

# CHEP 2018

## TRACK#8 - NETWORKS AND FACILITIES

### SUMMARY TALK


*Sang Un Ahn (KISTI), Josep Flix (PIC/CIEMAT), Oksana Shadura (University of Nebraska Lincoln), Wei Yang (SLAC)*



**23RD INTERNATIONAL CONFERENCE ON**  
**COMPUTING IN HIGH ENERGY AND NUCLEAR PHYSICS**

**9-13 July 2018**  
**National Palace of Culture**  
**Sofia, Bulgaria**

# Overview of Track 8: Networks and Facilities



# Disclaimer slide

- **Difficult** to credit all Track 8 contributions in detail in a 20' summary talk
- We highlight the most important points for each of the Oral sessions and Poster sessions
  - It might be incomplete in some cases, **sorry for this!** [in particular for Posters]
  - Look into the talks or chat with the authors if you are interested in specific issues

**I am extremely grateful to all my convener colleagues for helping in this talk preparation**

# Track 8: selected plenary talk

## *Disaster recovery of the INFN Tier-1 data center: lesson learned*



Luca dell'Agnello is Director of Technology Research at INFN and currently coordinates the INFN Tier-1 data center located at INFN CNAF in Bologna, and he also represents INFN in the WLCG Management Board.

He graduated in Physics at the University of Firenze in 1992 and worked on computing since then.


He joined INFN in 1996 working, at the beginning on the GARR project (Italian NREN) and then in the European Projects Datagrid and Datacloud.

Since 2003 he has been involved in the start-up of the INFN Tier-1.

He was member of the GARR technical committee.



# T8: Orais



# T8: S2 Session - Networking [Mon. Aft. / 8 talks / ~40 attendees]

- Improving WLCG Networks Through Monitoring and Analytics - Edoardo Martelli
- IPv6 in production: its deployment and usage in WLCG - David Kelsey
- Using Kytos SDN platform to enhance international big data transfers - Beraldo Costa Leal
- Integration and evaluation of QUIC and TCP-BBR in long-haul WLCG data transfers - Raul Cardoso Lopes
- Integrating Networking into ATLAS - Shawn McKee
- Entry Stage for the CBM First-level Event Selector - Dirk Hutter
- Netbench - large-scale network device testing with real-life traffic patterns - Stefan Stancu
- Long-term experiences in keeping balance between safety and usability in research activities in KEK - Tadashi Murakami



Improving WLCG Networks Through Monitoring and Analytics  
Edoardo Martelli, Marian Babik, Shawn McKee  
on behalf of WLCG Network Throughput WG





IPv6 in production: its deployment and usage in WLCG  
David Kelsey  
STFC UK Research and Innovation  
At CHEP18, Sofia, 9 July 2018



Using Kytos SDN platform to enhance international big data transfers  
Beraldo Costa Leal  
on behalf of Kytos WG  
July 9, 2018

Integration and evaluation of QUIC and TCP-BBR in long-haul data transfers  
Raul H. C. Lopes<sup>1</sup>  
Tim Chien<sup>2</sup>  
Oscar Hall<sup>3</sup>  
<sup>1</sup>Law and Digital Analytics London, UK  
<sup>2</sup>Law and Internet College, UK  
<sup>3</sup>Law  
July 9, 2018



Integrating Networking into ATLAS  
Shawn McKee(speaker), Marian Babik,  
Hironori Ito, Mario Lassnig, Ilija Vukotic  
for the ATLAS Collaboration  
Track 8, Networks and Facilities  
CHEP 2018, Sofia, Bulgaria, July 9, 2018



Netbench  
Testing network devices with real-life traffic patterns  
CHEP 2018  
stefan.stancu@cern.ch

Long-term experiences in keeping balance between safety and usability in research activities in KEK  
Tadashi Murakami, Fukuoka Yusaku, Ryoichi Baba, Tetsu Nakamura, Kiyoharu Hoshino, Saki Y. Sasaki, Masao Kitaguchi, and Toshiaki Kamei  
High Energy Accelerator Research Organization (KEK)  
Computing Research Center, Security Group  
Speaker: Tadashi Murakami,  
on behalf of Security Group



# T8: S2 Session - Highlights

- OSG and WLCG successfully operate a **network performance monitoring** platform that help sites and experiments understand and fix issues
- **IPv6 deployment** is being successful throughout WLCG and a fraction of data transfers goes over IPv6, which shows better performance and efficiency than IPv4
- An open source **SDN platform** (Kytos - SPRACE) and its gfal2 plugin can be used to create programmable networks (with QoS) before transfer starts
- Integration and evaluation of **QUIC (preliminary) and TCP-BBR** in long-haul data transfers - TCP BBR shows good performance values

# T8: S2 Session - Highlights 2

- Various activities underway to **integrate the networking information into ATLAS tools** and operations
- CBM experiment at FAIR will transmit data to the Green IT Cube DC (free streaming front-end electronics - **10 TBit/s over 700 m distance**). Testbed with Infiniband EDR, which is challenging for such distance
- **Netbench** is designed for large-scale testing of network device by means of a cost-effective framework
- **KEK Security team** presented their long-term experiences and difficult tradeoffs in keeping balance between safety and usability in research activities in KEK

