

22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016


# Security, Policy and Outreach

## Track 8 Highlights

*David South, DESY*

*Torre Wenaus, BNL*

*Hannah Short, CERN*



22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

# Security, Reproducibility, Outreach, Authentication and Authorisation, Collaborative Tools...

A.K.A. Misc.



# 17 talks


Spread across 3 sessions

# 22 posters

On display throughout the week



Thank you to all our  
contributors!



22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

Within HEP, Long Tail of Science & Citizen Science

**OUTREACH**



CERN  
(2000)



LHCb



CM  
S



ALICE



LIGO

## CERN Virtual Visits

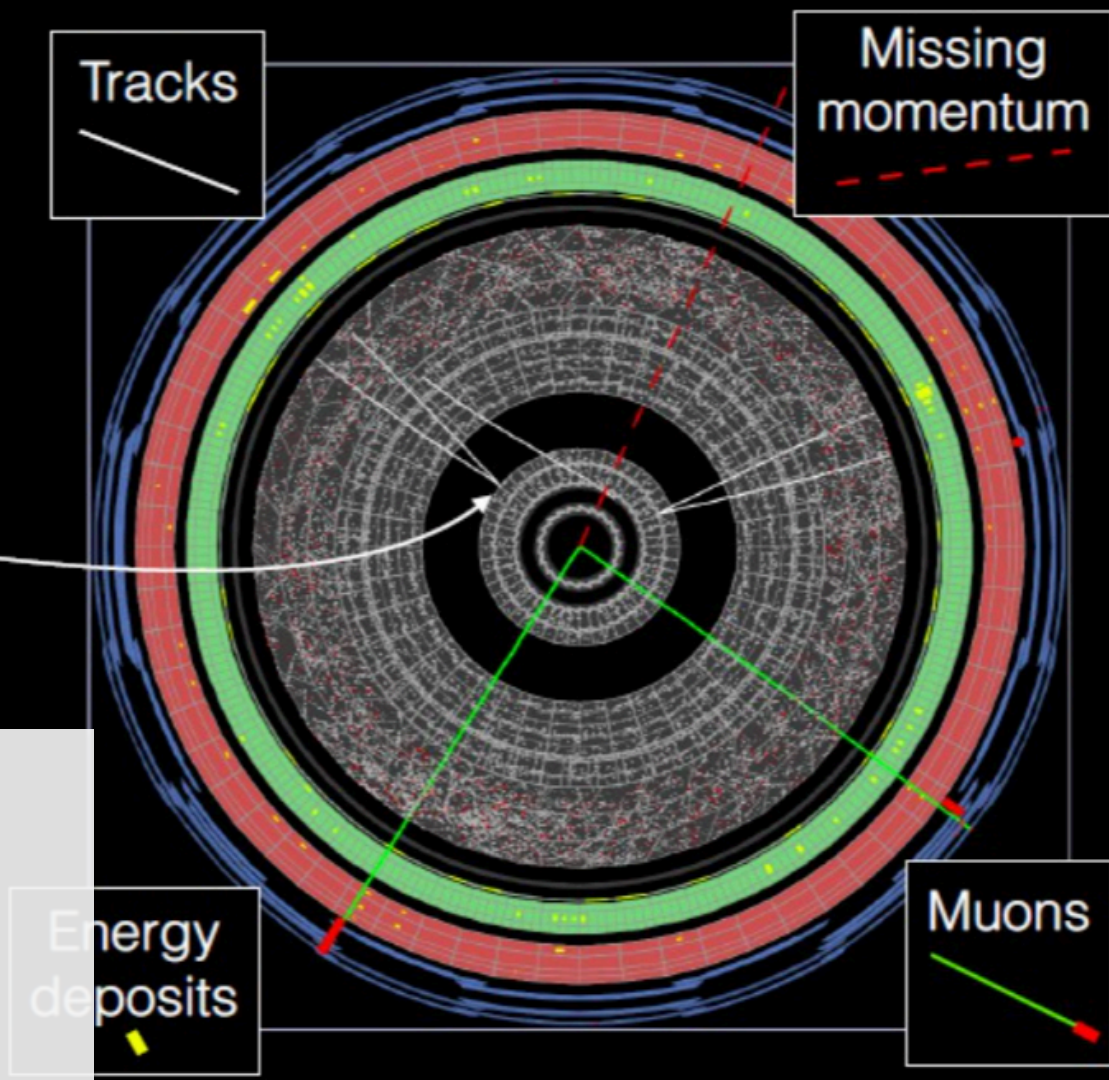
- Bringing CERN tours to the extended public
- Next step is CERN wide coordination

- We ask people to click on “tracks appearing 'out of thin air' away from the centre”
- ‘Off Centre Vertices’

- Here: simulation

### Higgs Hunters

- Citizen science project
- General public did pretty well compared to algorithms!





**Are you sure you want to permanently erase the items in the Trash?**

You can't undo this action.

Cancel

Empty Trash



# Software meets INSPIRE

The image shows two overlapping web interfaces. On the left is the Zenodo repository page for 'decouple software associated to arXiv:1401.0080'. It displays the repository's structure, including folders like 'BatchPlugins' and 'effectiveModel.py'. On the right is the INSPIRE page for the paper 'A Novel Approach to Higgs Coupling Measurements' by Kyle Cranmer, Sven Kreiss, David Lopez-Val, and Tilman Plehn. The INSPIRE page includes the title, authors, date (Dec 30, 2013), and abstract. It also features three plots showing Higgs coupling measurements and a 'Show more plots' link. The Zenodo page includes a 'Description' section that links to the paper and mentions a 'Makefile to recreate the plots in the paper'.

## Invenio

- Share and preserve (!) your data
- Zenodo – long tail of Science
- CERN Open Data – including outreach projects
- DOIs registered and permanent, helping InSPIRE to identify data sets

# The solution

- The **Starterkit** team provides...

- Online **tutorials**
- Interactive **workshops**

- **Goals**

- Improve software **literacy**
- Teach **good practices**



- **Socialisation** amongst **collaboration members**


## LHCb Starterkit

- Teach computing essentials for new joiners
- Interest in reusing the content at other experiments

# Why Simulate Brain Development


---

- Neuro scientific insights
  - How does the brain develop?
  - How do genetic and environmental cues interact?
  - How do neurons and brain regions communicate?
  - How does cognition, planning and memory work?
- Medicine
  - Understanding of brain diseases (epilepsy)
  - Tumor growth
  - Drug development



## CERNopenlab Knowledge Transfer

- Sharing HEP expertise with non-HEP domains
- Genomics and Brain Simulation



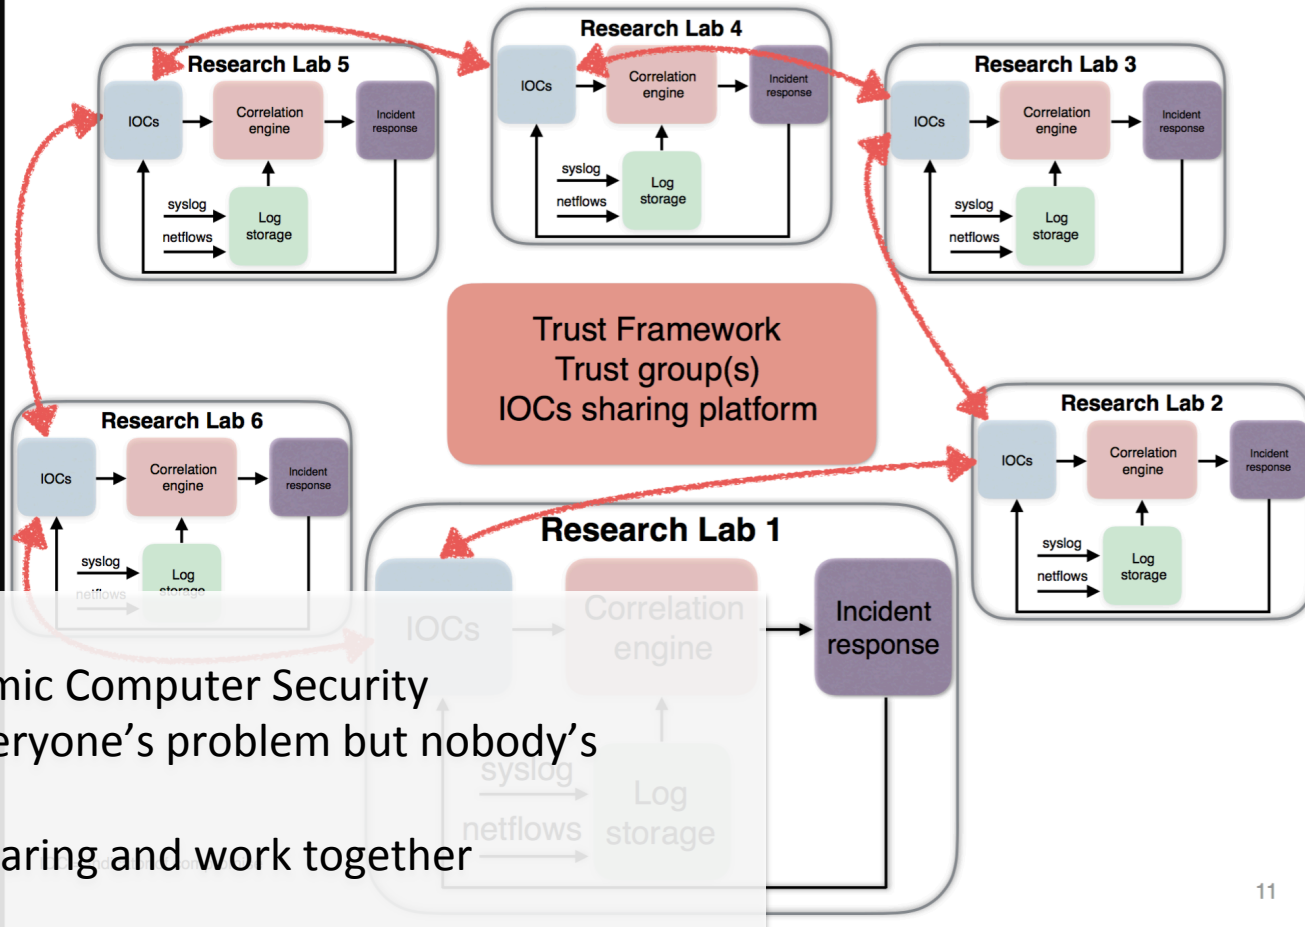
22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

