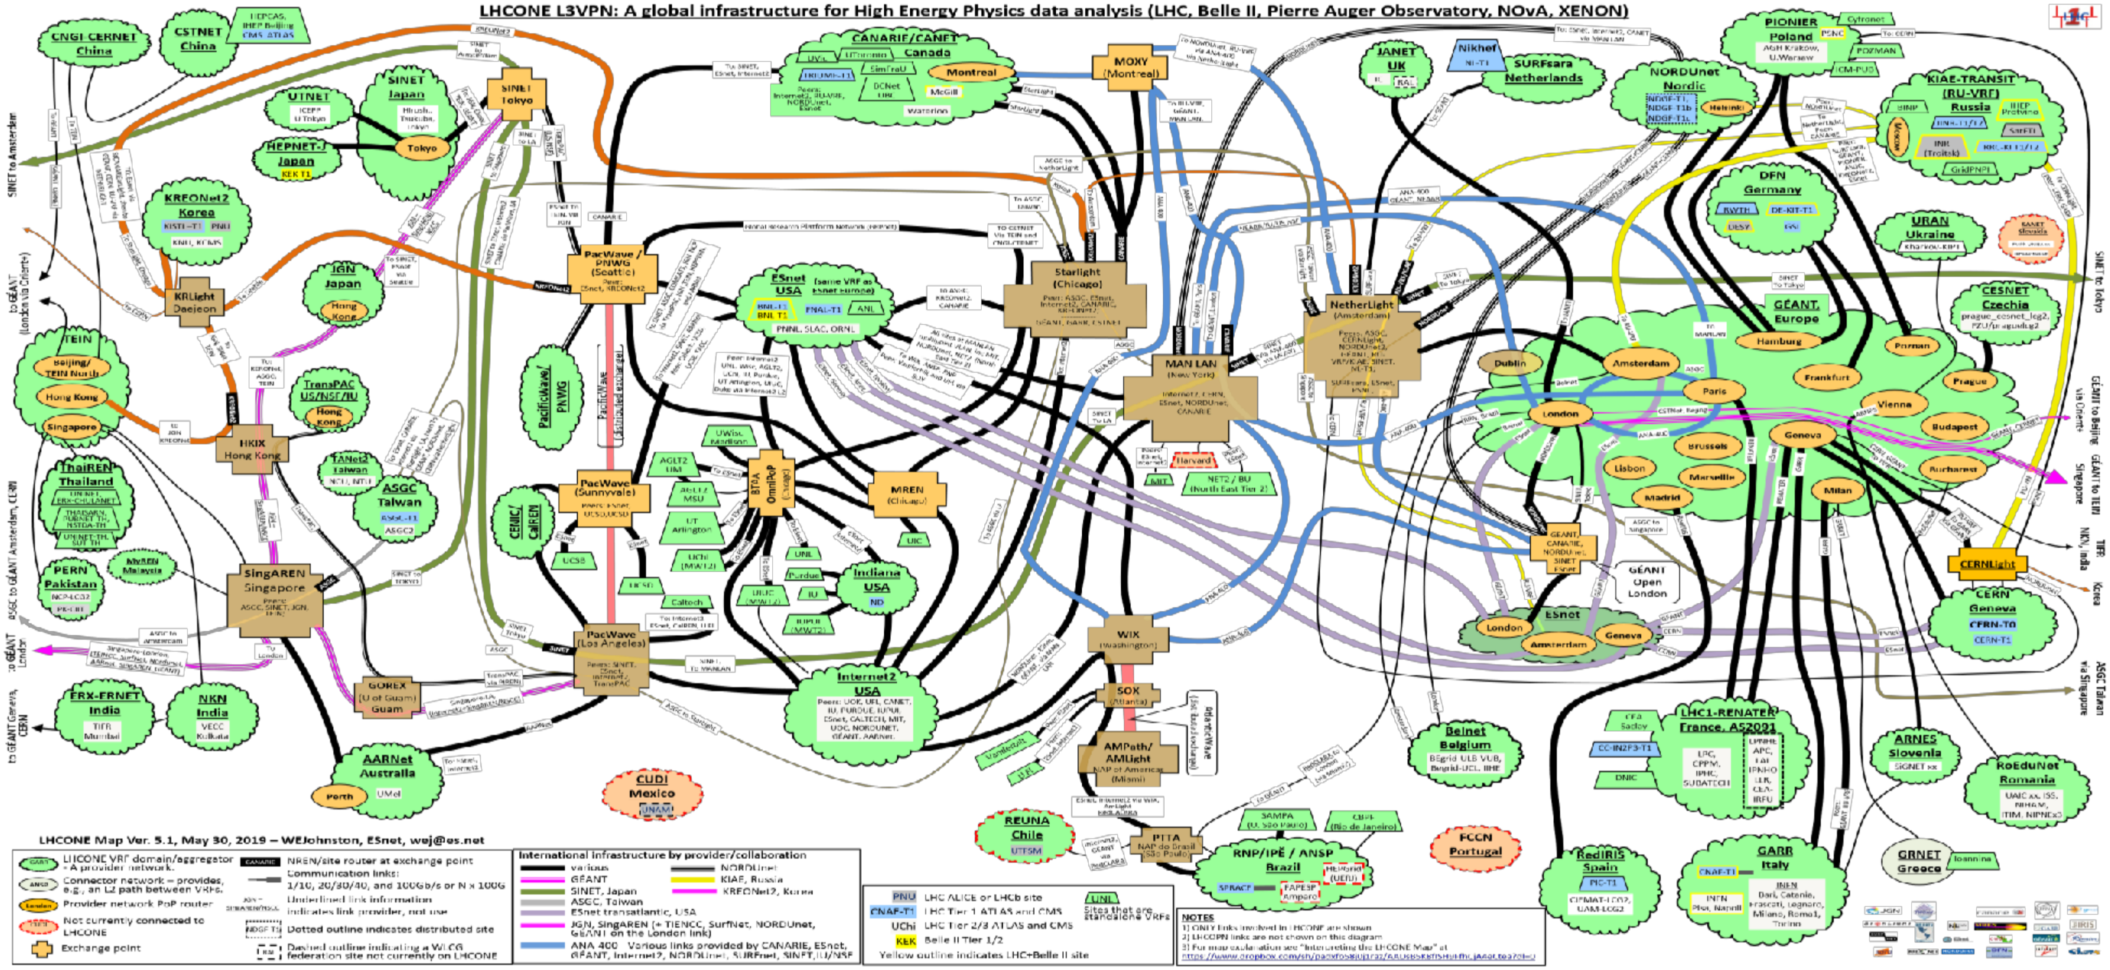

LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger Observatory, NOvA, XENON)
CNGI-CERNET China
CSTNET China
SINET Japan
SINET Tokyo
UTNET
HEPNET-J Japan
KEK T1
KREONet2 Korea
KRLight Daejeon
JGN Japan
TEIN
Beijing/TEIN North
Hong Kong
Singapore
HKIX Hong Kong
TransPAC US/NSF/IU
ThaiREN Thailand
TANet2 Taiwan
ASGC Taiwan
PERN Pakistan
MyREN Malaysia
SingAREN Singapore
ERX-ERNET India
NKN India
GOREX (U of Guam) Guam
AARNet Australia
Perth
CANARIE/CANET Canada
Montreal
MOXY (Montreal)
PacWave / PNWG (Seattle)
PacificWave / PNWG
ESnet USA
Starlight (Chicago)
MAN LAN (New York)
NetherLight (Amsterdam)
PacWave (Sunnyvale)
CENIC/CalREN
PacWave (Los Angeles)
OmniPoP
MREN (Chicago)
Indiana USA
Internet2 USA
WIX (Washington)
SOX (Atlanta)
AMPath/AMLight (Miami)
CUDI Mexico
REUNA Chile
PTTA NAP do Brasil (São Paulo)
RNP/IPÊ / ANSP Brazil
FCCN Portugal
JANET UK
Nikhef
SURFsara Netherlands
NORDUnet Nordic
PIONIER Poland
KIAE-TRANSIT (RU-VRF) Russia
DFN Germany
URAN Ukraine
GÉANT Europe
CESNET Czechia
Dublin
Amsterdam
Hamburg
Frankfurt
Poznan
Prague
Paris
London
Vienna
Budapest
Geneva
Brussels
Lisbon
Marseille
Madrid
Milan
Bucharest
CERNLight
CERN Geneva
GÉANT Open London
ESnet
Belnet Belgium
LHC1-RENATER France, AS2091
ARNES Slovenia
RoEduNet Romania
RedIRIS Spain
GARR Italy
GRNET Greece
LHCONE Map Ver. 5.1, May 30, 2019 – WEJohnston, ESnet, wej@es.net