# T8: S4 Session - Monitoring++ [Tue. Aft. / 8 talks / -50 attendees]

- Next Generation of HEP CPU Benchmarks - Domenico Giordano
- Sharing server nodes for storage and compute - David Smith
- MONIT: Monitoring the CERN Data Centres and the WLCG Infrastructure - Alberto Aimar
- Challenges, solutions and lessons learnt in 7 years of Service Management at CERN - David Martin Clavo
- Notifications workflows using the CERN IT central messaging infrastructure - Zhechka Toteva
- Deploying a “push” model Prometheus - Hristo Umaru Mohamed
- Detection of erratic behavior in load balanced clusters of servers using a machine learning based method - Martin Adam
- Monitoring System of the AMS Science Operation Centre - Baosong Shan



# T8: S4 Session - Highlights

- HS06 is not scaling with some of LHC experiments workload and demand on faster benchmarks emerged. **SPEC CPU2017** has been investigated and shown correlated results with HS06. A suite of HEP workloads could be an alternative *[in progress]*
- Proof of concepts presented to **scavenging CPU resources in CERN EOS storage** system without impact on I/O operations
- MONIT: A comprehensive **new monitoring system** for CERN and experiment services that has been setup using open source technologies
- Efforts on managing and simplify **service processes** at CERN have been conducted

# T8: S4 Session - Highlights 2

- **CERNMegabus** is designed to propagate state of services promptly to the provisioning system and is being used for storage, DNS load-balancing, and power management
- Telegraf adopted to Kafka based MQ system and enabled to push metrics to **Prometheus**, a Time-series DB
- A preliminary work is ongoing to **analyze metrics from CERN MONIT embracing ML methods**
- **Monitoring for AMS experiment** relies on CERN MONIT and a future work planned to enrich metrics and visualization

# T8: S6 Session - Facilities [Thu. Mor. / 6 talks / -60 attendees]

- INFN Tier-1: a distributed site (Luca dell'Agnello)
- Dynamic Integration and Management of Opportunistic Resources for HEP (Matthias Jochen Schnepf)
- A Feasibility Study about Integrating HTCondor Cluster Workload with SLURM Cluster Workload (Ran Du)
- LHCb High Level Trigger in remote IT datacentre (Edoardo Martelli)
- A prototype for the ALICE Analysis Facility at GSI (Kilian Schwarz)
- CMS Site in a Box: Deploying a Tier-3 Site using Local Resources and Central Services via a Centrally Managed Server (Carl Lundstedt)






# T8: S6 Session - Highlights

- **Elastic integration of external CPU resources** in a transparent way for users @**CNAF**. Cloud&HPC (Aruba, Azure, HNSciCloud, Bari-RECAS, CINECA)
- **KIT** presented the **integration of opportunistic CPU resources** (non-HEP) on-demand to absorb peaks, managed by ROCED (KIT development). Cloud (Exoscale, T-Systems, HNSciCloud) and HPC (Freiburg, KIT)
- **IHEP** experience in migrating **jobs from HTCondor cluster to SLURM cluster**, by means of Overlap, Flocking, and HTCondor-C (preferred)


# T8: S6 Session - Highlights 2

- **LHCb High Level Trigger in remote IT DC** (Meyrin, CERN, atm) for Run3: 7.2 Km fiber distance - 40 Tbps. Testbeds and initial setup deployed, in which a fraction of the online is now done @CERN (since 10th July 2018)
- Prototype and first tests of the **ALICE Analysis Facility at GSI**, with the goal to process ~5 PB of data within ½ day timescale (~115 GB/s)
- **Site in a Box**: Deploying Centrally Managed Servers to easily exploit Tier-3 CMS resources elsewhere. US examples and more sites to be added soon


# T8: Posters




# T8: Posters (the actors)



 Pier Paolo Ricci (INFN CNAF)




 David Smith (CERN)




 Zhechka Toteva (CERN)




# T8: Posters (the actors)



 Tadashi Murakami (KEK)