Security Awareness, Future Landscape

**SECURITY**



# A global response



## Future of Academic Computer Security

- Security is everyone's problem but nobody's responsibility
- Enable IOC sharing and work together

# “We are there to help you”

**Computer Security**  
**Incident Response**

Emergencies  
Self-mitigation portal

**Security Consulting, Audits & System Reviews**


CERN Computer Security Team is available to assist you with:

- **Threat modelling and risk assessment** - to make sure that risks are correctly managed, and no major threat is neglected;
- **Designing system security architecture** - when starting a new system or software project;
- **Security code reviews** - before deploying developed code;

Please do not hesitate to contact [Computer.Security@cern.ch](mailto:Computer.Security@cern.ch) if you think that your

## CERN Security Consulting

- Improving communication between security and developers before poor software decisions are made
- Involving users by teaching them penetration testing, Whitehats




22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016


Moving away from certificate based AUTHN

# **AUTHENTICATION AND AUTHORISATION**

## Quick recap



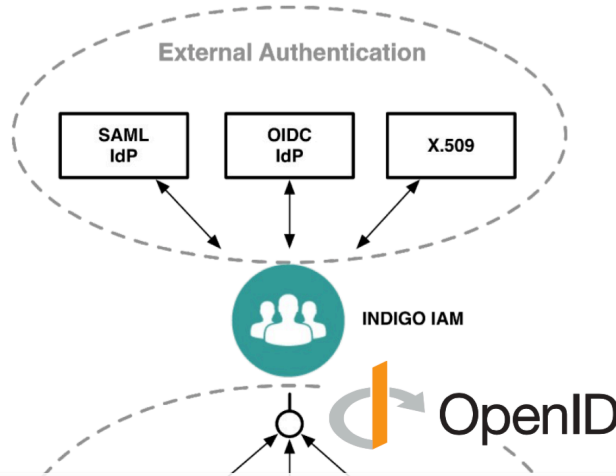
**Authn**



**Authz**



# INDIGO-DataCloud AAI framework



## Federated Identities and Anonymized Delegation

- Funnel identities into a central IdP and use OpenID for authentication and delegation under the hood
- “Macaroons” (like cookies but better) control validity of authorisation tokens

**The INDIGO-DataCloud AAI**  
A. Ceccanti, M. Hardt, B. Wögh, P. Millar, S. Licehammer, M. Caberletti, E. Vianello

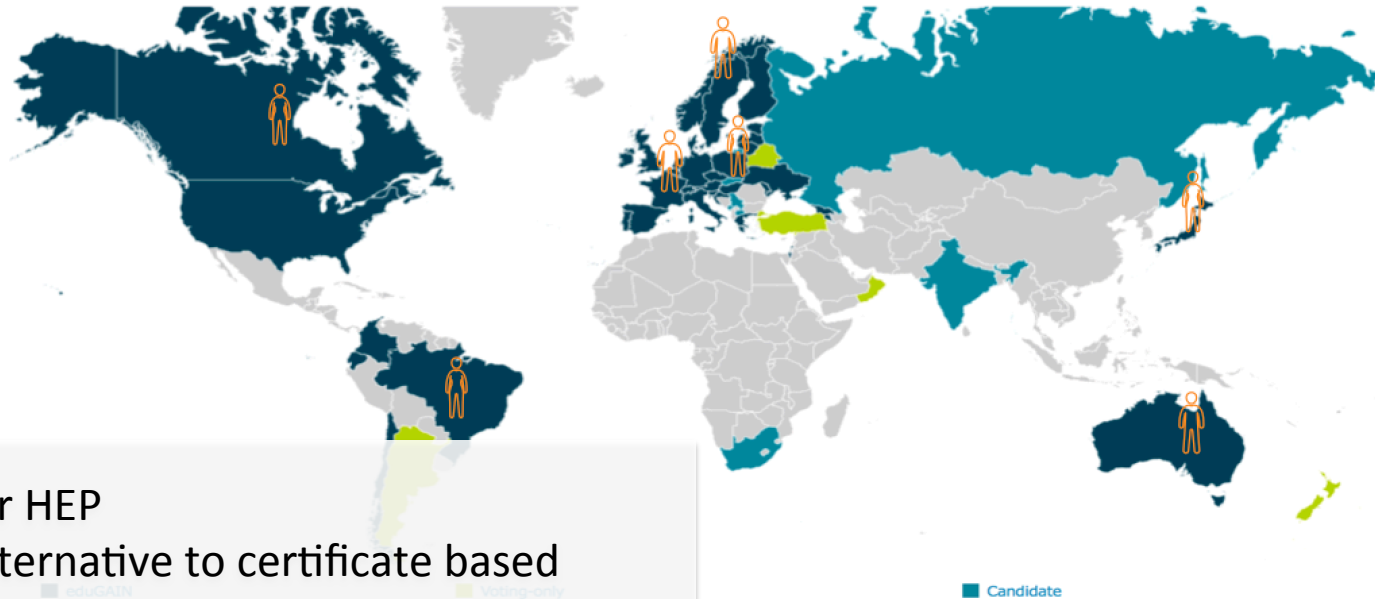
The INDIGO Authentication and Authorization Infrastructure (AAI) leverages and extends existing standards (OpenID Connect, OAuth2, SCIM) to enable secure composition of resources from multiple providers in support of scientific applications. The central **Identity and Access Management (IAM)** service provides tools to implement brokered user authentication, identity harmonization, account linking as well as Virtual Organization (VO) management. Identity information is provided to relying services via **OpenID Connect**, which allows simpler integration in off-the-shelf software. The **Token Translation Service (TTS)** integrates services that do not directly support OpenID Connect, by creating credentials on-demand from the identity information provided by the IAM. This poster introduces the main INDIGO AAI components and highlights the main architectural decisions that were taken during the first year of the INDIGO project.

- 1. Identity = OpenID Connect**  
The INDIGO IAM Identity layer leverages the **OpenID Connect [2]** standard to provide user authentication and identity information to INDIGO services. This approach provides several advantages:
  - Standardized and widely adopted solution
  - Accommodation of several authentication mechanisms (via identity brokering)
  - Simplified integration in relying services
  - Native support for delegation and offline access
  - Native support for dynamic client registration
- 2. Identity harmonization and brokering**  
The INDIGO Identity and Access Management (IAM) [3] service deals with user authentication and identity brokering, allowing users to authenticate with local username and password, SAML, X.509 certficates, or via a remote OpenID Connect provider (e.g., Google). The different identities are linked to a single INDIGO user profile, which provides each user with a **unique, persistent identifier**. This identifier is then used for auditing, accounting and authorization at all relying services. Other attributes, like group membership information, can be linked to the identifier. **Indigo-DataCloud** is a project of the European Union's Horizon 2020 research and innovation programme under grant agreement No 101019723.
- 3. Authorization = OAuth 2 + XACML**  
INDIGO services require authorization through HTTP. The **Authorization and Access Management (AAM)** [4] service provides a **XACML 2.0** policy engine, allowing a user to be authorized to access a resource. The user's identity is translated into a **unique, persistent identifier**, which is then used for auditing, accounting and authorization at all relying services. Other attributes, like group membership information, can be linked to the identifier. **Indigo-DataCloud** is a project of the European Union's Horizon 2020 research and innovation programme under grant agreement No 101019723.
- 4. Identity provisioning**  
INDIGO IAM leverages the standard **System for Cross-Domain Identity Management (SCIM) [5]** version 2.0 to implement identity provisioning, de-provisioning and management. The SCIM APIs provides means to **propagate identity and group information to relying services**, to implement, for instance, **dynamic account creation and other resource lifecycle management at various levels of the INDIGO Infrastructure** depending on events related to user identity status.
- 5. Delegation and offline access**  
OAuth and OpenID Connect support **delegation and offline access natively**, to support **chained delegation across services**, in which a component can act both as a service and as a client for another downstream service. The INDIGO IAM provides a partial implementation of the **OAuth token exchange draft standard [6]**. The token exchange is used in particular to implement **controlled delegation of access rights across services** (i.e., the ability to execute on behalf of a user while the user is not connected).
- 6. Token translation and non-HTTP services**  
ID AAI integrates services that rely on different authentication protocols (e.g., X.509 certficates, SSH keys, Amazon S3 keys) via the **ID Token Translation Service (TTS) [7]**. **TTS** translates an OpenID Connect authentication assertion, like the one issued by the IAM service, into one of the targeted downstream service credentials. The TTS currently supports the generation of:
  - HTTP keys
  - D9 certficates
  - arbitrary username/password details

**References**

[1] <https://www.indigo-aa.org/>  
 [2] [https://openid.net/specs/openid-connect-core-1\\_0.html](https://openid.net/specs/openid-connect-core-1_0.html)  
 [3] <https://indigo-aa.org/indigo-aa-identity-and-access-management/>  
 [4] <https://indigo-aa.org/indigo-aa-authorization-and-access-management/>  
 [5] <https://tools.ietf.org/html/rfc7642>  
 [6] <https://tools.ietf.org/html/draft-ietf-oauth-token-exchange-02>  
 [7] <https://indigo-aa.org/indigo-aa-token-translation-service/>

