

# Database ecosystem at RACF/SDCC

---

CARLOS FERNANDO GAMBOA

ERIC LANCON

MAGNET TEST DATABASE WORKSHOP

# Scientific Data and Computing Center



- Service Operation For:

RHIC, LHC ATLAS, BER ARM, LQCD, RIKEN, BES Center for Functional Nano Materials, National Synchrotron Light Source II, National Nuclear Data Center, Simons Foundation,...

- ~1500 users from 20 projects (<10 to 100+ users/project)

**24/7 availability**

# SDCC and large experiments

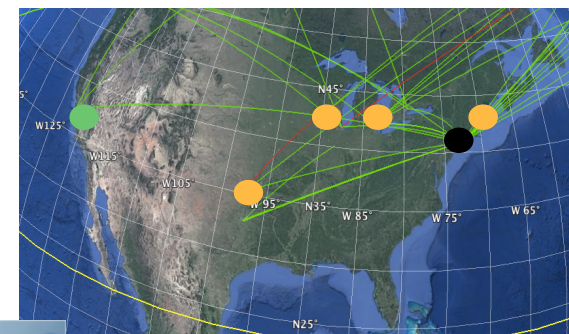

## ◉ The RHIC Tier 0

- Store and process data from RHIC experiments
- Provide analysis means for 1'200 users
- Long term data preservation
- Simulation resources for future programs (sPHENIX & EIC)

## ◉ The US ATLAS Tier 1

- ~25% of ATLAS Tier 1 computing capacity worldwide
- Store RAW data from LHC and from simulation
- Distribute data to the 4 US Tier 2 sites + analysis site (SLAC)
- **Analysis center for US physicists**
  - From 41 institutes (incl. 4 Nat. Labs)
  - 600 physicists, 190 PhDs

## ◉ The Belle II data center outside Japan




**BROOKHAVEN**  
NATIONAL LABORATORY

Scientific Data and  
Computing Center

# SDCC in numbers

- 90+k CPU cores — 4 PFlops
  - 3 HPC Institutional Clusters (GPU, KNL, Skylake)
- ~90 PB of disk storage
  - of various technologies
- 130+ PB of tape storage
  - Largest HPSS tape library in the US, 3rd worldwide [1]
- 2x100 Gbps connection to ESnet
  - Onsite ESNet support




[1] <http://www.hpss-collaboration.org/customersT.shtml>



# Databases in the SDCC

---

- ◉ Usability/technology driven by the application
- ◉ Relational databases most popular technology
- ◉ Commercial and open source databases. i.e DB2, MySQL and postgresQL



# SDCC Critical services databases

---

## ◉ Storage

- dCache Storage Element's namespace uses PostgreSQL to store metadata (filenames, file ids, location,...), 432GB size .
- Hosted in a powerful hardware (cpu, memory and storage) to ease application requirements.

## ◉ Tape Storage

- HPSS metadata manage all attributes and usages of files, and tapes, configuration of libraries, tape drives, 110 GB size.
- Commercial technology supported by IBM DB2

# Databases as part of computer infrastructure management

## PostgresSQL

**FOREMAN** Carlos Gamboa

Monitor Hosts Configure Infrastructure Administrator

Puppet classes

Filter ... Search Import environments from PuppetMaster Documentation

Class name	Environments	Host groups	Hosts	Parameters	Variables	Actions
afs	atlasdcache30, belicdb, blcondup, condor, c...		8	1	0	Delete
afs:server	condor, cvdef, cvmfs, development, future, ...		5	1	0	Delete
apache	atlasdcache30, belicdb, blcondup, condor, c...		0	8	0	Delete
apache:params	atlasdcache30, belicdb, blcondup, condor, c...		0	0	0	Delete
apache:ssl	atlasdcache30, belicdb, blcondup, condor, c...		0	6	0	Delete
atlasdcache	atlasdcache30, belicdb, blcondup, condor, c...		1	0	0	Delete
atlasdcache:admin	atlasdcache30, belicdb, blcondup, condor, c...		0	3	0	Delete
atlasdcache:atlas_post	diskpart, hpss, icgpfnsfs, img, linktest, open...		0	0	0	Delete
atlasdcache:atlas_post:conf_client_connect...	diskpart, hpss, icgpfnsfs, img, linktest, open...		0	3	0	Delete
atlasdcache:atlas_post:conf_connection	diskpart, hpss, icgpfnsfs, img, linktest, open...		0	5	0	Delete
atlasdcache:atlas_post:conf_error_handling	diskpart, hpss, icgpfnsfs, img, linktest, open...		0	2	0	Delete
atlasdcache:atlas_post:conf_error_reportin...	diskpart, hpss, icgpfnsfs, img, linktest, open...		0	7	0	Delete

