Outlook for Accelerator Databases

…for Runs 3 & 4

Chris Roderick, BE-CO-DS

with inputs and content from:

V. Baggiolini, L. Burdzanowski, R. Gorbonosov, P. Le Roux,
M. Peryt, N. Tsvetkov, J. Wozniak

Database Futures Workshop 2017, 29-30 May 2017
https://indico.cern.ch/event/615499/

Accelerator Databases & CERN Installations

Accelerator Databases are directly tied to the physical accelerator installations (including supporting infrastructure)

and the way they are operated

Both of which change relatively slowly….

CERN Installations Change Slowly

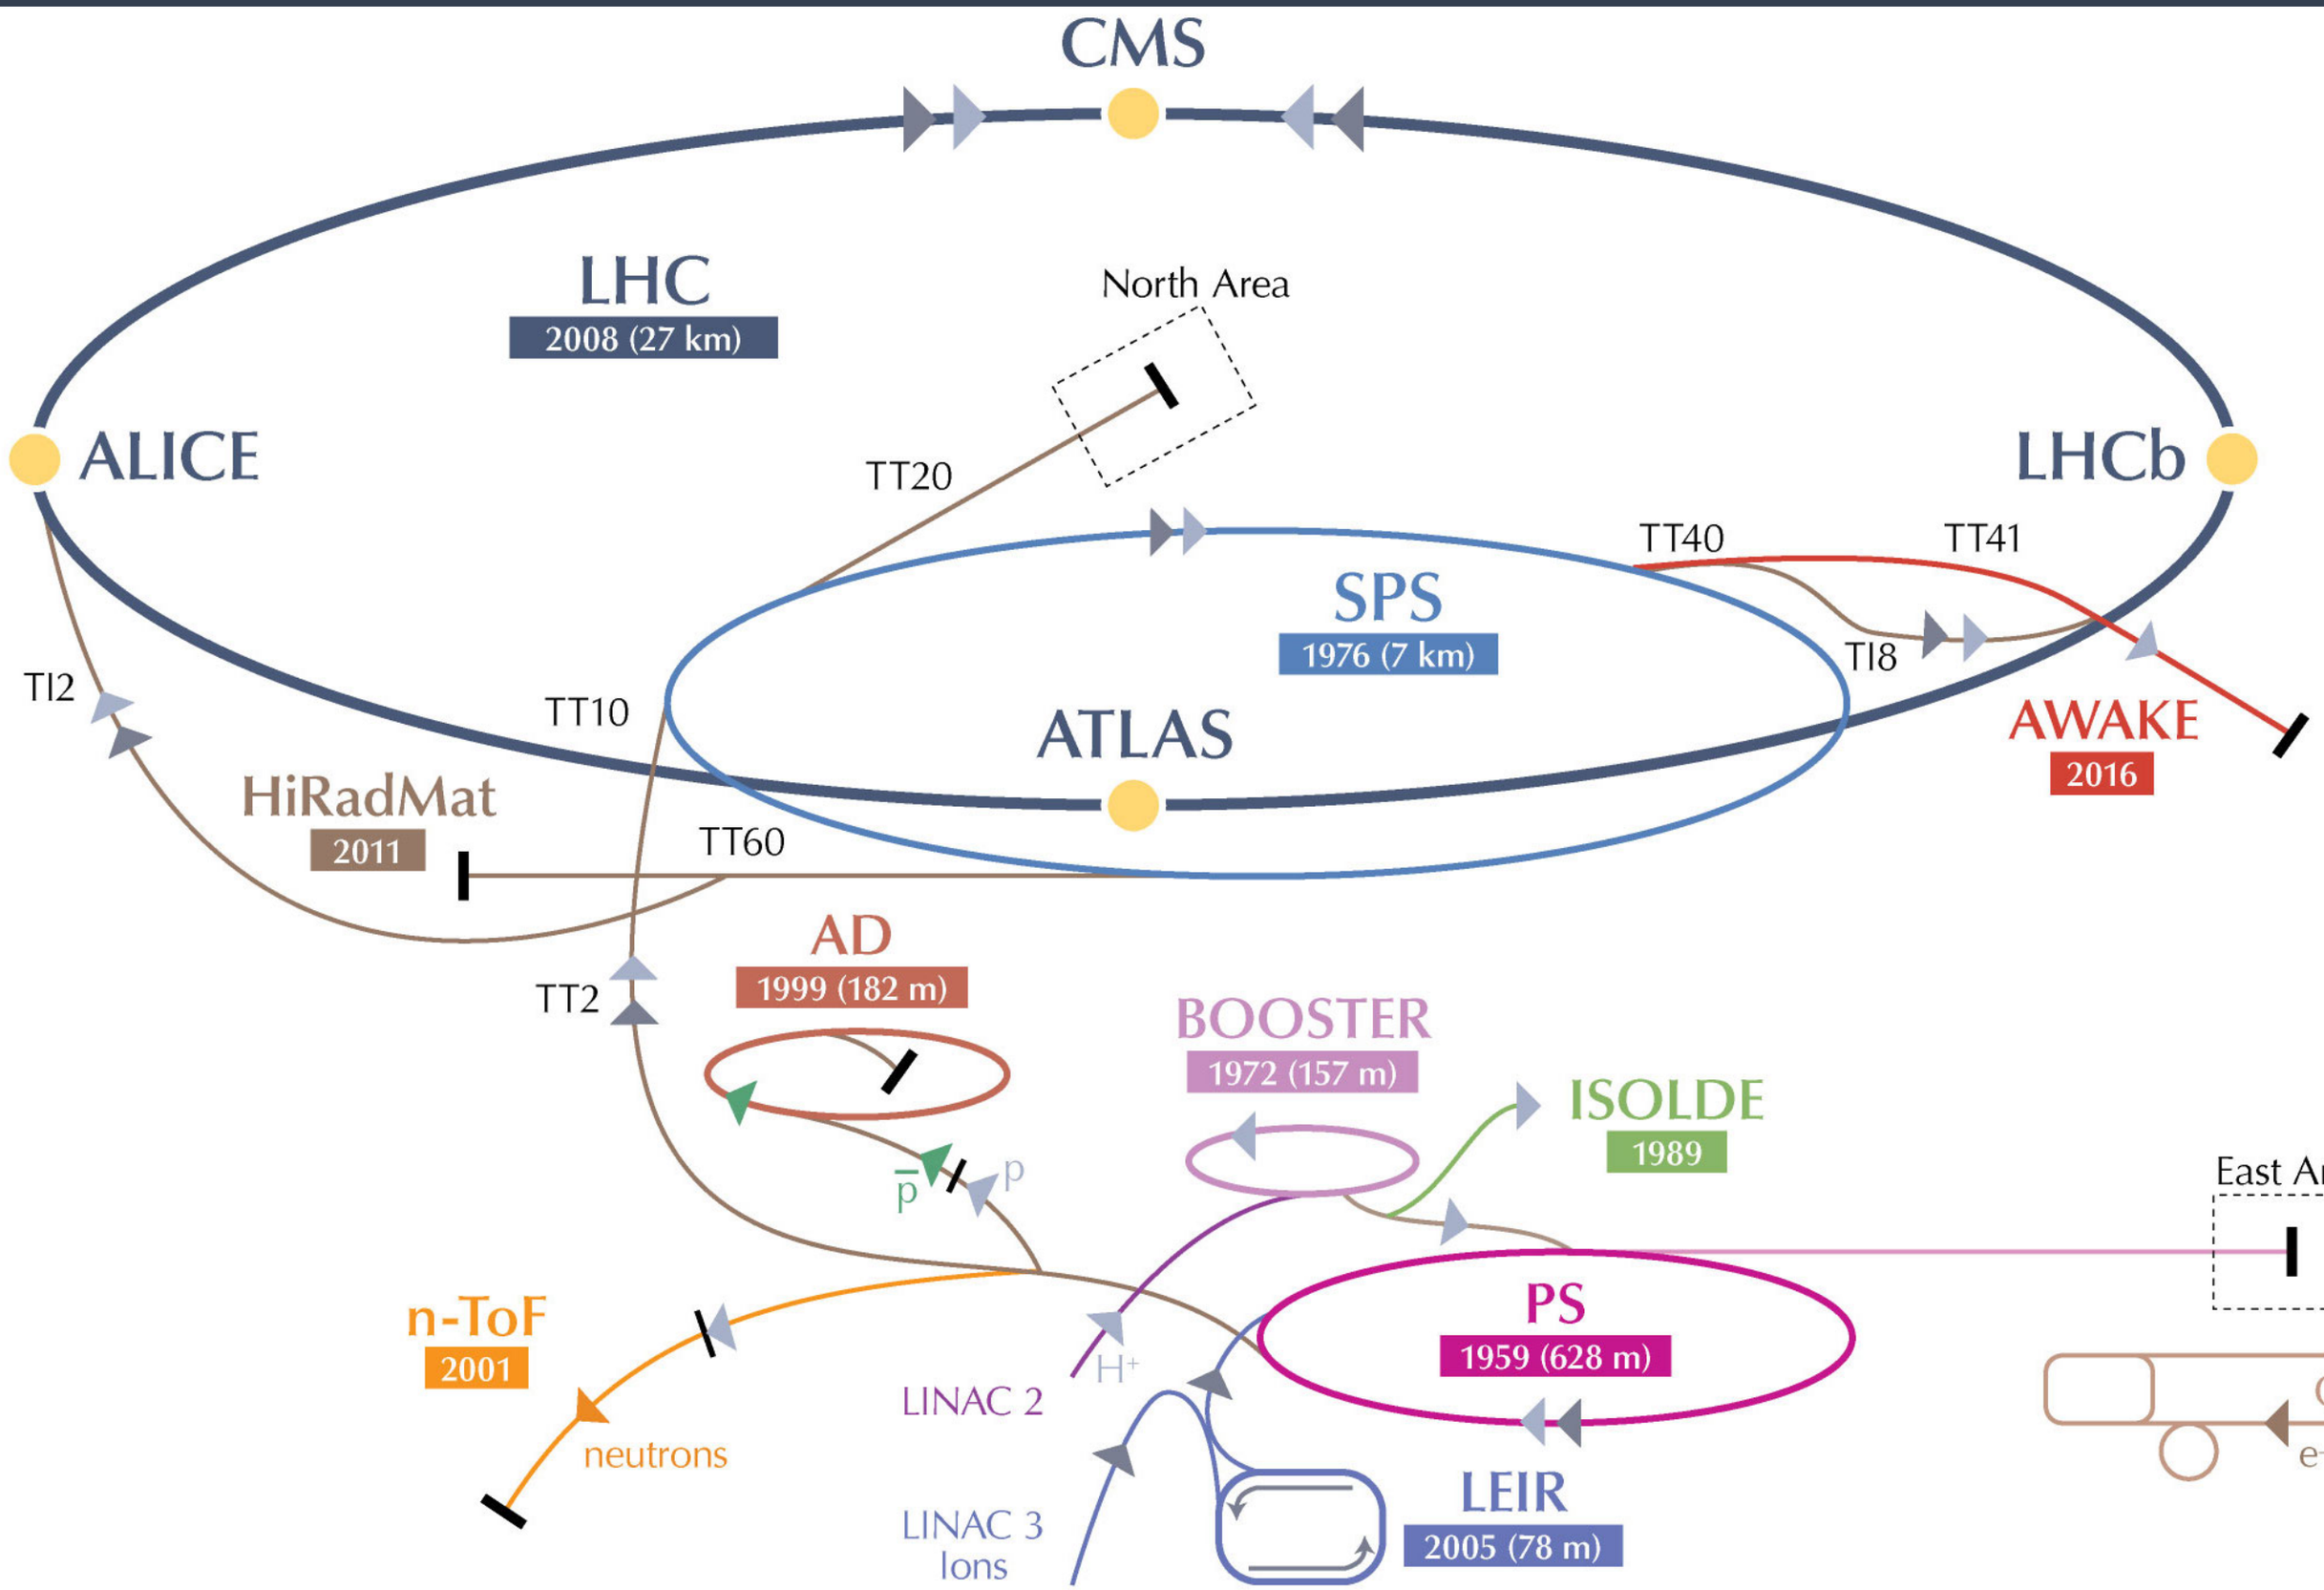