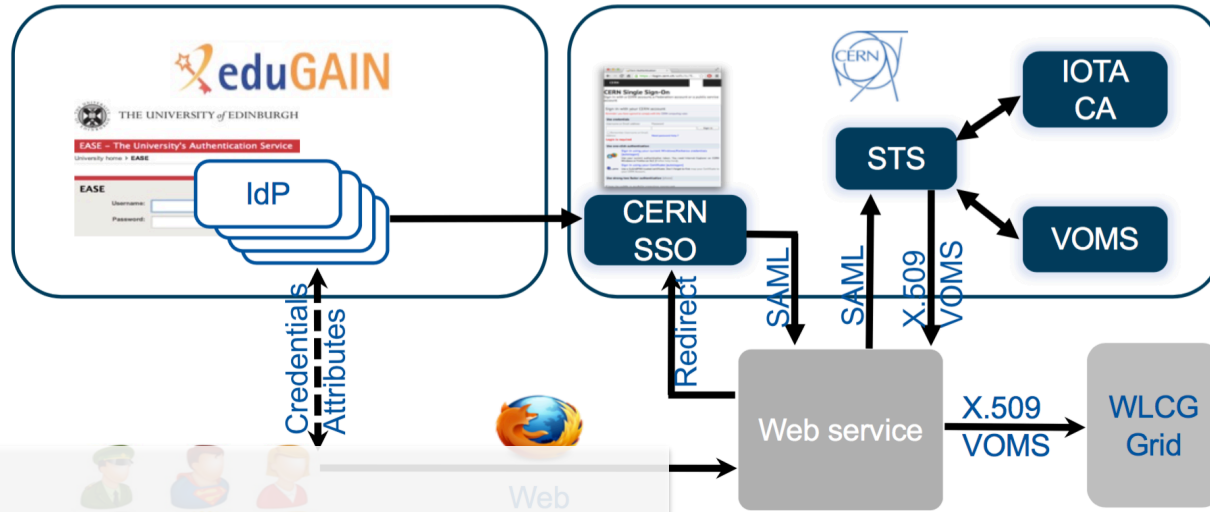
## An alternative source of identities, federated identity via eduGAIN



### eduGAIN for HEP


- Viable alternative to certificate based federation
- Work to ensure that eduGAIN is as secure as current access methods

# How can I implement X509 free?



## X509 Free Pilot

- Transparent eduGAIN access to grid services
- Ready to go for some Vos
- Currently being implemented for PanDA




22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

How are we sharing knowledge, and enabling collaboration?

# **COLLABORATIVE TOOLS**

## Detailed view with quick access to all information


HEPData

ⓘ About
⚙ Help
👤 Sign in

---

Aad, Georges et al.

Last updated on 2016-04-08 15:56:50
📄 Accessed 10 times
📄 Cite

← Hide Publication Information Information

### Charged-particle distributions in $\sqrt{s}=13$ TeV $pp$ interactions measured with the ATLAS detector at the LHC

Aad, Georges , Abbott, Brad , Abdallah, Jalal , Abdinov, Ovsat , Abeloos, Baptiste , Aben, Rosemarie , Abolins, Maris , AbouZeid, Ossama , Abraham, Nicola , Abramowicz, Halina  
ATLAS

No Journal Information, 2016

<http://dx.doi.org/10.17182/hepdata.72205>

**Abstract (data abstract)**

CERN-LHC. Measurements of charged particle distributions in proton-proton collisions at a centre-of-mass energy of 13 TeV. A data sample of nearly 9 million events recorded by the ATLAS detector during a special LHC fill with low beam currents, and thus giving a low expected mean number of interactions, is used. The charged-particle multiplicity, its

Download All

Filter 18 data tables

momentum for...

**Table 9**

Data from Auxiliary Material 10.17182/hepdata.72205.v1/t9  
 Extrapolated charged-particle multiplicity distributions in proton-proton collisions at a centre-of-mass energy of 13000 GeV for events with the number of...

**Table 10**

Data from Auxiliary Material 10.17182/hepdata.72205.v1/t10  
 Extrapolated average transverse momentum in proton-proton collisions at a centre-of-mass energy of 13000 GeV as a function of the number...

**Table 10**

Extrapolated average transverse momentum in proton-proton collisions at a centre-of-mass energy of 13000 GeV as a function of the number of charged particles  $>=1$  having transverse momentum  $>500$  MeV and absolute(pseudorapidity)  $<2.5$ .

[10.17182/hepdata.72205.v1/t10](http://www.hepdata.net/r/10.17182/hepdata.72205.v1/t10)

**observables**

PT

**phrases**

Inclusive Proton-Proton Scattering

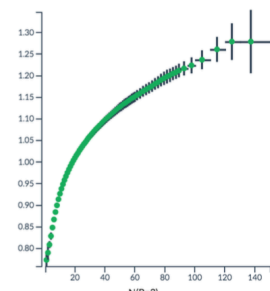
**reactions**

P P --> CHARGED X

Showing 50 of 81 values Show All 81 values

ETARAP(P=3)	-2.5-2.5
Extrapolated to include strange baryons	
N(P=3)	$>= 1$
PT(P=3)	$> 500$ MEV
RE	P P --> CHARGED X
SQRTS(5)	13000.0 GeV
N(P=3)	MEAN(NAME=PT(P=3)) [GEV]
0.50 - 1.50	0.7737 $\pm 0.0008$ stat $\pm 0.0155$ sys
1.50 - 2.50	0.7904 $\pm 0.0007$ stat $\pm 0.0158$ sys
2.50 - 3.50	0.809 $\pm 0.001$ stat $\pm 0.008$ sys

Visualize




Sum errors  Log Scale (X)  Log Scale (Y)

## HEPData

- Evolving into a more inclusive tool
- Allows flexible data format and easy conversion to common output

CHEP 2015, Okinawa



21st International Conference on Computing in High Energy and Nuclear Physics **CHEP2015** Okinawa Japan: April 13 - 17, 2015

See the [Proceedings introduction](#) for an overview of the conference.

CHEP 2015 attracted a very high number of oral and poster contribution, 535 in total, and hosted 450 participants from 28 countries.

CHEP 2015 hosted contributions on online computing; offline software; data store and access; middleware, software development and tools, experiment frameworks, tools for distributed computing; computing activities and computing models; facilities, infrastructure, network; clouds and virtualization; performance increase and optimization exploiting hardware features.

2015-04-13

- Website
- Summary timetable
- Full indico agenda
- Topics and tracks
- Online computing track contributions
- Offline software track contributions
- Data store and access track contributions
- Distributed computing track contributions
- Computing activities and models track contributions
- Facilities, infrastructure, network track contributions
- Clouds and virtualization track contributions
- Exploiting hardware features track contributions
- Proceedings

Q Search

Search the app:  Clear

title  content

Search

21 results by 1 1/2

<> ALFA

- CHEP 2013 DAQ, trigger and controls track proceedings
- CHEP 2013 Data stores, databases and storage systems track proceedings
- CHEP 2013 Distributed processing and data handling track proceedings
- CHEP 2013 Event processing, simulation and analysis track proceedings
- CHEP 2013 Facilities, production infrastructures, networking and collaborative tools track proceedings
- CHEP 2013 Software engineering, parallelism and multi-core programming track proceedings
- CHEP 2013, Amsterdam
- CHEP 2015 Clouds and Virtualization track proceedings
- CHEP 2015 Computing Activities and Computing Models track proceedings
- CHEP 2015 Computing Facilities, Infrastructure, Network track proceedings
- CHEP 2015 Data Store and Access track proceedings

**Objective: draw CHEP content into the cross-referenced knowledge base, have it accessible, visible, discoverable, integrated**

Start made with CHEP 2015, 2013 but only to the session & proceedings track level

Add your own talks & proceedings as links and associate them with their session, sw category etc

## HEP Software & Computing Knowledge Base

- Ongoing effort to share information on software in use in HEP
- <http://hepsoftware.org>

<http://indico-software.org>




## Indico 2.0

- Ongoing migration for 3 years, front end changes are only the tip of the iceberg
- Complex not complicated!
- 78% of legacy code removed



Indico is used every day at CERN to manage more than 300.000 events of different complexities and 200 meeting and conference rooms.



22nd International Conference on Computing in High Energy and Nuclear Physics, Hosted by SLAC and LBNL, Fall 2016

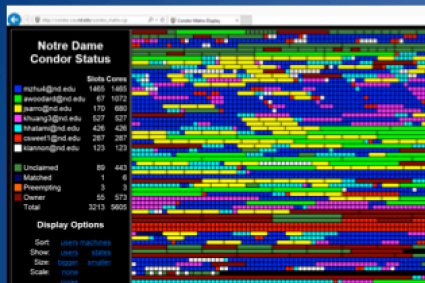
Can we get the same result twice, 10 times, 100 times?

**REPRODUCIBILITY**




# What we usually work on:

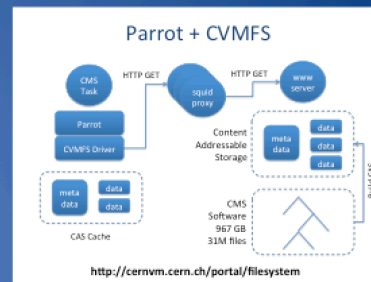
HTCondor



Lobster Data Analysis



Global Filesystems



# Today, a different question:


## Reproducibility

- Where do differences come from? Non-deterministic code
- ROOT-diff tool developed to do bitwise comparison
- By artificially warping machine time, can get bitwise equality... to a point

*What happens if we attempt to run the same thing twice?*

## The big picture

*Providing the information on data production and the software, and linking the monitoring system Dirac and bookkeeping to the CERN analysis preservation (CAP) portal, users and containerised software in Docker*



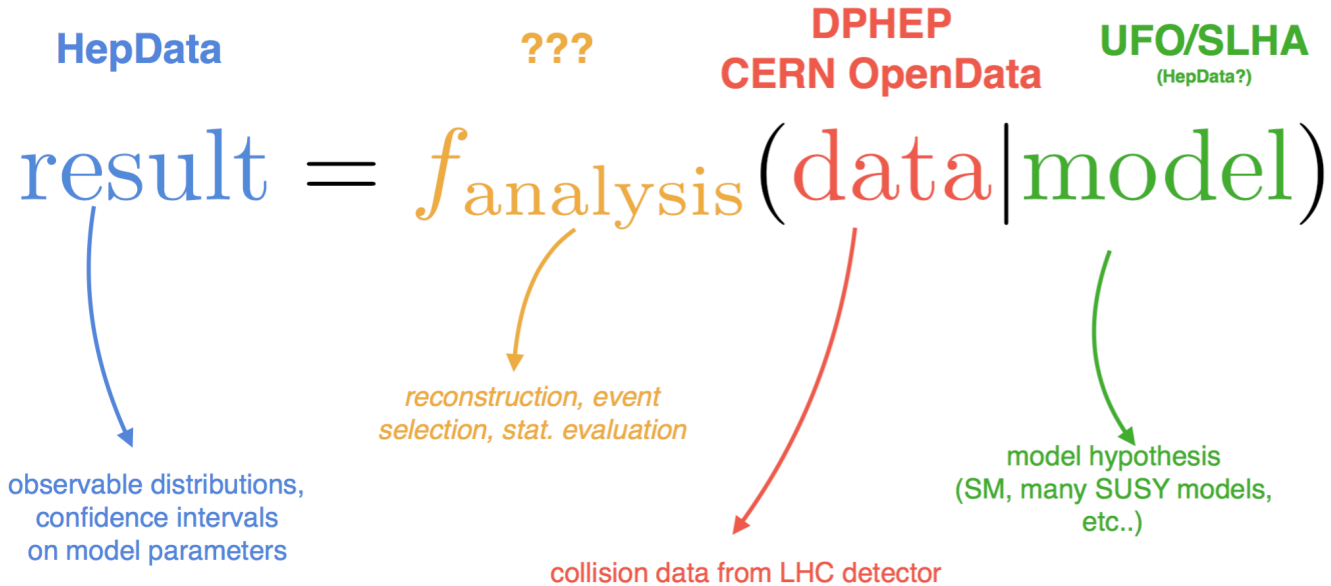
### LHCb Data Preservation

- Consolidating metadata on system and data production into graph database
- Can mine graph for statistics on software usage and help when troubleshooting

## Current Standards / Preservation / Archival efforts:

Efforts like DASPOS and CERN Analysis Preservation try to fill a gap in the preservation efforts.