Networking

- Network session was a rather technical one
- We can be grateful to our teams of experts to address issues like
  - Enable us to connect or mobile devices anywhere (including tunnels 100m deep in the rocks)
  - Identifying misrouted traffic in a complex network like LHC1
  - Getting familiar with ISC Kea as replacement for DHCP
  - Understanding the network traffic in and towards our sites
  - Preparing for IPv6 (only)
    - FTS close to 60% on IPv6
    - Fractions of WLCG storages with dual stack increasing (T1: 96%, T2: 73%)
- Belle II network build very much on existing LHCOPN and LHC1
  - All RAW data centers on LHC1 and 80% of storage and 80% of compute power reachable via LHC1
  - Side note: Also SKA networking plans very much to include existing structures (non-LHC session)
- Preparations for the networking in the HL-LHC domain
  - “Programmable networks”
  - Network aware data transfer services/brokers

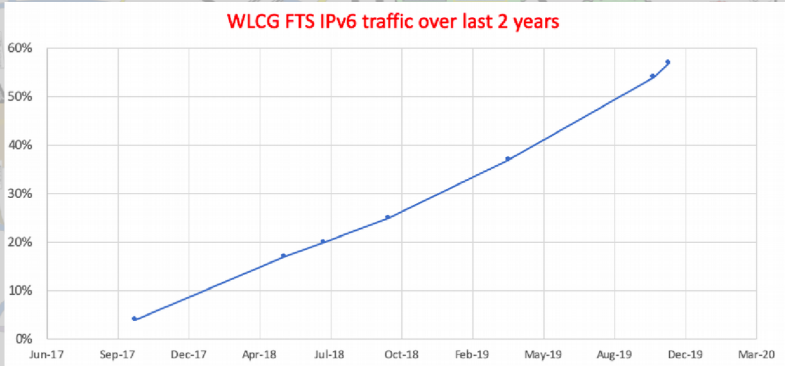

And many more nice posters … complementing the talks

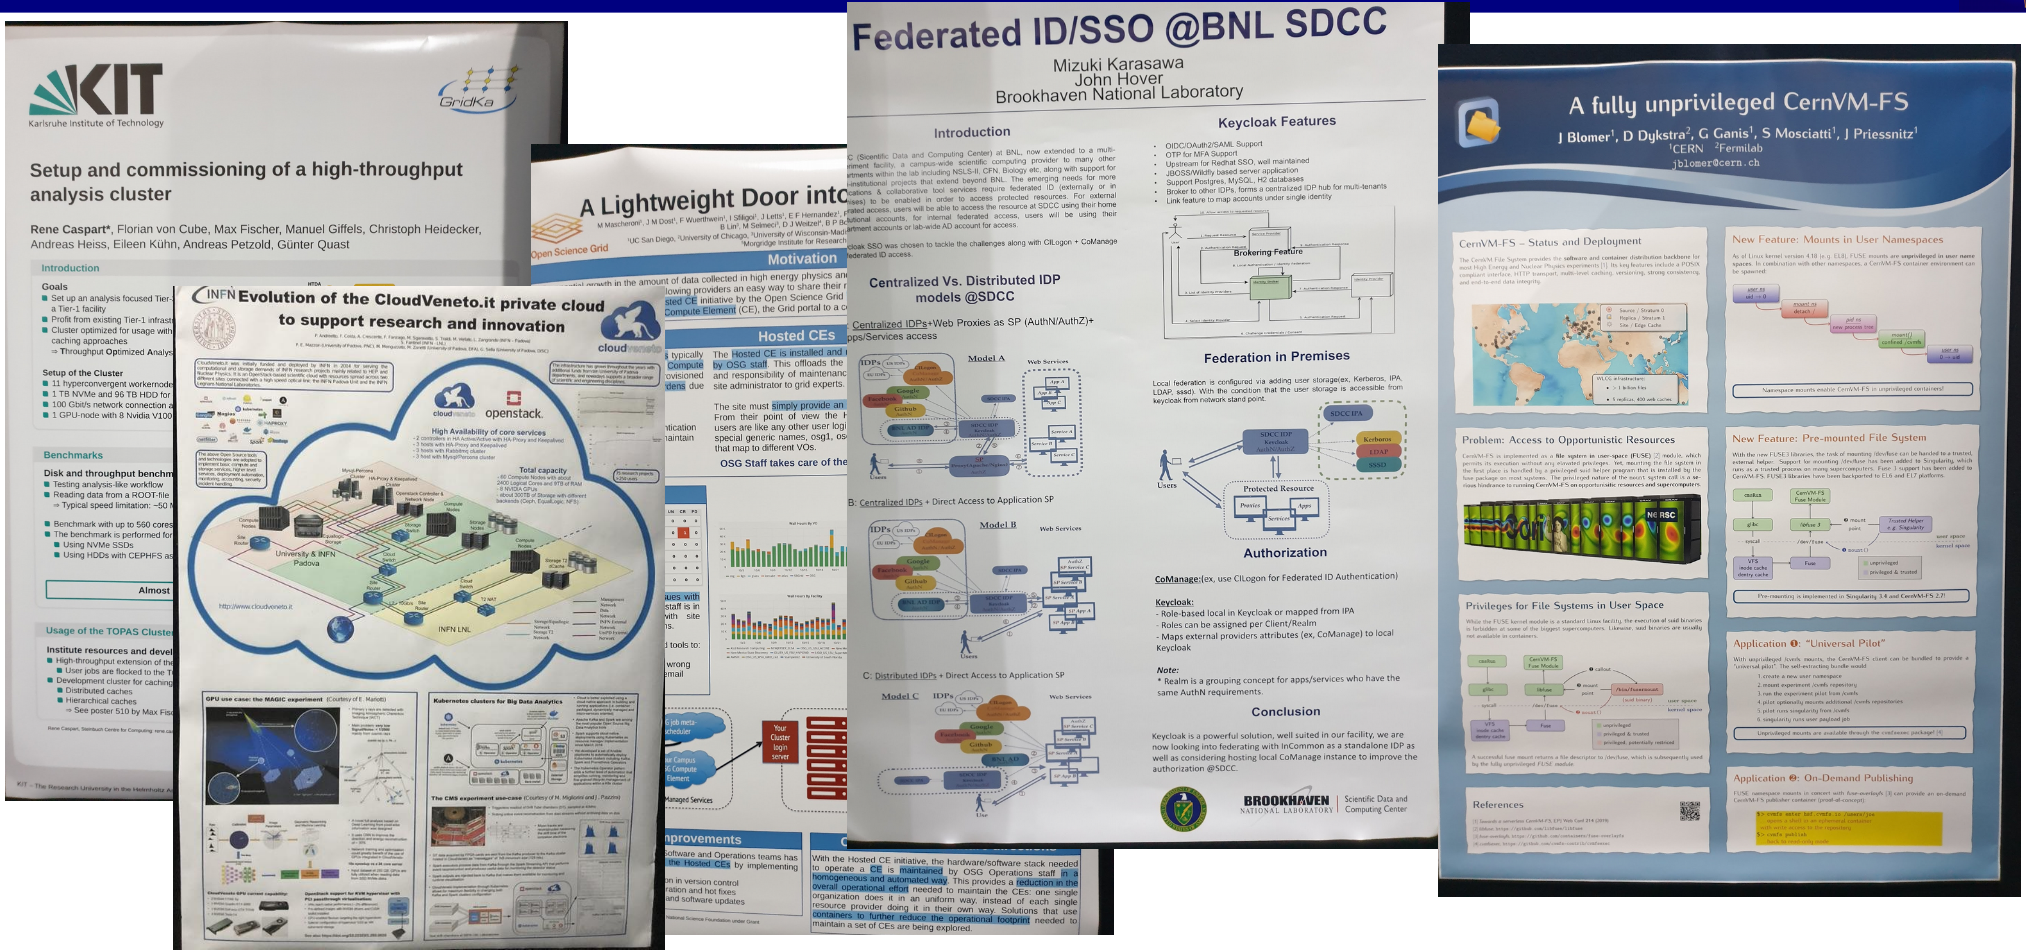

Final Remarks

- Our community migrates away from home grown tools towards industry standard products
  - Modern container solutions
  - Cloud architecture
  - HEP community is power user, sometimes contributor but almost never the main developer
    → Poses some challenges in collaboration, typically less of technical nature
- Many sites are no longer HEP only
  - Synergies with neighboring sciences: SKA and LIGO/VIRGO using HEP data management tools
  - Challenges: Dealing efficiently with a diverse user community

Closing the loop to the first plenary talk by H. Schellman:
Do not pay too much attention to the summary speaker (not a 1:1 quote)

Final Remarks

- Our community migrates away from home grown tools towards industry standard products
  - Modern container solutions
  - Cloud architecture
  - HEP community is power user, sometimes contributor but almost never the main developer
    → Poses some challenges in collaboration, typically less of technical nature
- Many sites are no longer HEP only
  - Synergies with neighboring sciences: SKA and LIGO/VIRGO using HEP data management tools
  - Challenges: Dealing efficiently with a diverse user community

Closing the loop to the first plenary talk by H. Schellman:
Do not pay too much attention to the summary speaker (not a 1:1 quote)

...so instead of listing...