CERN Installations Change Slowly

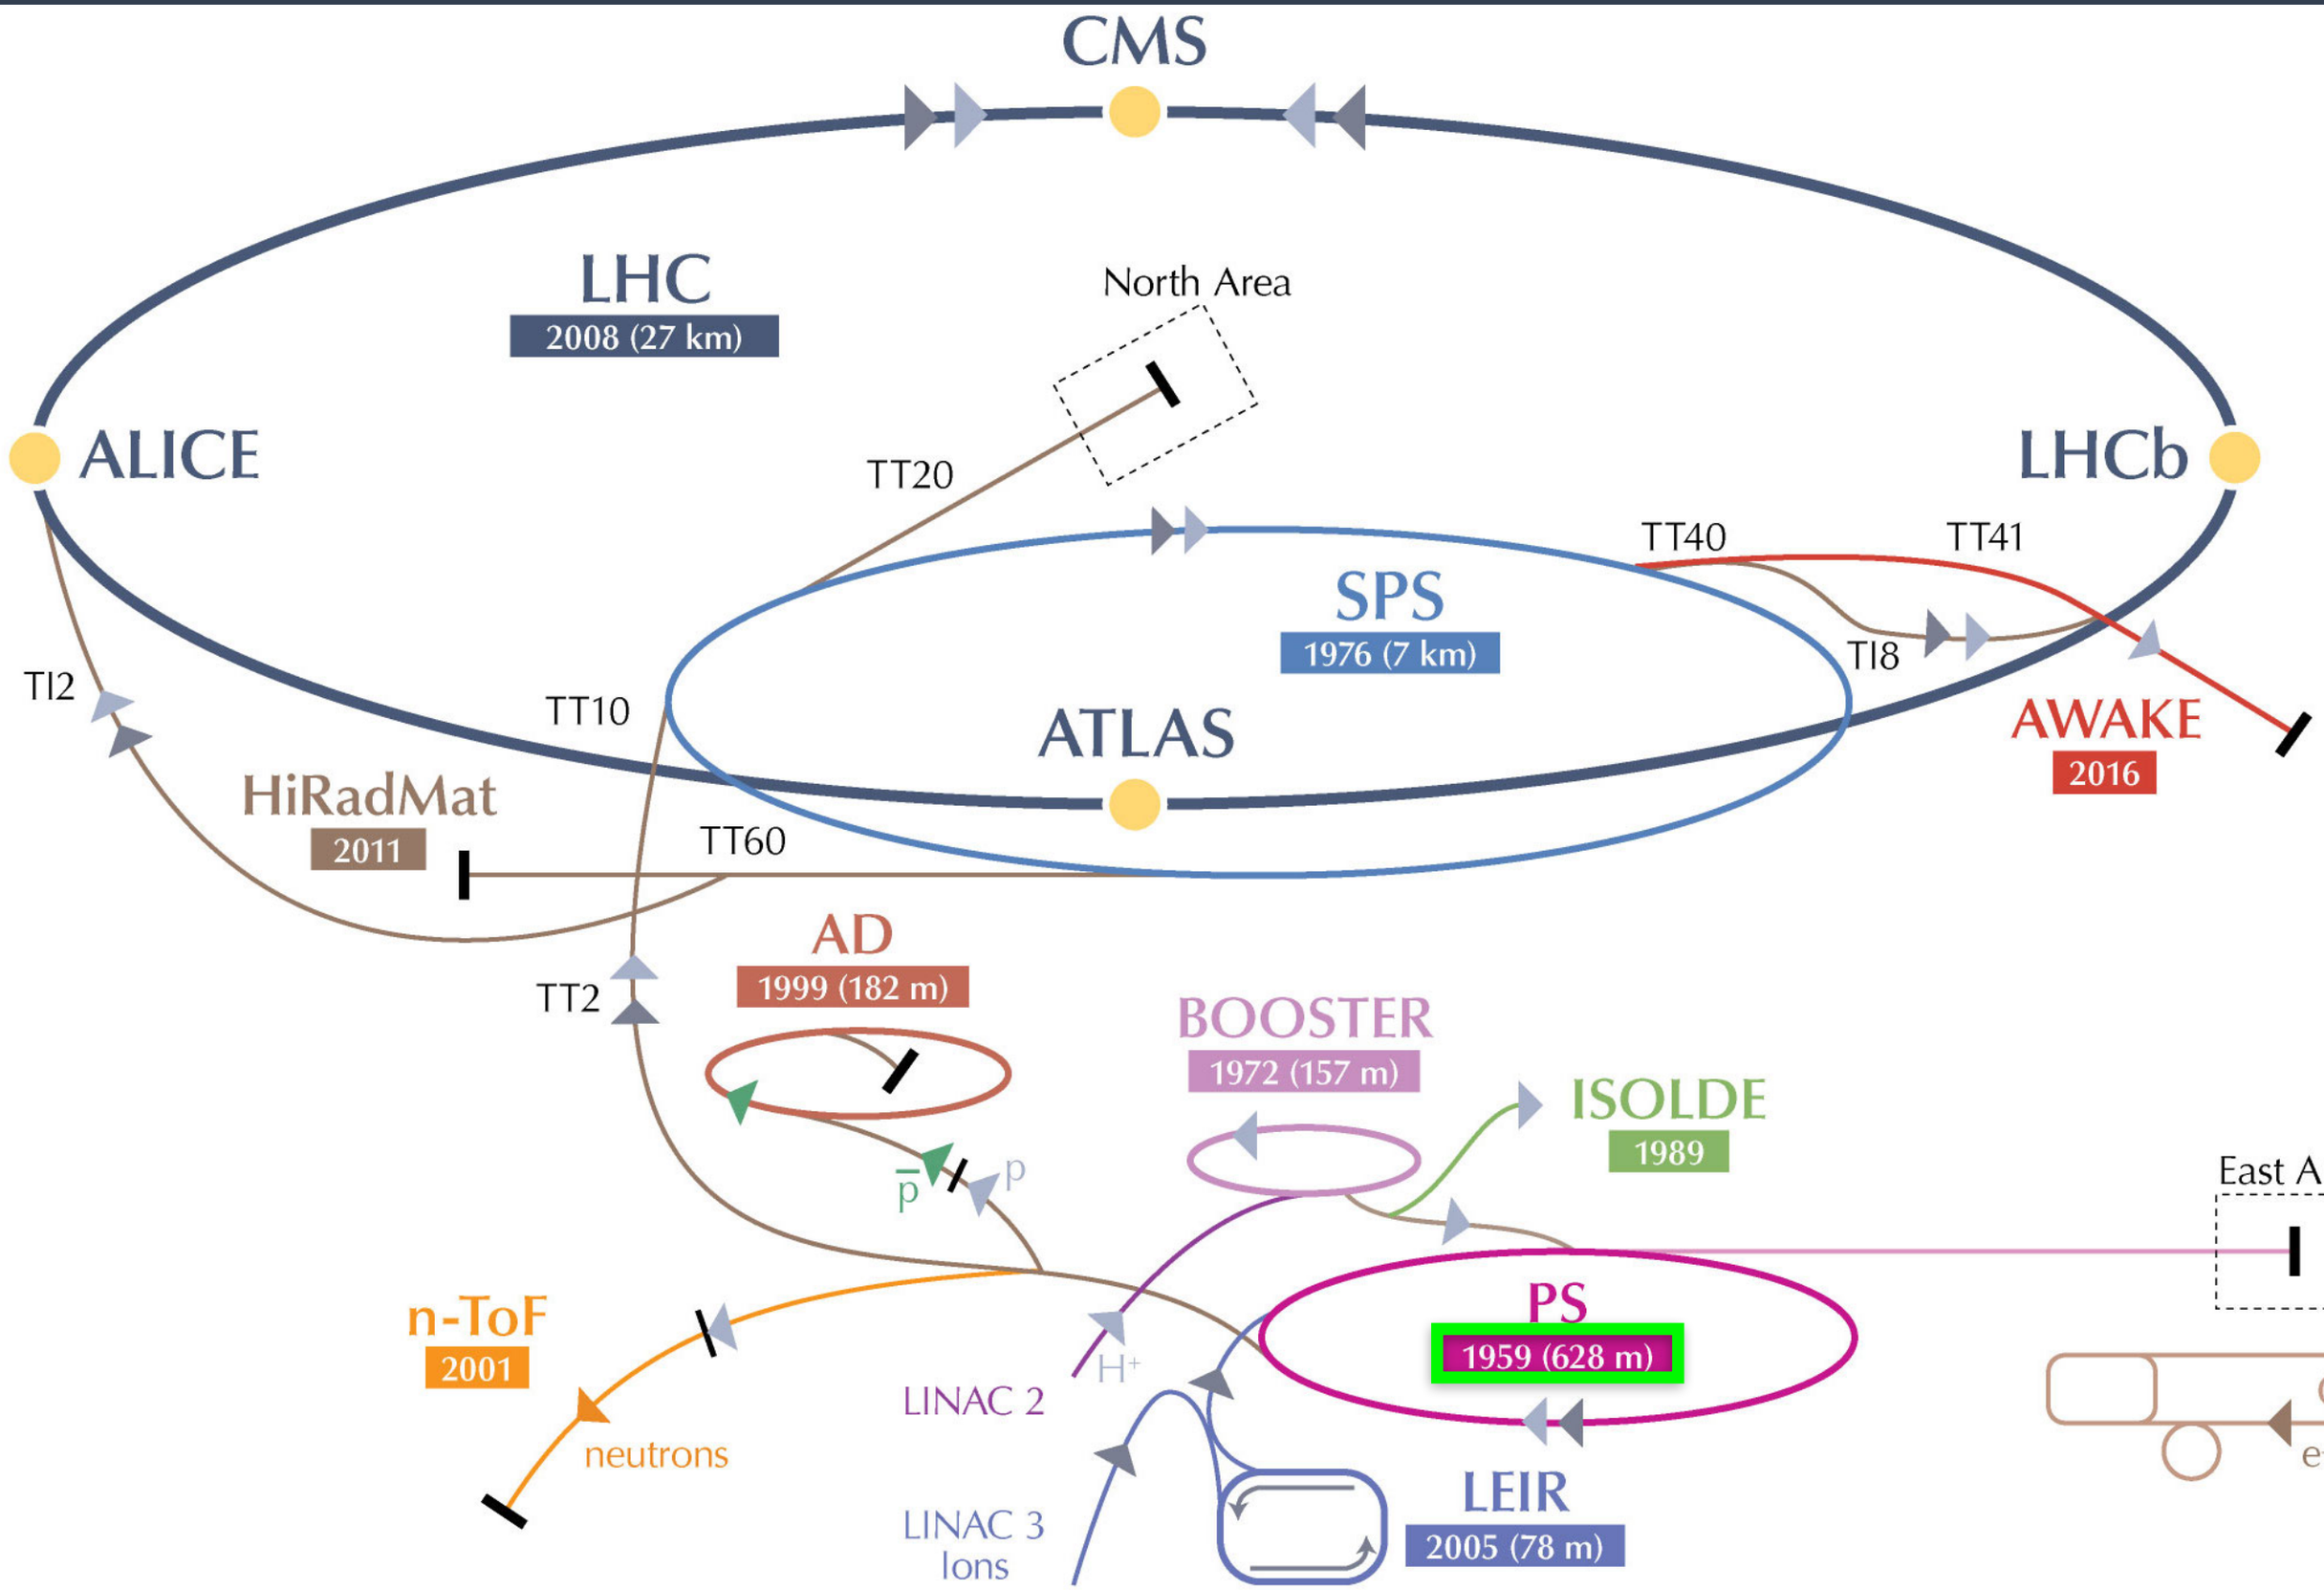

CERN Installations Change Slowly

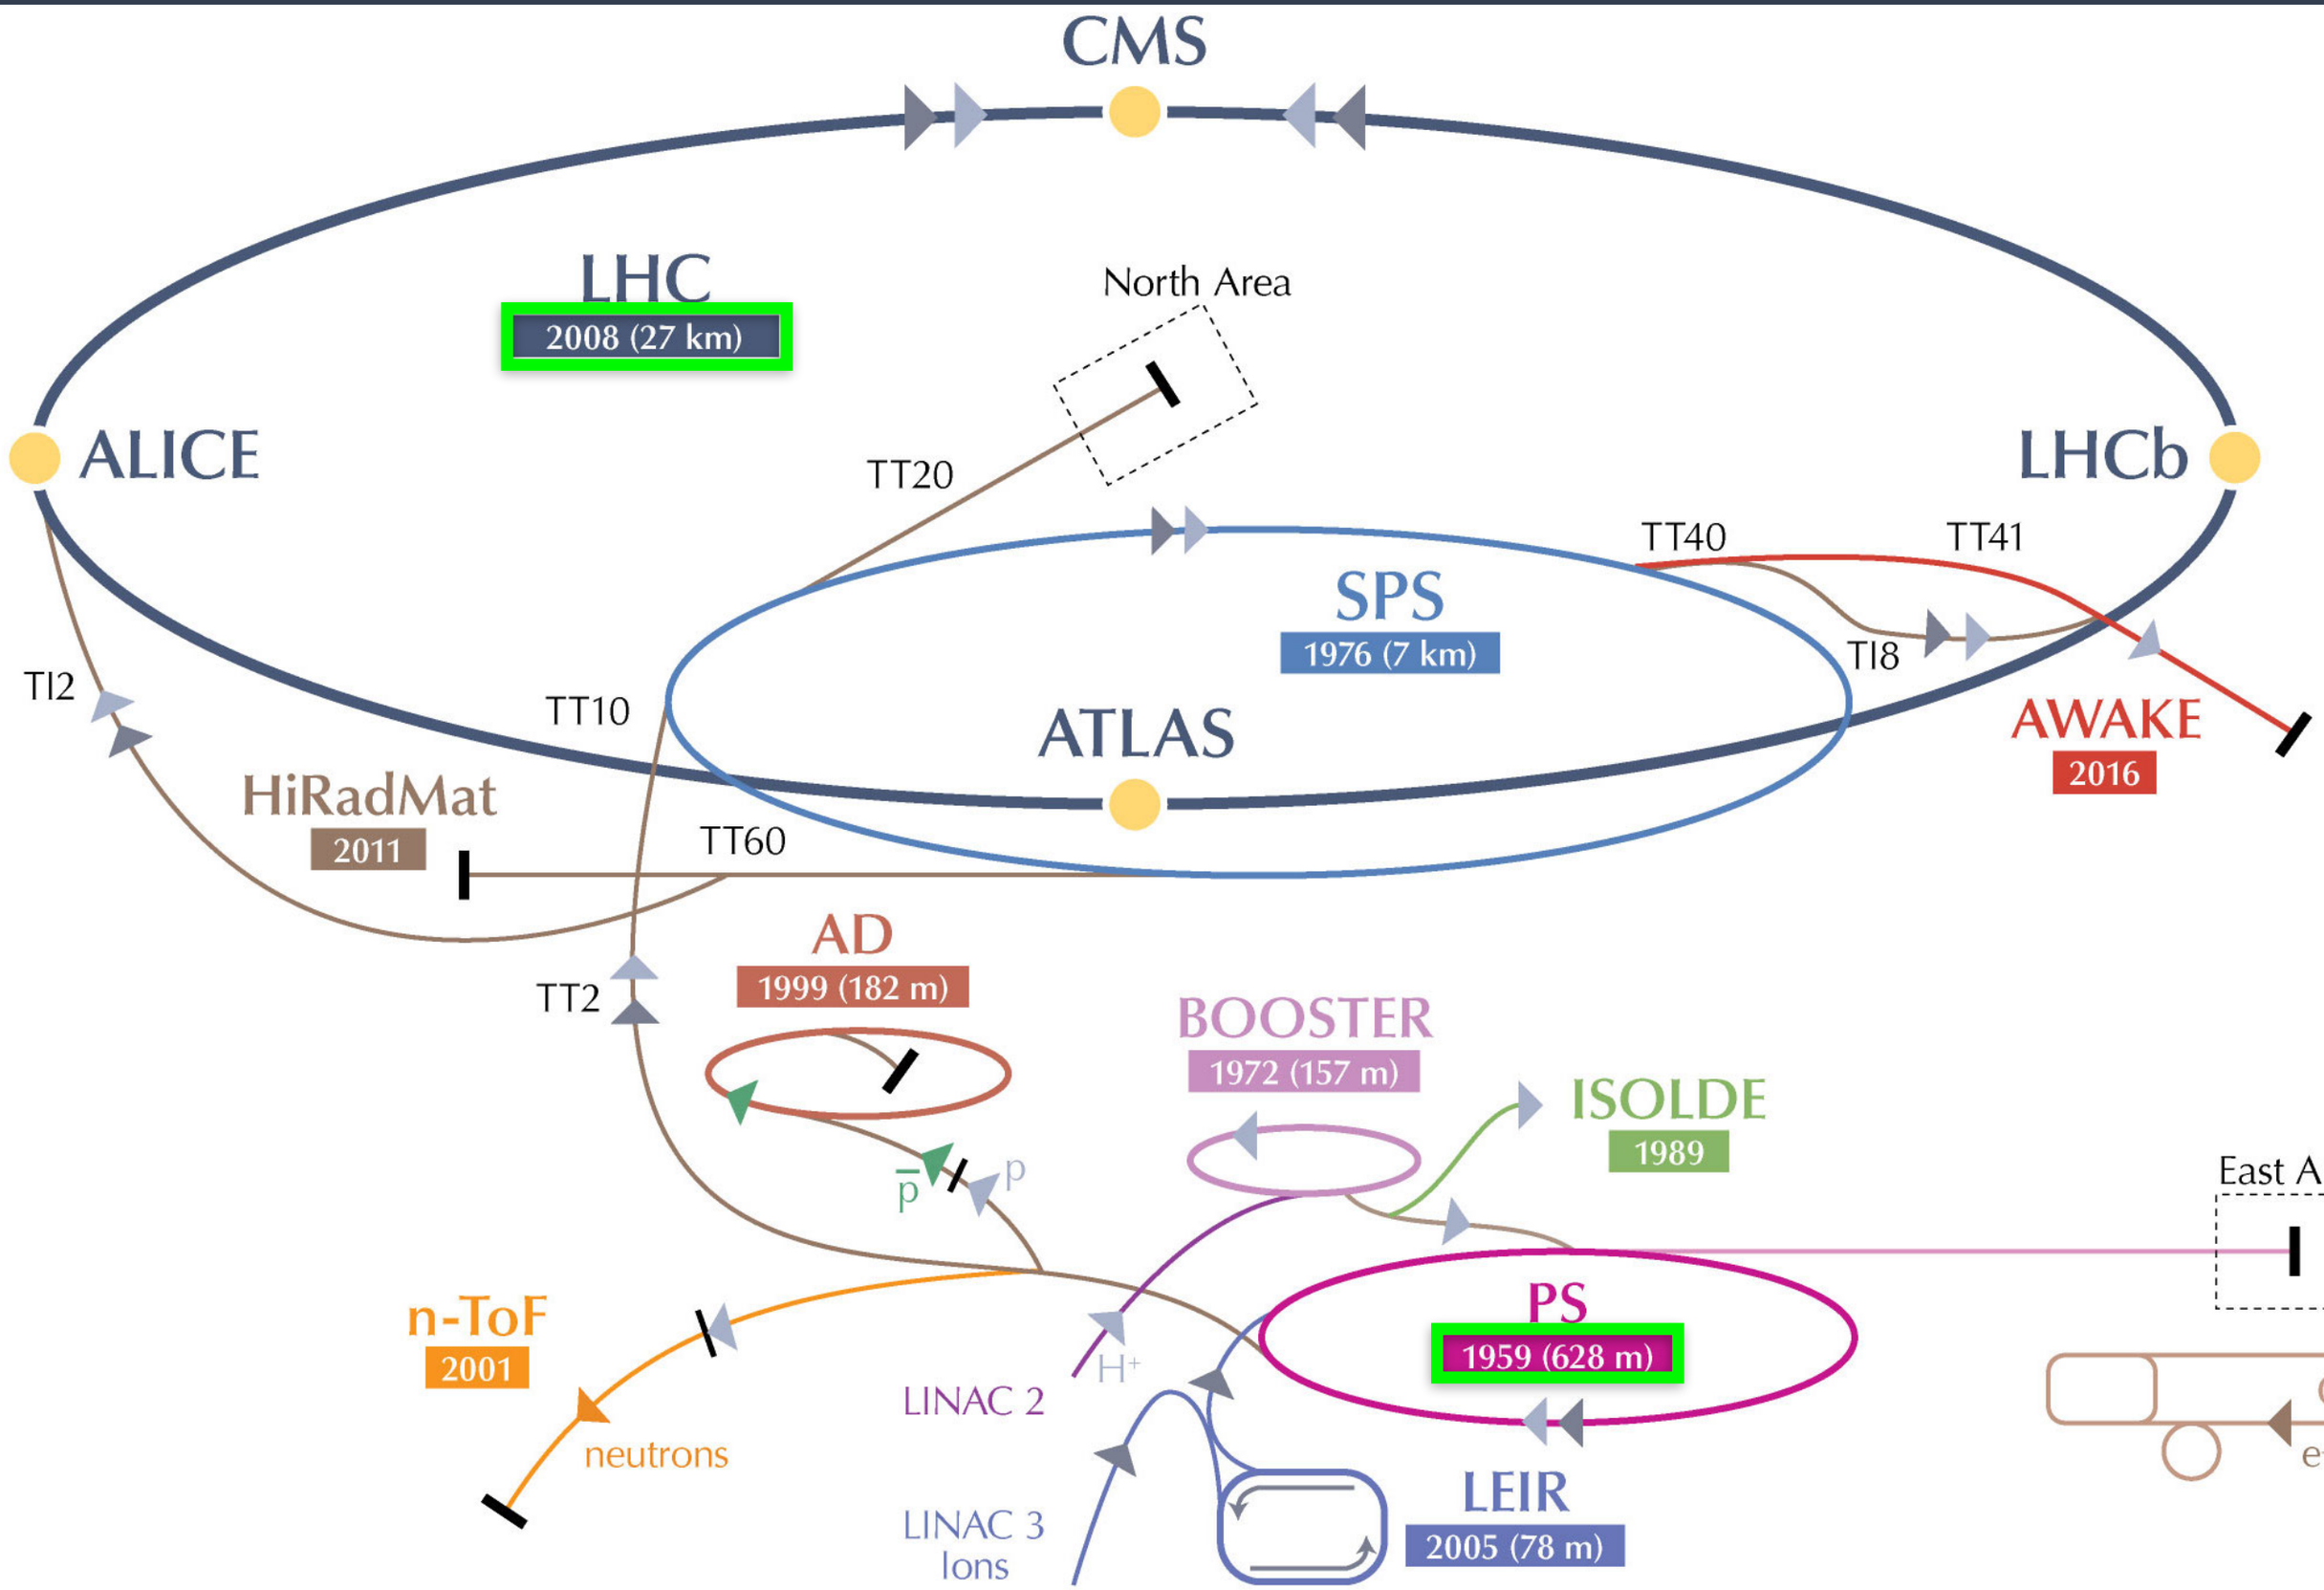

What is on the Horizon?

What is on the Horizon?

Run 3 (2021-2024) & Run 4 (2026-2030)

What is on the Horizon?

Run 3 (2021-2024) & Run 4 (2026-2030)

LIU - LHC Injectors Upgrade (Run 3)

“Delivering reliably to the LHC the beams required for reaching the goals of the HL-LHC. This includes LINAC4, the PS booster, the PS, the SPS, as well as the heavy ion chain.”

What is on the Horizon?

Run 3 (2021-2024) & Run 4 (2026-2030)

LIU - LHC Injectors Upgrade (Run 3)

“Delivering reliably to the LHC the beams required for reaching the goals of the HL-LHC. This includes LINAC4, the PS booster, the PS, the SPS, as well as the heavy ion chain.”