**Foreman and GLPI database size  
10GB/each**

## MySQL

**GLPI** English My settings Help Logout (Carlos Fernando Gamboa)

Assets Assistance Management Tools Plugins Administration Setup Search

Home > Computers

Computers

Domain	Status	Manufacturer	Model	Serial number	Inventory number	Location	Rack enclosures - Name	Operating system	Type	Technician in charge of the hardware	User	Last update	Fuelinv - Last inventory
cc.bnl.gov	belicdb	Dell Inc.	PowerEdge R430	DXP2JL2				Red Hat Enterprise Linux Workstation release 7.4 (Majuro)	Rack Mount Chassis			2016-04-30 13:30	2016-04-30 21:31
cc.bnl.gov	belicdb	Dell Inc.	PowerEdge R430	DXP3JL2				Red Hat Enterprise Linux Workstation release 7.4 (Majuro)	Rack Mount Chassis			2016-04-30 13:33	2016-04-30 20:00

Unmanaged devices  
**BROOKHAVEN NATIONAL LABORATORY**

RT at a glance  
Home - RACF SLA - RACF Tickets - GDS Tickets - OSG GDS Tickets

Bookmarked Tickets 10 highest priority tickets I own 10 newest unowned tickets My Requests

Calendar

Tue	Wed	Thu	Fri	Sat	Sun	Mon
1	2	3	4	5	6	7

Reminders

Quick search

Genre	new	open	stalled
AFS	0	0	0
AtlasODM	0	2	0
BatchSystems	1	2	0
CFN	0	0	0
Databases	0	0	0
GCE	0	3	0
General	0	1	0
GLTier2	0	1	0
GridServices	0	1	0
HPSS	0	1	0
LinuxFarms	0	1	0
NETTier2	0	0	0
NETTier3	0	0	0
Network	0	1	0
NFS	0	0	0
SDCCApplications	0	3	0
SDCCCluster	0	3	0
Software	0	1	0
StorageManagement	2	10	0
SWTier2	1	0	0
Test	0	0	0
TicketSystem	0	0	0
USAtlasSharedT3	0	4	0
USAtlasTier3	1	2	0
UserAccounts	1	2	0

rt ticketing system  
3 GB DB size

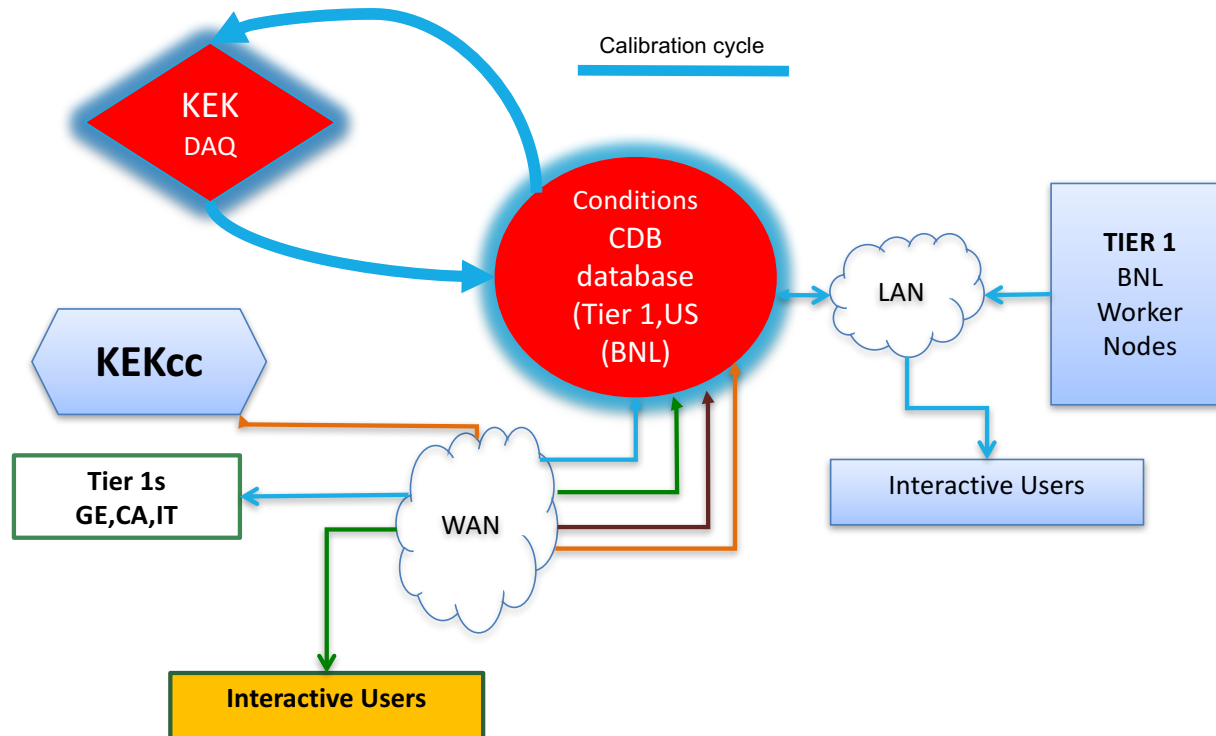
# Databases in the SDCC, as a service for scientific experiments