Data only is not enough. We need software and environment as well.




# How to preserve $f_{\text{analysis}}(\cdot)$ ?

## 2. Problem: Preserve Parametrized Workflow

**Therefore:** Sequentially build up graph, as sufficient information becomes available, using a number of stages that add nodes and edges


**To capture analysis workflow, capture the stages.**



**Example:**  
Parametrized  
Map-Reduce




**Stage 1:**  
unknown number of files. e.g.  
download & unpack archive with a  
p priori unknown # of files



**Stage 2:**  
for each file in the archive, add node  
to process it  
(only possible after first node done)

**Stage 3:**  
add a node that merges results of  
the map nodes  
node/edge can be added before  
execution of map nodes



Par. Set 2

- To capture analysis workflow, capture the individual stages
- Preserve in a parameterized way; source code is not enough, need full environment




And the posters...

# CERN Document Server

Towards the Next Generation CERN Institutional Repository

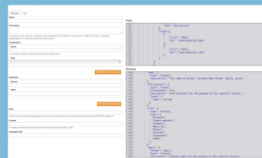
## Retrieval

The new version of CDS significantly improves the information retrieval system. By switching to Elasticsearch as the underlying search engine and redesigning our complex data model, the new CDS provides up to 3 times faster search results and up to 50 times faster indexing. Another determining change is a completely new design that offers a clearer, more intuitive way of searching, and real time filtering and previewing of data.



## Submission

The CDS submission system has undergone a complete redesign based on our extensive experience working with different communities at CERN. The new version, based on a flexible data model, offers the possibility to create tailored content for different user communities. Together with customizable submission and publication workflows comes a new, richer user interface and significant speed improvements that enhance user experience.



## Community driven

The new CDS aims to be the CERN's document hub, acting as an aggregator over specialized repositories, each having its own software stack, with features enabled based on the repository's content.

The aim is to enable each content producer community to have its own identity, both visually and functionally, as well as increased control on the data model and the submission, curation, management, and dissemination of the data.

powered by  
**INVENIO**

### Research

All research documents produced at CERN are stored on this repository. Metadata is validated according to strict standards in order to facilitate reporting and service interoperability.

### Books

Repository for books, standards, proceedings, and reports. Curated by the CERN Library. With circulation features enables this vast online catalog is available to all CERN users.

CDS



### Photos

Repository for images produced by CERN and all the Experiments. It comes with photo-related features like content-based image retrieval and image tagging.

\*\*\*

Repositories for all submission and publication workflows tailored for the Experiments and CERN Departments, with the emphasis on user empowerment, covering a variety of document types.

## Videos



### Submission

- Video upload
- Metadata form
- Data validation

### Processing

- Metadata extraction (FFmpeg)
- Subformats creation (Sovon)
- Thumbnails extraction (FFmpeg)

### Record creation

- OAI indexing
- File storage (D0G)
- Metadata storage (PostgreSQL)

### Presentation

- Video playback
- Downloading
- Embed & Export



<http://cds.cern.ch> © CERN CC-BY-SA 4.0



<http://cds.cern.ch>



[cds.support@cern.ch](mailto:cds.support@cern.ch)



[/CERN/DocumentServer](https://github.com/CERN/DocumentServer)

# CERN Document Server


Towards the Next Generation CERN Institutional Repository

## Retrieval

The new version of CDS significantly improves the information retrieval system. By switching to Elasticsearch as the underlying search engine and redesigning our complex data model, the new CDS provides up to 3 times faster search results and up to 50 times faster indexing. Another determining change is a completely new design that offers a clearer, more intuitive way of previewing of data.

## Submission

The CDS submission system has undergone a complete redesign based on our extensive experience working with different communities at CERN. The new version, based on a flexible data model, offers the possibility to create tailored content for different user communities. Together with customizable submission and publication workflows comes a new, richer user interface and significant speed improvements that enhance user experience.



## IPv6 Security

M. Böhler, J. Chabbert, A. Deschamps, T. Haverly, T. Iny, C. Gargano, K. Hübner, B. Hurrell, D. P. Kelly, J. Laine-Matijevic, S. Mennelli, B. Nandakumar, K. Oberholzer, F. Pösch, D. Radd, J. A. Schell, D. Trötschel, U. Tregenna-Reyer and R. Ward

arXiv:1601.06105v1 [hep-th] 12 Jun 2016

IPv6 network addresses are running out and the deployment of IPv6 networking in many places is now well underway. The work of the HEPX IPv6 Working Group, a growing number of sites in the Worldwide Large Hadron Collider Computing Grid (WLCG) have deployed dual-stack IPv6/IPv4 services. The aim of this is to support the use of IPv6-only clients, i.e. worker nodes, virtual machines or containers. The IPv6 networking protocols while they do contain features aimed at improving security also bring new challenges for operational IT security. We have spent many decades understanding and fixing security problems and concerns in the IPv4 world. Many WLCG IT support teams have only just started to consider IPv6 security and they are far from ready to follow best practice, the guidance for which is not easy to find. The lack of maturity of IPv6 implementations together with the increased complexity of the protocol standards and the fact that the new protocol stack allows for pretty much the same attack vectors as IPv4, raise many new issues for operational security teams. The HEPX IPv6 Working Group is producing guidance on best practices in this area. We consider some of the security concerns for WLCG in an IPv6 world and present the HEPX IPv6 working group guidance both for the system administrators who manage IT services on the WLCG distributed infrastructure and also for their related security and networking teams.

### Checklist for Administrators

- I. Ensure all security/network monitoring/logging are IPv6-capable
- II. Filter IPv6 packets that enter and leave your network/system
- III. Filter/disable IPv6-on-IPv4 tunnels
- IV. Deploy RA Guard or otherwise deal with Rogue RA
- V. Filter ICMPv6 messages wisely
- VI. Allow special-purpose bundles only if needed
- VII. Make an addressing plan
- VIII. Decide whether to use DHCPv6 or SLAAC-DynDNS
- IX. Use synchronized IPv4/IPv6 access rules
- X. Do not be tempted by transition technologies

### Checklist for Developers

- I. Code that replaces IPv4 transport with IPv6 is expected to behave as well and to be tested at least as well as existing code/plan for extensive testing
- II. Make sure that the choice/coloring/preference of source and destination IP address follows what is administratively chosen and configured at the OS level
- III. Existing IPv4 security measures should not be removed, worked around or simply forgotten when porting code for IPv6

## New IPv6 features

Longer IP addresses  
 • They may slow down brute force scans  
 • But no bad guy is that crude...

Cannot generate IPv6 on-site  
 • Minimum ATU (see RFC 6502)  
 • Still have small fragments of IPv6 on wide area local networks at least

Transition technologies (e.g. tunnels)  
 • Have various vulnerabilities on their own  
 • Need to be three horses...

## Business as usual

Not really a feature of IPv6: proper, but much of the network stack and application code is entirely IPv4-fresh!

Broadcasts and Multicasts are still there, with a vengeance  
 • Can still run a rogue DHCP server

Can still pollute Ethernet address discovery (ND instead of ARP)  
 • Can still use IP headers for out-of-band communications

Can still try forging and injecting packets into the local network  
 • Upper-layer protocols did not change!

\* As long as all network monitoring and administration tools are up-to-date and (therefore) aware of IPv6.

CHEP 2016 San Francisco - October 10-14

CDS

### Research

All research documents produced at CERN are stored on this repository. Metadata is validated according to strict standards in order to facilitate reporting and service interoperability.

### Books

Repository for books, standards, proceedings, and reports. Curated by the Library. With circulation features enables this vast online catalog is available to all CERN users.

### Photos

Repository for images produced by CERN and all the Experiments. It comes with photo-related features like content-based image retrieval and image tagging.

### Submission

- Video upload
- Metadata form
- Data validation

### Processing





- Metadata extraction (Pfmpeg)
- Subformats creation (Sovisart)
- Thumbnails extraction (Ffmpeg)

### Record creation

- OAI metadata
- File storage (S3)
- Metadata storage (Pfmpeg/S3)

### Presentation

- Video embed
- Download
- Embed & Export


http://hepdx.cern.ch © CERN CC-BY-SA 4.0

http://cds.cern.ch

cds.support@cern.ch

CERN Document Server



M Domaracky, N Juszcak, M Peksa, T Baron

Retrieval

The new version of CDS significantly improves the information retrieval system. By switching to Elasticsearch as the underlying search engine and redesigning our complex data model, the new CDS provides up to 3 times faster search results and up to 50 times faster indexing. Another determining change is a completely new design that offers a clearer, more intuitive way of previewing data.

Sub

The CDS submission system has our extensive experience working in new version, based on a flexible tailored content for different user submission and publication work and significant speed improve



For almost 10 years, CERN has been providing live webcast of events using Adobe Flash technology. This year is finally the year when flash died at CERN! With Flash being slowly phased out on most streaming platforms, the CERN streaming service moved as well from Flash to HTTP streaming.