Final Remarks

...just look at us saying:

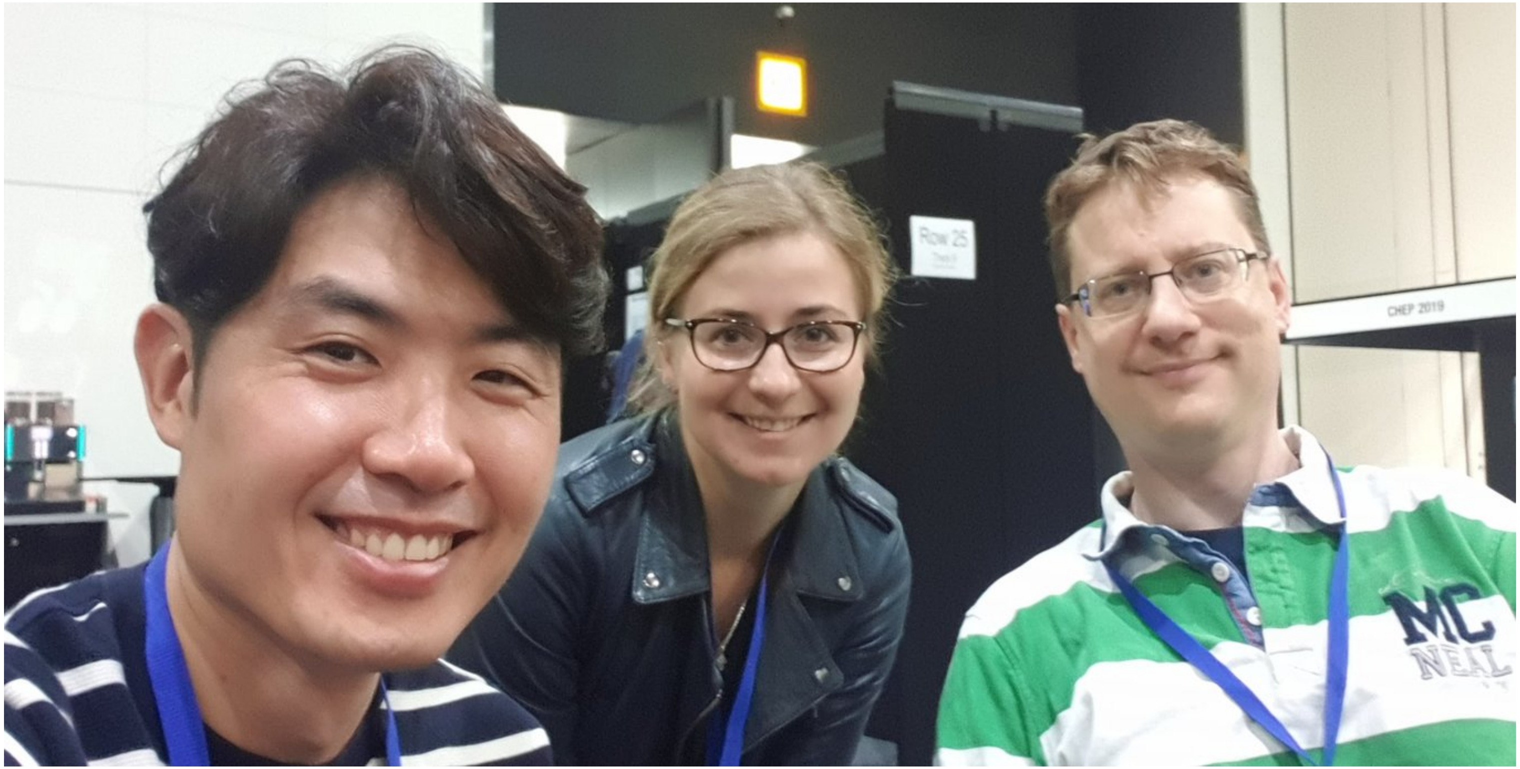

THANKS a lot to all contributors to our track, the local organizers and the conference chair team