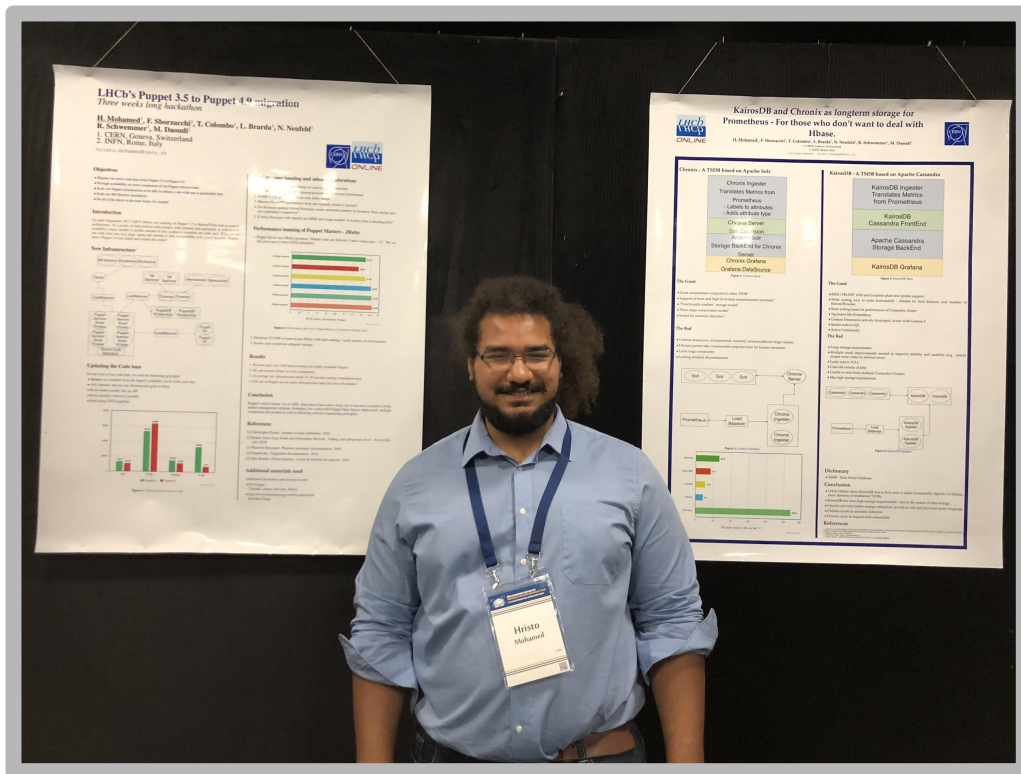


 Dr. Geonmo Ryu (Korea Institute of Sci...)




 Pablo Martin Zamora (CERN)

# T8: Posters (the actors)




 Hristo Umaru Mohamed (CERN)



 Stefan Nicolae Stancu (CERN)



# T8: Posters (the actors)



 Bruno Heinrich Hoeft (KIT - Karlsruhe Instit... )



 David Martin Clavo (CERN)

And many  
more!

Sorry for not be able to take all poster  
presenters.  
Thank you all for the hard work!

# T8: Posters list

- Centralizing Elasticsearch resources (CERN)
- Trident - monitoring for analysing node utilisation such as data IO, CPU core and memory (Tier-0, CERN)
- Concurrent Adaptative Load Balancing at CERN (CERN)
- Design and development of vulnerability management portal for DMZ admins powered by DBPowder (KEK)
- Experience of migration from Puppet 3.5 to Puppet 4.9 migration (LHCb, CERN).
- KairosDB and Chronix as longterm storage for Prometheus (LHCb Online)
- Monitoring, accounting and alerting services at INFN-CNAF Tier1 (INFN)
- IT Service Management at CERN: Data Centre and Service monitoring (CERN)
- Power Usage Effectiveness analysis and optimization in the INFN CNAF Tier-1 data center infrastructure (INFN, CNAF)



# T8: Posters list 2

- Increasing Windows security by hardening PC configurations at CERN (CERN)
- Evolving CERN's Network Configuration Management System (CERN)
- Establishment of new WLCG Tier Center using HTCondorCE on UMD middleware (KISTI)
- CDS Videos - The new platform for CERN videos (CERN)
- Simulation approach for improving the computing network topology and performance of the China IHEP Data Center (IHEP, China)
- Equal cost multipathing in high power systems with TRILL (JINR)
- The GridKa WLCG Tier-1 Center: Status and Plans (GridKa)
- GridKa network development towards LHC run 3 and run 4 (GridKa)
- Developing a monitoring system for Cloud-based distributed datacenters (Bari, INFN)
- Service monitoring system for JINR Tier-1 (JINR)

# (some) Trends from the track

## Networking

- Network resource optimisation, increment, and integration into experiment tools and services
- IPv6 full deployment and use in WLCG context
- Expanding DC networking via L3VPNs to exploit additional CPU resources
- Data caches co-located with network hubs in a similar way as on commercial CDNs
- SDN/NFV approaches - currently looked at by HEPiX NFV WG
- HLT-free contexts with high throughput transmissions from experiment to distant DC

## Facilities

- Security is something we shouldn't overlook
- Automate, automate, automate - better tools and techniques
- Wide adoption of new monitoring tools/techniques that ease the life of users
- Application of ML techniques to analyze the performance of systems and services
- More work on setting next HEP CPU Benchmarks and Fast CPU Benchmarks
- Increase of experiences in elastic integration of external resources into existing DC
- Prototypes focused on (big) Analysis Facilities for users
- Improvements on Cooling systems at the DC facilities

# Thanks from Track 8 conveners

- To the speakers and poster presenters for the interesting topics
  - We acknowledge the speaker efforts in being on time! (12' talk – not easy)  
... and for not bringing firearms, ammunition or explosives to the conference
- To the audience of the sessions for their interest and the vital discussions
- To the LOC team, helpers, and all the scientific secretaries on-site
- To the conference organizers for the great CHEP 2018 in Sofia!



# Благодаря ви



T8 live notes