The biggest challenge for providing pure HTML5 video streaming goes with the support of different streaming protocols across browsers and OS platforms. Thanks to THEOPlayer (www.theoplayer.com) we are able to stick with the HTTP Live Streaming (HLS) protocol for Android and play it on all modern browsers on desktops and mobile devices. We are able to deliver the same experience as we did with Adobe Flash based players. Our users can still enjoy video of the speaker synchronized with video of the presentation, so they have the same experience as sitting in the auditoria.


DVR – Digital Video Recorder

We will soon introduce DVR functionality for all our live webcasts. Users that arrived late on the webcast website, will have a possibility to go back to the beginning of the webcast or if they missed something they can seek back to watch it again. With the DVR feature we will also be able to publish the recording right after the webcast is finished.




New HTML5 viewer for recorded lectures

The CERN Recording Service provides recordings from 19 CERN rooms in a web lecture format (with synchronized speaker and slides). To give users a better control on the player side, we developed a new interactive HTML5 viewer with a built-in THEOPlayer.



Central Encoding System

With 19 CERN rooms capable of webcast and recording, about 300 live webcasts and 1200 lectures recorded every year, we needed a tool for our operators to easily start webcasts and recordings. We developed a Central Encoding Interface, from which our operators see all the planned events for a given day and with one click can start webcasting and/or recording. With this new interface we manage to faster publish the recorded lecture by integrating the Sorensen (http://www.sorensenmedia.com/queue/squeeze-server/) transcoding cluster to Micala. The communication with users is now done with the automatic creation of the ServiceNOW ticket for each recording.



Service statistics for last years



http://cern.ch/IT ©CERN CC-BY-SA 4.0

## IPv6 Security

M. Böhler, J. Chabada, A. Deshotel, T. Hasegawa, T. Inoue, C. Legrand, K. Nihei, R. Nozji, D. S. Raley, J. Laine-Mattar, S. Nourou, M. Nankamaki, K. Okamoto, F. Pöhl, D. Rudi, J. A. Sola, D. Traylor, U. Tegenrod and R. Ward

arXiv:1601.04561v1 [hep-ph] 16 Jan 2016

IPv6 network addresses are running out and the deployment of IPv6 networking in many places is now well underway. Following the work of the HEP-X IPv6 Working Group, a growing number of sites in the Worldwide Large Hadron Collider Computing Grid (WLCG) have deployed dual-stack IPv6/IPv4 services. The aim of this is to support the use of IPv6-only clients, i.e. worker nodes, virtual machines or containers. The IPv6 networking protocols while they do contain features aimed at improving security also bring new challenges for operational IT security. We have spent many decades understanding and fixing security problems and concerns in the IPv4 world. Many WLCG IT support teams have only just started to consider IPv6 security and they are far from ready to follow best practice, the guidance for which is not easy to find. The lack of maturity of IPv6 implementations together with the increased complexity of the protocol standards and the fact that the new protocol stack allows for pretty much the same attack vectors as IPv4, raise many new issues for operational security teams. The HEP-X IPv6 Working Group is producing guidance on best practices in this area. We consider some of the security concerns for WLCG in an IPv6 world and present the HEP-X IPv6 working group guidance both for the system administrators who manage IT services on the WLCG distributed infrastructure and also for their related security and networking teams.

Challenges for Administrators

- I. Ensure all security for network monitoring/logging are IPv6-capable
- II. Filter IPv6 packets that enter and leave your network/system
- III. Filter/disable IPv6-to-IPv4 tunnels
- IV. Deploy KA-Guard or otherwise deal with Rogue KA
- V. Filter ICMPv6 messages wisely
- VI. Allow special-purpose features only if needed
- VII. Make an addressing plan
- VIII. Decide whether to use DHCPv6 or SLAAC-DynDNS
- IX. Use synchronized IPv4v6 access rules
- X. Do not be tempted by transition technologies

Challenges for Developers

- I. Code that replaces IPv4 transport with IPv6 is expected to behave as well and to be tested at least as well as existing code/plan for minimum testing
- II. Make sure that the choice/coloring/preference of source and destination IP address follows what is administratively chosen and configured as the OS level.
- III. Existing IPv4 security measures should not be removed, worked around or simply forgotten when porting code for IPv6

New IPv6 features

New methods for auto-configuring addresses, routes, DNS

Good for the end-user

Must do something against Rogue Router Advertisements (see RFC6964)

Longer IP addresses

Mitigates known issues that "they may slow down brute force scans when no bad guy is that crude..."

Cannot prevent malware on-site

Mitigates APTs (Advanced Persistent Threats) still that "stealed data and stole good guys at the same time"

Business as usual

Broadcasts and Multicasts are still there, with a vengeance

Can still run a rogue DHCP server

Can still pollute Ethernet address discovery (ND instead of ARP)

Can still use IP headers for out-of-band communications

Can still try forging and injecting packets into the local network

Upper-layer protocols did not change!

\* As long as all network monitoring and administration tools are up-to-date and (therefore) aware of IPv6.

CHEP 2016 San Francisco – October 10–14

CDS

Document

Software stack, repository

Research

All research documents produced at CERN are stored on this repository. Metadata is validated according to strict standards in order to facilitate recording and service interoperability.

Book

Repository of circuitry, circuitry cataloging

Content producer

with visually controlled on, duration, the data.

Photos

Repository for images produced by CERN and all the Experiments. It comes with photo-related features like content-based image retrieval and image tagging.

Micala

Michigan and CERN Automated Lecture Archiving (Micala) is an automated tool developed at CERN with a collaboration from the University of Michigan. We improved the workflow to faster publish the recorded lecture by integrating the Sorensen (http://www.sorensenmedia.com/queue/squeeze-server/) transcoding cluster to Micala. The communication with users is now done with the automatic creation of the ServiceNOW ticket for each recording.

Record creation

Full screen  
Flash message (D00)  
Metadata storage (PmgPQJ)

Presentation

Download content  
Embed & Export

http://cern.ch/IT ©CERN CC-BY-SA 4.0

http://cds.cern.ch

cds.support@cern.ch

CERNDocumentServer





# Web Application Detection (WAD)

for asset inventory and vulnerability management

Sebastian Łopieński / CERN Computer Security Team  
e-mail: Sebastian.Lopienki@cern.ch





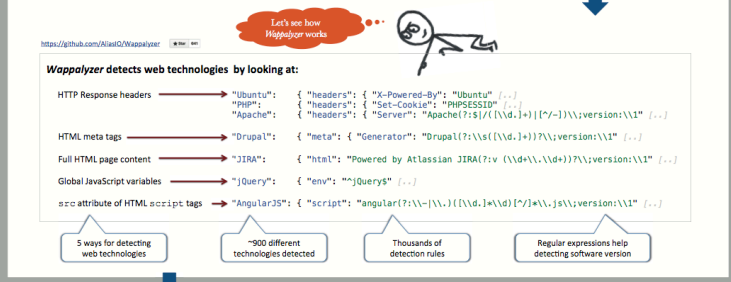
# CERN Document Server

Towards the Next Generation CERN Institutional Repository



# Flash is Dead. Finally.

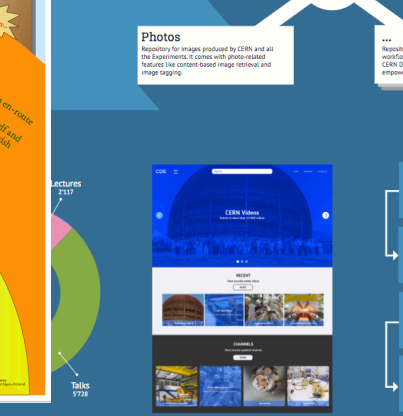
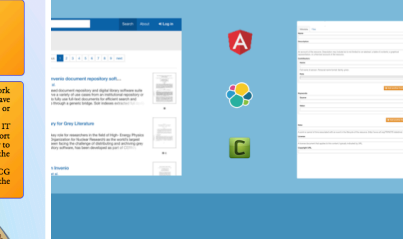
M Domaracky, N Juszcak, M Peksa, T Baron



## Retrieval

significantly improves the information retrieval (at least as far as the underlying search engine and data model, the new CDS provides up to 3 times p to 50 times faster indexing. Another determining w design that offers a clearer, more intuitive way of previewing of data.

The CDS submission system has our extensive experience working new version, based on a flexible tailored content for different use submission and publication work and significant speed improve




For almost 10 years, CERN has been providing live webcast of events using Adobe Flash technology. This year is finally the year when flash died at CERN. With Flash being slowly phased out on most streaming platforms, the CERN streaming service moved as well from Flash to HTTP streaming.



The biggest challenge for providing pure HTML5 video streaming goes with the support of different streaming protocols across browsers and OS platforms. Thanks to THEOplayer (www.theoplayer.com) we are able to stick with the HTTP Live Streaming (HLS) protocol from Apple and play it on all modern browsers on desktops and mobile devices. We are able to deliver the same experience as we did with Adobe Flash based players. Our users can still enjoy video of the speaker synchronized with video of the presentation, so they have the same experience as sitting in the auditoria.


## DVR – Digital Video Recorder

We will soon introduce DVR functionality for all our live webcasts. Users that arrived late on the webcast website, will have a possibility to go back to the beginning of the webcast or if they missed something they can seek back to watch it again. With the DVR feature we will also be able to publish the recording right after the webcast is finished.



## Micala

Michigan and CERN Automated Lecture Archiving (Micala) is an automated tool developed at CERN with a collaboration from the University of Michigan. We improved the workflow to faster publish the recorded lecture by integrating the Sorensen (http://www.sorensenmedia.com/squeeze/squeeze-server/) transcoding cluster to Micala. The communication with video is now done with the automatic creation of the ServiceNOW ticket for each recording.



http://cern.ch/IT ©CERN CC-BY-SA 4.0

## Record creation


- Log in to the system
- Flash message (DVR)
- Metadata storage (MySQL)

## Presentation

- View presentation
- Download presentation
- Embed & Export


## New HTML5 viewer for recorded lectures

The CERN Recording Service provides recordings from 19 CERN rooms in a web lecture format (with synchronized speaker and slides). To give users a better control on the player side, we developed a new interactive HTML5 viewer with a built-in THEOPlayer.