---

BELLEII CONDITIONS DATABASE



# Belle II CDB Remote service accessibility (WAN/LAN)



9

# Databases in the SDCC

## Belle II Conditions DataBase (CDB)

---

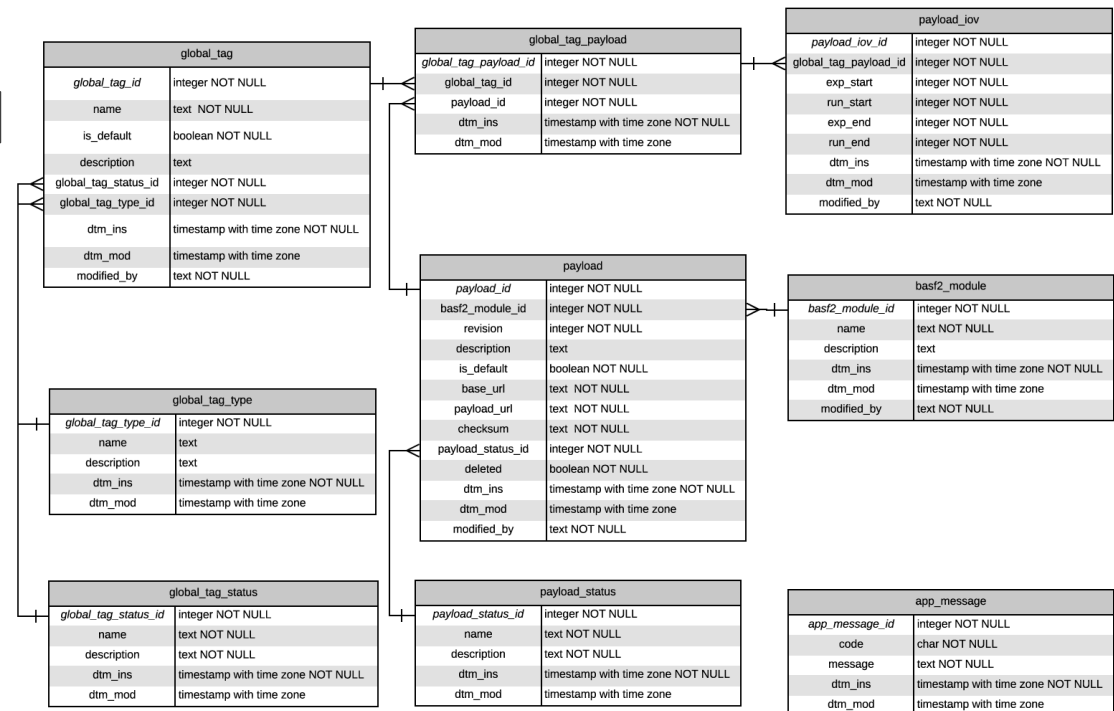
- ◉ Events recorded by the Belle II detector are grouped in *runs*[2].
- ◉ A *run* sets a data taking period with stable operating conditions.
- ◉ The database manages conditions data on run granularity
- ◉ CDB metadata stored/accessed in the database:
  - Allows identify payloads(Payload Binary objects) accessible is an external service.