## Central Encoding System

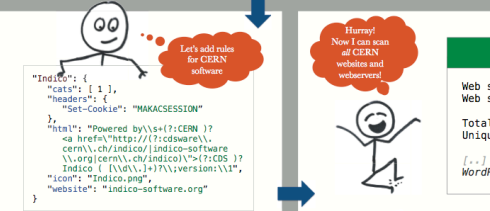


With 19 CERN rooms capable of webcast and recording, about 300 live webcasts and 1200 lectures recorded every year, we needed a tool for our operators to easily start webcasts and recordings. We developed a Central Encoding Interface, from which our operators see all the planned events for a given day and with one click can start webcasting and/or recording. With this new interface we manage to almost eliminate issues where operators forget to start the webcast and with an automatic stop, we now support webcasts and recordings which finish out of standard working hours without additional manpower expenses.



## Service statistics for last years










**Evolution of CERN Print Services**  
Retrieval  
Sub  
Vincent Bippus, Natalie Kane, Tim Smith (CERN, IT-CDA)

**Golden rule: everyone at CERN can print** - ALL staff, visitors, conference participants, anyone!  
**Demanding environment: approx. 1000 networked print devices in 230 buildings**  
**Our challenge: evolve Print Services into a robust and user-focused service**






For almost 10 years, CERN has been providing live webcast of events using Adobe Flash technology. This year is finally the year when flash died at CERN! With Flash being slowly phased out on most streaming platforms, the CERN streaming service moved as well from Flash to HTTP streaming.

The biggest challenge for providing pure HTML5 video streaming goes with the support of different streaming protocols across browsers and OS platforms. Thanks to THEOPlayer (www.theoplayer.com) we are able to stick with the HTTP Live Streaming (HLS) protocol from Apple and play it on all modern browser devices. We are able to deliver with Adobe Flash based pl video of the speaker's presentation, so they have the the auditoria.



**BelleII@home : Integrate volunteer computing resources into DIRAC in a secure way**  
Wenjing Wu, Takano Hara, Hideaki Miyake, Kou Ueda, Wenzhao Kan, Phillip Urrujo  
IHEP 138 Yuquan Road, Beijing, 100049 China  
\*KEK, 1-1 Oho, Tsukuba, Ibaraki, 305-0801 Japan  
\*University of Melbourne, Parkville VIC 3010, Australia



**BelleII@home prototype** is developed and deployed, we tested its full workflow which proves the feasibility and stability of this approach. This approach can also be applied on HPC systems whose nodes do not have the payload (without credential), upon finishing the payload, it compresses the data outband communication to interact with the DIRAC server in general.

**Validation:** This work is sponsored by the KEK short term visiting fellow program.

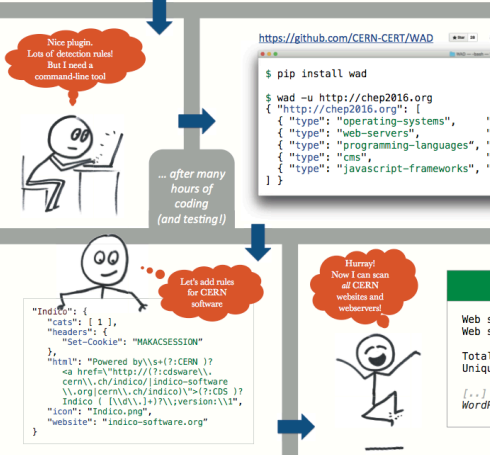
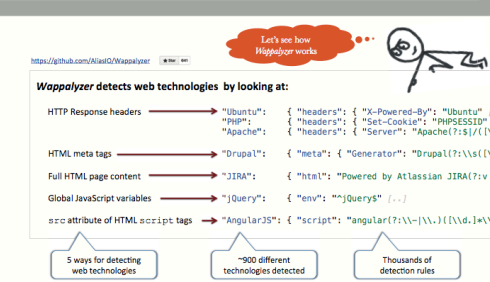
- Fetch and cache the job's input data/scripts from the DIRAC BOINC Site.
- Wrap the job with a script which controls the payload running.
- Submit the job to the BOINC Job Queue.
- BOINC Client fetches the job, and starts a virtual machine to run the payload (without credential), upon finishing the payload, it compresses the data outband communication to interact with the DIRAC server in general.
- Transfer the work directory back to the BOINC Server.
- Validates the results, and mark a successful result.
- Assimilator uploads the root files to the Grid Storage Element.
- Assimilator uploads send/updates job status to the DIRAC Server.



# Web Application Detection (WAD)

for asset inventory and vulnerability management

Sebastian Łopieński / CERN Computer Security Team  
e-mail: Sebastian.Lopienki@cern.ch



# CERN Document Server

Towards the Next Generation CERN Institutional Repository

## Evolution of CERN Print Services

Vincent Bippus, Natalie Kane, Tim Smith (CERN, IT-CDA)

**Golden rule: everyone at CERN can print – ALL staff, visitors, conference participants, anyone!**  
Demanding environment: approx. 1000 networked print devices in 230 buildings  
Our challenge: evolve Print Services into a robust and user-focused service

### Legacy issues

- New strategy was needed to deal with legacy issues:
  - Over 100 different printer models
  - Support for only half of models onsite
  - Photocopiers evolved into multifunction copier machines
  - But we were not involved in specifications
  - Huge administrative overhead costs
  - No synergies between support providers

### Standardisation

- First steps centred around standardisation:
  - Standardised printer models (one supplier not two)
  - Standardised multifunction copier models (one supplier not three)
  - Closely monitored current and future on-site support


### Invitation to Tender

- A tender was issued including specifications for:
  - Standardised multifunction copier models with set prices (just 4 models)
  - Plus on-site support
  - Single per page cost across all models
  - Contractor had to be manufacturer
  - Scored tests carried out for functionality and usability
  - Included option to cover onsite support for whole Printer Park



### Positive initial results

- After 3 years of operation:
  - Reduced lease costs: by average of 33% per model
  - Reduced downtime with onsite support for multifunction copier machines
  - Increased proactive support: now 56% of tickets being dealt with proactively
  - Same interface on all machines: all multifunction copier machines updated
  - Simplified support: all multifunction copier machines connected to network
  - Same support team: covering whole of CERN Printer Park
  - Canceled printer-only support contract: reducing support costs by 35%



### Improving user experience

We provide search and filtering options to organise list of printers:

- By functionality:
  - Color
  - A3
  - Multifunction Copier
- By geolocation:
  - Printers available in your building
  - Printers available in buildings around your actual position



### Easy printer installation

Printer Status page gives basic information including map

- Windows:
  - For Windows domain members only
- Mac:
  - Signed applescript application
  - Launched through URL
  - Gets printer settings and information
  - Downloads appropriate drivers package (if not already present)
  - Configures CUPS (Common Unix Print System) using ipadm command also:
  - Self Service
  - User-friendly interface, based on JAMF Casper Suite
  - Installs appropriate driver
  - Multiple queue configuration in one go
  - Proposes nearby printers, complete list also available

### Simple print solution for visitors

MultiPrint Services: any device registered on CERN network can send email. Can print email body and/or attachments.

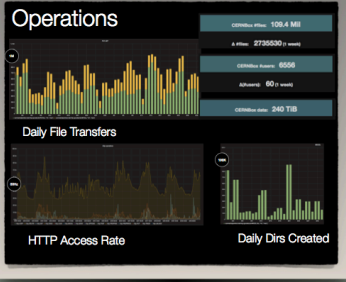

Validation link emailed to user, opens website to select:

- Printer to use
- Print options, such as colour or double-sided



### Future plans

- Reduce printer numbers: find solution for areas with multiple devices
- Increase proactive support: extend to printers? Identify owners & budget cost
- Extend leasing to printers? pay per usage but minimum charge, which would allow:
  - Remove printers from areas left with projects with time-limited budget
  - Reduce admin overhead costs as no need to order & hold stocks
- Easier to update models when required
  - Automatically add to machines with proactive support!
- Enable 'my favourite printers' on Windows?
- New dedicated app: enabling printer configuration on visitor devices



### Big Data Repositories at the Fingertips

Share

Open Source Storage

XRootD

Shared File Spaces for All Platforms

# Data hub for analysis, shared file spaces and more

- G. Adde
- B. Chan
- D. Dunst
- A. Fiorot
- N. Gurnev
- H. G. Labrador
- J. Lopez
- M. Lammiman
- L. Mascetti
- P. Mato
- J.T. Mosicki
- A.J. Peters
- D. Piparo
- E. Tejedor

# CERNBox

CERN IT-IST  
Contact: jakub.mosicki@cern.ch

### Cloud data analysis storage backend

#### Scientific & Educational Notebooks

SWAN is an integrated analysis environment which allows to write data acquisition in CERN's low-level environment (e.g. GEANT4), but may also operate on data that are synchronised and shared with other repositories.

Notebooks may be viewed or opened by clicking directly from CERNBox.

### Education & Research Community

#### Cloud Services for Synchronisation and Sharing (CSS)

Novel applications, cloud storage technology, collaborations

SURFSara Amsterdam  
30 Jan - 1 Feb 2017

PROGRAMME COMMITTEE  
Guido Allan (AAIN), Massimo Lammiman (CERN), Luca Mascetti (CERN), Jakub Mosicki (CERN), Tom Springer (IT-IST, High Temperature Lab/AWI)

https://css3.surfsara.nl

### Integration Architecture

Clients: CERNBox, EOS, Spark

Server: CERNBox, EOS, Spark

Storage: EOS Open Source Storage

### Federation

Multi-vendor API to connect on-premise clouds and sync/share solutions.

To enable secure, open, frictionless file sharing everywhere



# Web Application Detection (WAD)

for asset inventory and vulnerability management

Sebastian Łopieński / CERN Computer Security Team  
e-mail: Sebastian.Lopienki@cern.ch



https://github.com/AliaxiQ/Wappalizer

**Wappalizer detects web technology**

- HTTP Response headers
- HTML meta tags
- Full HTML page content
- Global JavaScript variables
- src attribute of HTML script tags


5 ways for detecting web technologies



```

"Indico": {
  "cats": [ {} ],
  "headers": {
    "Set-Cookie": "MAKACSESSIO
    ...
  }
}

```




# CERN Document Server

Towards the Next Generation CERN Institutional Repository

## Print Services

For visitors, conference participants, anyone!  
Worked print devices in 230 buildings  
A robust and user-focused service




**Improving user experience**  
Provides search and filtering options to organise list of printers:  
By functionality:  Color  A3  Multifunction Copier  
By geolocation:  Process available in your building  
Printers available in buildings around your actual position



- Easy printer installation**  
Printer Status page gives basic information including map  
With easy Web Add Printer buttons:
- Windows:
    - For Windows domain members only
  - Mac:
    - Signed appscript application
    - Launched through URL
    - Gets printer settings and information
    - Downloads appropriate drivers package (if not already present)
    - Configures CUPS (Common Unix Printer System) using ipadm command
  - Android:
    - Self Service
    - User-friendly interface, based on JAAM Casper Suite
    - Installs appropriate driver
    - Multiple queue configuration in one go
    - Proposes nearby printers, complete list also available

**Simple print solution for visitors**  
AllPrint Services: any device registered on CERN network can send email to print email body and/or attachments  
Printer link emailed to user, opens website to select:  
Printer to use  
Print options, such as colour or double-sided

**Future plans**  
Reduce printer numbers: find solution for areas with multiple devices  
Increase proactive support: extend to printers? Identify owners & budget cost  
Extend leasing to printers? pay per usage but minimum charge, which would help  
Reduce printers from across list & high projects with time-limited budget  
Reduce admin overhead costs as no need to order & hold stocks  
Easier to update models when required  
Automatically add to machines with proactive support!  
Enable "my favourite printers" on Windows?  
Web dedicated app: enabling printer configuration on visitor devices





# An educational distributed Cosmic Ray detector network based on ArduSiPM using microcontrollers as data acquisition node NTP protocol as time distribution and IoT technology for data aggregation.

Valerio Bocci, Giacomo Chiodi, Paolo Fresch, Francesco Iacoangeli, Luigi Recchia


INFN Roma, Piazzale A.Moro, 2 - 00185 Roma  
valerio.bocci@roma1.infn.it

The advent of microcontrollers with enough CPU power and with analog and digital peripherals give the possibility to design a complete acquisition system in one chip. The existence of a world wide data infrastructure as internet allows to think at distributed network of detectors capable to elaborate and send data or respond to settings commands. The internet infrastructure allow us to do things unthinkable a few years ago, like to distribute the absolute time with tens of milliseconds precision simple devices far apart from a few meters to thousands of kilometers and to create a Crowdsourcing experiment platform using simple detectors.



The terms of IoT (Internet of Things) define a set of data communication protocols and the capability of single embedded electronics objects to communicate using the internet.  
The MQTT (Message Queue Telemetry Transport) is one of the main protocol used in IoT. Device for data transmission over TCP/IP, the client version can run easily in nowadays microcontrollers, the MQTT broker (the server version) can run also in credit card-sized single-board computers as well in big server.

## ESP8266 MQTT, NTP and WiFi processor



The ArduSiPM sends data over n232, the WiFi Processor elaborate the data and send them using the MQTT protocol to the server. We use as network processor the Espressif ESP8266 low-cost W80P-chip with full TCP/IP stack and a 32-bit RISC CPU running at 80 MHz. The ESP8266 can be used to send and configure MQTT packets, NTP request and configure ArduSiPM device.

## ArduSiPM web configuration pages.



Using Network Time Protocol (NTP) the absolute time from the network, with a precision of tens milliseconds. The network time can be used from a cloud of ArduSiPMs to detect offline coincidence events linked to Ultra High Energy Cosmic Ray

## Users

6K Users, 20K Shares

ICHEP Chicago

July August

## Operations

6K Users, 20K Shares

2735500 (7 weeks)

8556

60 (7 weeks)

240 TIB

## Daily File Transfers

HTTP Access Rate

Daily Dirs Created

## Big Data Repositories at the Fingertips

Share

XRoad

## Shared File Spaces for All Platforms

## Data hub for analysis, shared file spaces and more

CERN IT-IST  
Contact: jakub.mroscicki@cern.ch

- G. Adde
- B. Chan
- D. Duert
- A. Fiorot
- N. Gussone
- H.G. Labrador
- J. Lopez
- M. Lamsina
- L. Mascetti
- P. Mato
- J.T. Moscicki
- A.J. Peters
- D. Piparo
- E. Tejedor

## Cloud data analysis storage backend

Scientific & Educational Notebooks

Sharing of SWAN notebooks via CERNBox

Notebooks may be viewed or opened by hitting directly from CERNBox.

## Education & Research Community

Cloud Services for Synchronisation and Sharing (CSS)

Novel applications, cloud storage technology, collaborations

SURFSara Amsterdam  
30 Jan - 1 Feb 2017  
https://ics3.surfsara.nl

PROGRAMME COMMITTEE

Guido Allen (SURF), Massimo Lamsina (CERN), Luca Mascetti (CERN), Jakub Moscicki (CERN), Tom Soper (CERN), Rick Teegen (SURF)

## Integration Architecture

Clients

REST API

Server

EOS Open Source Storage

## Federation

Multi-vendor API to connect on-premise clouds and sync/share solutions.

To enable secure, open, frictionless file sharing everywhere




# Web Application Detection (WAD)

for asset inventory and vulnerability management

Sebastian Łopieński / CERN Computer Security Team  
e-mail: Sebastian.Lopienki@cern.ch



Another usual day at CERN...  
Fortunately they've already published a security patch  
Aargh, WordPress is vulnerable again!



Nice plugin. Lots of detection rules! But I need a command-line tool

... after hour code (and tea)

```



{
  "Indico": {
    "cats": [ "1" ],
    "headers": {
      "Set-Cookie": "MAKACSESSIO
    }
  },
  "url": "Powered by s+(?<CERN
  cern.ch/indico/indico-ide
  .org/cern.ch/indico/ide
  indico/indico.org/ver
  icon": "indico.org",
  "website": "indico-software.o
  
```

# An educational distributed Cosmic Ray based on ArduSiPM using microcontroller acquisition node NTP protocol as time technology for data aggregation

Valerio Bocci, Giacomo Chiodi, Paolo Fresch, Francesco...


INFN Roma, Piazzale A.Moro, 2 - 00185 f...  
valerio.bocci@roma1.infn.it

The advent of microcontrollers with enough CPU power and with analog and digital peripherals gives the pass in one chip. The existence of a world wide data infrastructure as internet allows to think at distributed net data or respond to settings commands. The internet infrastructure allow us to do things unthinkable a few with tens of milliseconds precision to simple devices far apart from a few meters to thousands of kilometer perform using simple detectors.



The ArduSiPM [1] is an easy hand-held battery operated data acquisition system based with an Arduino board, which is used to detect cosmic rays and nuclear radiation. The ArduSiPM uses an Arduino DUE (an open Software/Hardware board based on an ARM Cortex-M3 microcontroller) as processor board and a piggyback custom designed board (shield), these are controlled by custom developed software and interface. The Shield contains different electronics features both to monitor, to set and to acquire the SiPM signal using the microcontroller. The SiPM photon counting detector can be coupled to a cheap plastic scintillator to realize a cosmic ray detector (mainly muon particles). An ArduSiPM channel give informations about rate of events, arrival time and number of photons produced by muons. It contains all the features from controls to data acquisition typical of High Energy Physics channel at a cost affordable for single user or school.

[1] The ArduSiPM a compact transportable Software/Hardware Data Acquisition system for SiPM detector V. Bocci et al. IEEE NSIC 2014 Roma, 1-4, DOI:10.1109/NSIC.2014.701132 arXiv:1401.7814



Using Network Time Protocol (NTP) the absolute time of the network, with a precision of tens milliseconds. The network time can be used from a cloud of ArduSiPMs to detect offline coincidence events linked to Ultra High Energy Cosmic Ray

The ArduSiPM sends data... send them using the MQTT protocol to the server. We use as network processor the ESP8266 a low-cost Wi-Fi chip with full TCP/IP stack and a 32-bit RISC CPU running at 80 MHz. The ESP8266 can be used to send and configure MQTT packets, NTP request and configure ArduSiPM device.

ArduSiPM web configuration pages.

ArduSiPM Home Page

ArduSiPM Web Configuration page

ArduSiPM MQTT settings

ArduSiPM Creator Monitor



# Update on CERN Search: SharePoint 2013

E. Alvarez, S. Fernandez, A. Lossent, I. Posada, A. Wagner [CERN]

CERN's enterprise Search solution CERN Search provides a central search solution for both users and CERN service providers. Public and protected documents from a wide range of documents are indexed and available for retrieval.

High Availability	Improved design	High Reliability
2000 daily queries	33 million documents	1100 daily unique visitors
29 servers	12000 daily page views	3 server farms

Content Sources: CDS, TWiki, Drupal, EDMS

Push model, Access control for each source.

Pull model, Dynamic URI list.

Stores data in protected space. Extracts metadata from binary files.

Sets ACL for each document.

Analyze and split documents in several stages. Custom processing. Different processing by content source. Dynamic property mapper. Property mapping by configuration. Auto load-balanced CEWS call.

Microsoft SharePoint 2013 search engine.

Redundancy of key components.

Staging farm as high availability failover.

Frontend: Prod, Stag

<5 seconds switch.

Search Engine: A, B, C, D, E, F, G, H, I, J, K, L, M, N, O, P, Q, R, S, T, U, V, W, X, Y, Z

Search: Frontend, REST

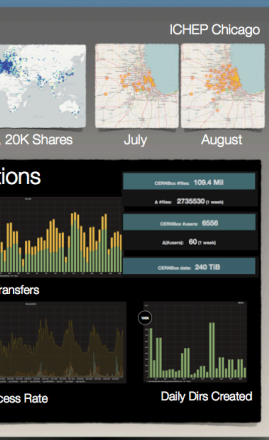
User interface upgraded. Improved user experience. Thesaurus for acronyms. Promoted CERN most relevant sites. Map integration for locations.