[2] <http://stacks.iop.org/1742-6596/898/i=4/a=042046>

# Belle II CDB data model relational database

The payloads are defined for an interval of validity (IoV) and grouped into *global tags*.

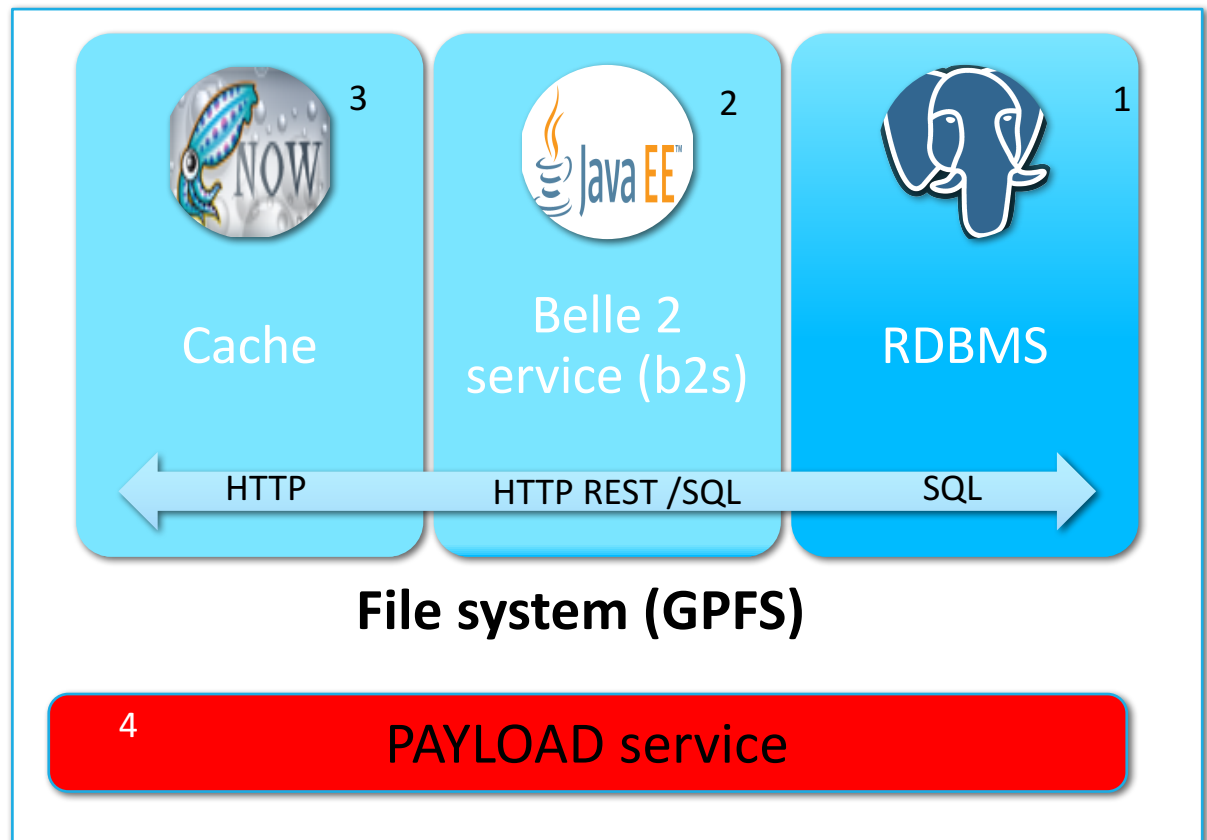
BELLE II CONDITIONS DB SCHEMA



DATE 20170928171657

# Belle II CDB general service architecture

1. Metadata stored in a persistent database repository
2. HTTP REST API presented by b2s
3. Caching technology to offload the b2s/database service
4. Payloads are stored in a shared file system



# Belle II CDB deployed service *belle2db* metadata service

- Architecture deployed using Kubernetes / docker framework
- Database replicated for reliability purposes
- Hardware:

2 Nodes Dell R730xd,

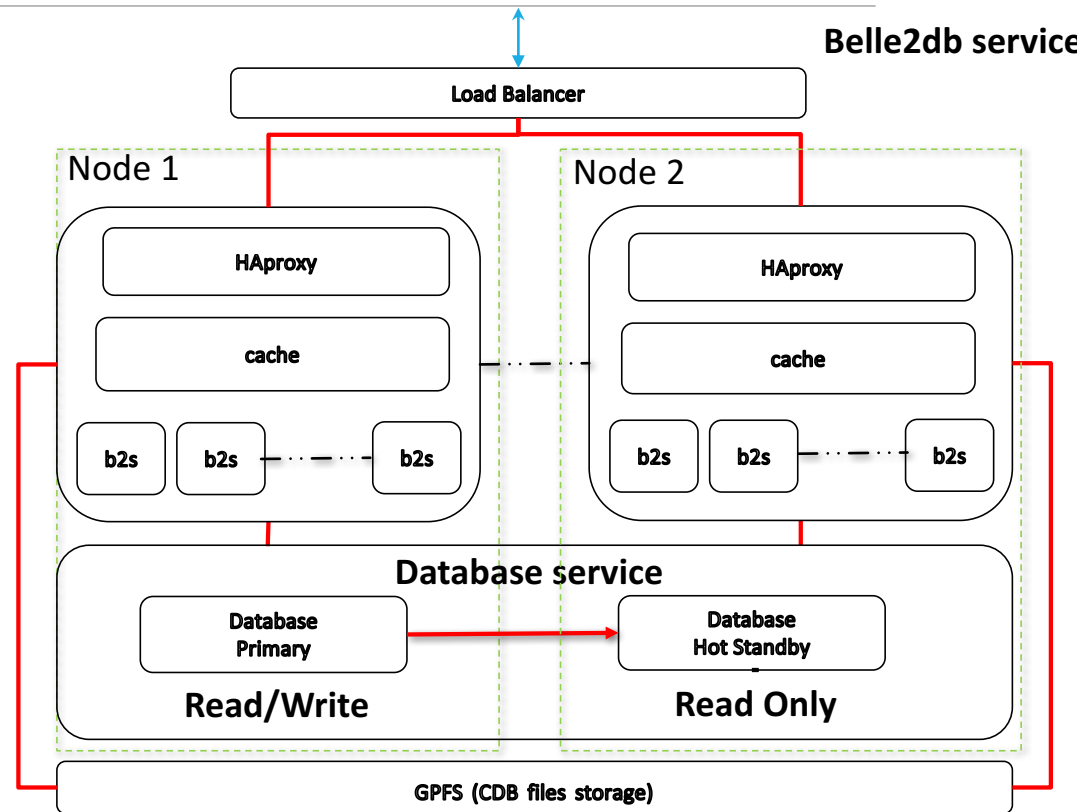
Two Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz -> total CPU thread 32

Memory 256 GB

Disk for Database Intel DC P3700

NVMe SSD on PCI Express 800 GB

20Gb/s Channel Bounded connectivity

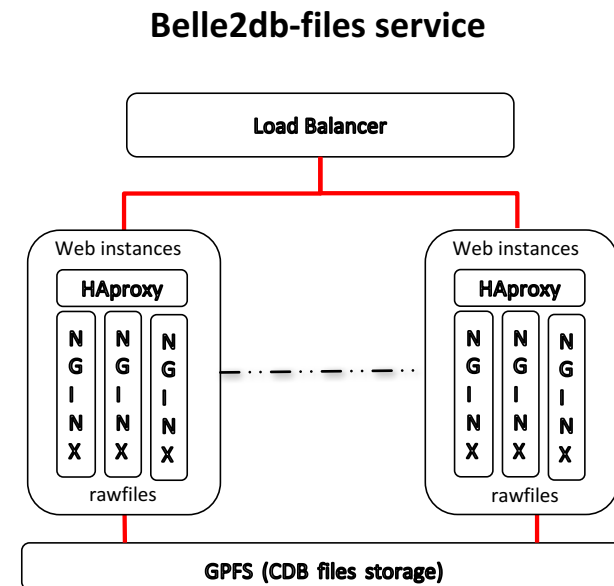


# Belle II CDB deployed service payload service

- General Parallel File System (GPFS) used to store CDB payload

## Hardware

- 2 Nodes
- Dell R430  
Two Intel(R) Xeon(R) CPU E5-2650 v4
- @ 2.20GHz total CPU threads 48  
Memory 256GB
- 20Gb/s Channel Bounded connectivity





# Summary

---

- ◉ Overview of the SDCC database services presented.
- ◉ SDCC expertise in deploying and operating different database technologies to host metadata/data critical for application performance and resiliency.
- ◉ SDCC supports databases core for HEP scientific experiment operation.