Integrated integration as search backend.

https://cern.ch/it © CERN CC-BY-SA 4.0

https://search.cern.ch

https://home.cern



Team: M. Lamanna, L. Mascetti, P. Mato, J.T. Moscioldi, A.J. PETERS, D. Piparo, E. Tejedor

# Big Data Repositories at the Fingertips

Sync, EOS Open Source Storage, XRootD

Shared File Spaces for All Platforms

Apple, Android, Windows, iOS, Linux

# Hub for analysis, shared file spaces and more

## CERNBox

CERN IT-IST  
Contact: jakub.moscioldi@cern.ch

# Data analysis storage backend

Integration Architecture: Clients, REST API, SWAN, Spark, EOS

# Education & Research Community

Cloud Services for Synchronisation and Sharing (CSS)

Novel applications, cloud storage technology, collaborations

SURFSara Amsterdam  
30 Jan - 1 Feb 2017  
https://cs3.surfsara.nl

PROGRAMME COMMITTEE: Guido Allen, Massimo Lamanna, Luca Mascetti, Jacek Moscioldi, Jacek Peters, Tommaso Piparo, E. Tejedor

# Federation

Multi-vendor API to connect on-premise clouds and sync/share solutions.

To enable secure, open, frictionless file sharing everywhere



Sebastian Łopieński / CERN Computer Security Team  
e-mail: Sebastian.Lopienki@cern.ch



# An educational distributed Cosmic Ray based on ArduSiPM using microcontroller acquisition node NTP protocol as time



data aggregation, Francesco

Moro, 2 - 00185 Roma, Italy  
atromal.infn.it



The terms of IoT (Internet of Things) protocols and the apps communicate using the MQTT (Message Queue Telemetry Transport) in order to be easily in nowadays microcontroller nodes also in credit card.

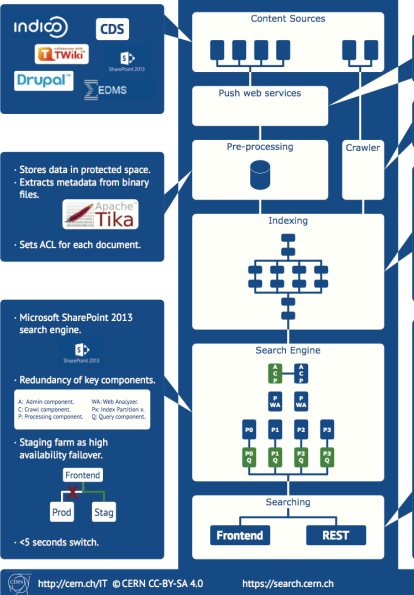


The ArduSiPM sends data to the server using the MQTT protocol to the server. We use as network processor the ESP8266 low-cost Wi-Fi chip with full TCP/IP stack and a 32-bit RISC CPU running at 80 MHz. The ESP8266 can be used to send and configure MQTT packets, NTP request and configure ArduSiPM device.

## ArduSiPM web configuration pages.



High Availability	Improved design
2000 daily queries	33 million documents
29 servers	12000 daily page views



## Future plans

- Reduce printer numbers: find solution for areas with more proactive support extend to printers? Identify (extend leaving to printers? pay per page, but minimum)
  - Remove printers from areas that hit projects well
  - Reduce admin overhead costs as no need to order if
  - Easier to update models when required
  - Automatically add to machines with proactive supp (enable "my favourite printers" on Windows)
- Dev dedicated: enabling printer configuration on v

OCERN, CC BY-SA 4.0

## Migrating the Belle II Collaborative Services and Tools

A. Gellrich\*, D. Jahnke-Zumbusch\*, D. Knittel\*, P. v. d. Reest\*, B. Venenmann (DESY IT), N. Braun†, D. Dossett†, O. Frost†, T. Hauth†, J. Grygler†, T. Kuhl†, L. Li†, N. Nakao†, M. Prim†, F. Schwennsen†, P. Urquijo† (Belle II)


(DESY, KEK, KIT, LMU, USTC, U Melbourne)

### Introduction

Collaborative services and tools are essential for any (HEP) experiment involving a large number of international partner. They help to integrate global virtual communities by allowing all members to share and exchange relevant information by way of web-based services.

### Motivation

In order to achieve stable and reliable services, the Belle II collaboration decided to migrate the current set of services into the existing IT infrastructure at DESY.



- ### Belle II at SuperKEKB
- From 1998 to 2010, KEK, the Japanese High-Energy Accelerator Research Organization, operated KEKB, a first-generation asymmetric electron-positron collider facility tracking the world record in instantaneous luminosity of 2.5x10^31 cm^-2. The Belle averages were chosen so that in the colliders large numbers of 10-15 million pairs were produced, and hence Belle II is also known as a B factory.
- ### Belle II Collaboration
- 960 countries
  - 100 institutes
  - 2 continents
  - 4 continents (Asia, Europe, North America)
  - 30 institutes (DESY, KEK, USTC, LMU, KIT, U Melbourne)

### User Registration

- Advanced security due to personal accounts
- Authorized access only (no common account)
- All Belle II members request credentials (account/password)
- One account for all services (Single-Sign-On (SSO))
- First certified based authentication
- Form based (paper) registration
- Belle II membership based authorization (CMS)
- Solutions for students or co-workers outside Belle II
- First Last@belle.org
- Multiple groups (primary / secondary)
- Identity and Access Management (IAM) is planned
- Status: ~450 registered Belle II members (of 681)

### Procedure and Experiences

- Migration is a technical and social challenge!
- KEK CS shutdown in August 2016 set the timeline
- Main goal: Don't lose any information!
- Integrate services and tools into an existing IT infrastructure
- Technology changes are necessary
- 1 to 1 copy of contents is usually not possible
- Migration tools are not always available
- Quite some manual work required
- ~10 Expert teams taking care of the various services
- Cooperation of collaboration members needed
- Wiki by most complicated
- Migration from SVN to Git most controversial within Belle II
- Maintaining infrastructure vs. maintaining contents
- For the migration: ~10 FTE / 6 month
- Further maintenance: ~1 FTE, on the last 2 years!

<b>Website</b> belle2.org	<b>Wiki</b> X.Confluence	<b>Issue Tracking</b> JIRA Software	<b>Mailing Lists</b> symfony	<b>Code Repository / Build Services</b> Stash / Bamboo	<b>Agenda Service</b> Indico	<b>Document Service</b> Invenio
Public / Internal pages (SSO)	ATLASSIAN tool Single Sign-On (SSO) State of the art wiki Connection to jira & stash Single-Sign-On (SSO)	Multiple groups (primary / secondary) Identity and Access Management (IAM) is planned Status: ~450 registered Belle II members (of 681)	Multiple groups (primary / secondary)	ATLASSIAN tools Single Sign-On (SSO) Connected to confluence	Invenio Commonly used Separate accounts Migration from Invenio Copy from KEK Invenio Selected categories only	Invenio Single Sign-On (SSO) Migration from Invenio Straight forward
Modern responsive design, CMS based, 2x62 (python, jquery, bootstrap)	By far the most work ~4000 page migrated No automatic migration Some content require content 3 IT/IT manual work per bootstrap	Workflows need adoption Archived wiki available	Cleaning and syncing	Virtual machines in Xen Built on VMs Build slaves for OS	No 'yet done'	



Stick Man Illustrations by Cath Noble (CERN), WAD















# Using the Detector Final State to describe a Physics Analysis

**Abstract**

How Do You Organize What You Know?

How do you organize what you know? This is a global effort.

How do you organize what you know? This is a global effort.

**What Would It Take to Build a Search Engine With Direct Knowledge of the Paper Content?**

Search the Archive (arXiv.org)

Search for the words "top quark" in the abstract

Search for the words "top quark" in the abstract

**What would it take to make this a wild search engine query?**

Search Engine

Search Engine

**The Detector Final State OWL Pattern**

OWL

Ontology

Knowledge

**.cern Timeline**

Timeline

Timeline

**Managing users that**

Managing users that

**.cern Governance**

Governance

Governance

**Status**

Status

Status

**Work supported by the U.S. Department of Energy**

Work supported by the U.S. Department of Energy

# Success with Federated Identities

Dave Dykstra, Mine Altunay, Jery Teheran, Tanya Levshina, Neha Sharma, Dennis Bo, Kenneth Heron, Amanda Gao. Scientific Computing Division, Fermilab, Batavia, IL



## Architecture

### The CMS Data Analysis School experience

Nicola De Filippis  
Dipartimento Interateneo di Fisica "M. Merlin", Politecnico and INFN Bari, Italy

on behalf of the CMS Collaboration  
The CMS Data Analysis School experience

CMAS are the official schools for learning about CMS Data Analysis; they are coordinated by the CMS Schools Committee.

Operations

Half of the first day devoted to plenary lectures on physics, detector and software tools

Short exercises

Citizen science projects

Long exercises

ACA visitors can either click on word bubbles or use keywords to access a set of project specific web sites.

Results from surveys

Participated level (level 5 good results)

Best Analysis Team prize

ACA visitors can either click on word bubbles or use keywords to access a set of project specific web sites.

Conclusions & Future work

WLCG Traceability & Isolation Working Group will continue:

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

# Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Search: P. S. Lossent, J. S. D.

Science & Technology Facilities Council

Developing the Traceability Model to meet the Requirements of an Evolving Distributed Computing Infrastructure

Introduction

Security incidents are an operational reality in distributed infrastructure.

Potential solution: Singularity

Containers are an operational reality in distributed infrastructure.

Simplifying site requirements

Containers would not need special UID switch.

Current GExec model

GExec manages authentication, authorization & isolation and logging.

Proposed execution and traceability model

Proposed execution and traceability model

Separating Isolation & Traceability

Containers (namespaces) can provide isolation between processes (jobs).

Proposed incident response workflow

Proposed incident response workflow

Conclusions & Future work

WLCG Traceability & Isolation Working Group will continue:

Work supported by the U.S. Department of Energy

Work supported by the U.S. Department of Energy

Fingertips

Sync

Storage

Platforms

More

Community

Sharing (CS3)

Collaborations

Summit

Massimo Lattuada (CERN)

Janus Moscardi (CERN)

Christoph Dittmann (CERN)

Innovation

Vendor API to protect on-premise sites and sync/solutions.









See you in 2018