HL-LHC - High Luminosity LHC (Run 4)

“A major upgrade of the LHC to extend its discovery potential, increasing its luminosity (rate of collisions) by a factor of 10 beyond the original design value”.

What is on the Horizon?

Run 3 (2021-2024) & Run 4 (2026-2030)

LIU - LHC Injectors Upgrade (Run 3)

“Delivering reliably to the LHC the beams required for reaching the goals of the HL-LHC. This includes LINAC4, the PS booster, the PS, the SPS, as well as the heavy ion chain.”

HL-LHC - High Luminosity LHC (Run 4)

“A major upgrade of the LHC to extend its discovery potential, increasing its luminosity (rate of collisions) by a factor of 10 beyond the original design value”.

FCC - Future Circular Collider (study phase)

What is on the Horizon?

Run 3 (2021-2024) & Run 4 (2026-2030)

LIU - LHC Injectors Upgrade (Run 3)

“Delivering reliably to the LHC the beams required for reaching the goals of the HL-LHC. This includes LINAC4, the PS booster, the PS, the SPS, as well as the heavy ion chain.”

HL-LHC - High Luminosity LHC (Run 4)

“A major upgrade of the LHC to extend its discovery potential, increasing its luminosity (rate of collisions) by a factor of 10 beyond the original design value”.

FCC - Future Circular Collider (study phase)

Database “Game Changers” or Business as Usual?

What is Business as Usual? - Core Accelerator Database Services

What is Business as Usual? - Core Accelerator Database Services

Layout Service

Controls Configuration Service

LSA Accelerator Settings Management

CERN Accelerator Logging Service

Layout Service Database

The Layout Service documents the accelerator topographical organisation by defining components (Functional Positions) & relationships between them:

- Mechanical assemblies (Physical structures)
- Powering connections & circuits
- Control connections & circuits
- MAD optic sequences

Layout Service Database

The Layout Service documents the accelerator topographical organisation by defining components (Functional Positions) & relationships between them:

- Mechanical assemblies (Physical structures)
- Powering connections & circuits
- Control connections & circuits
- MAD optic sequences

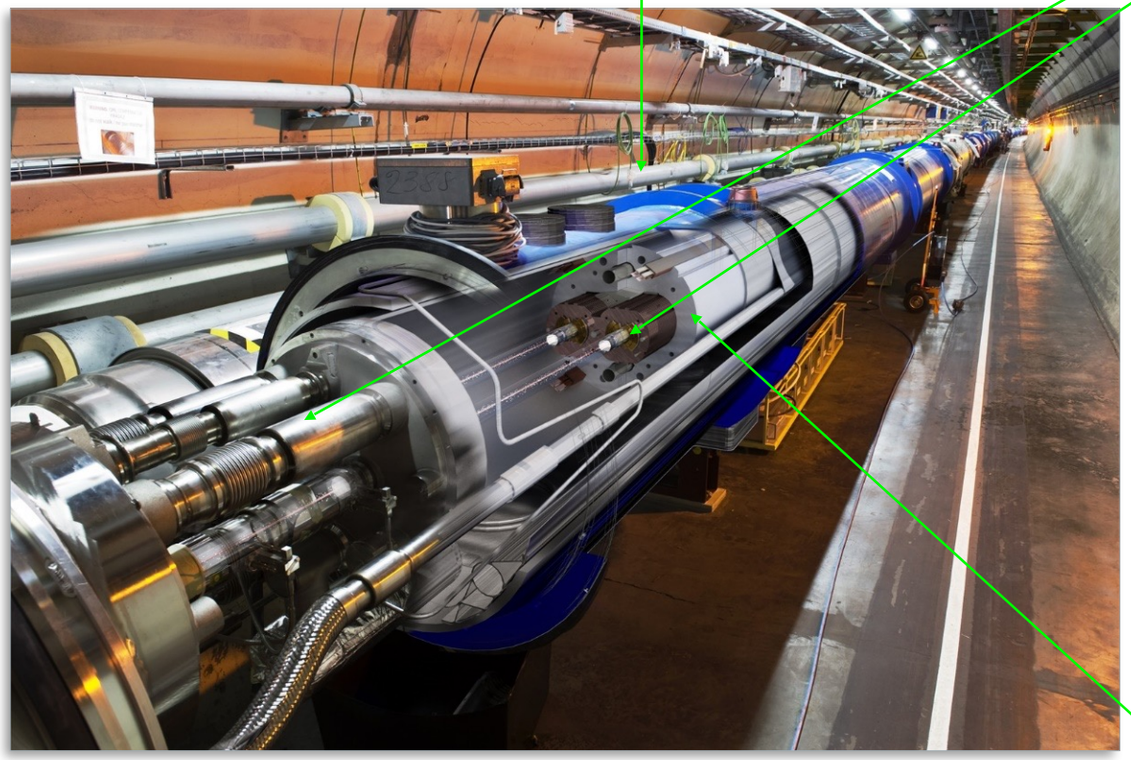

Layout Service Database

The Layout Service documents the accelerator topographical organisation by defining components (Functional Positions) & relationships between them:

- Mechanical assemblies (Physical structures)
- Powering connections & circuits
- Control connections & circuits
- MAD optic sequences

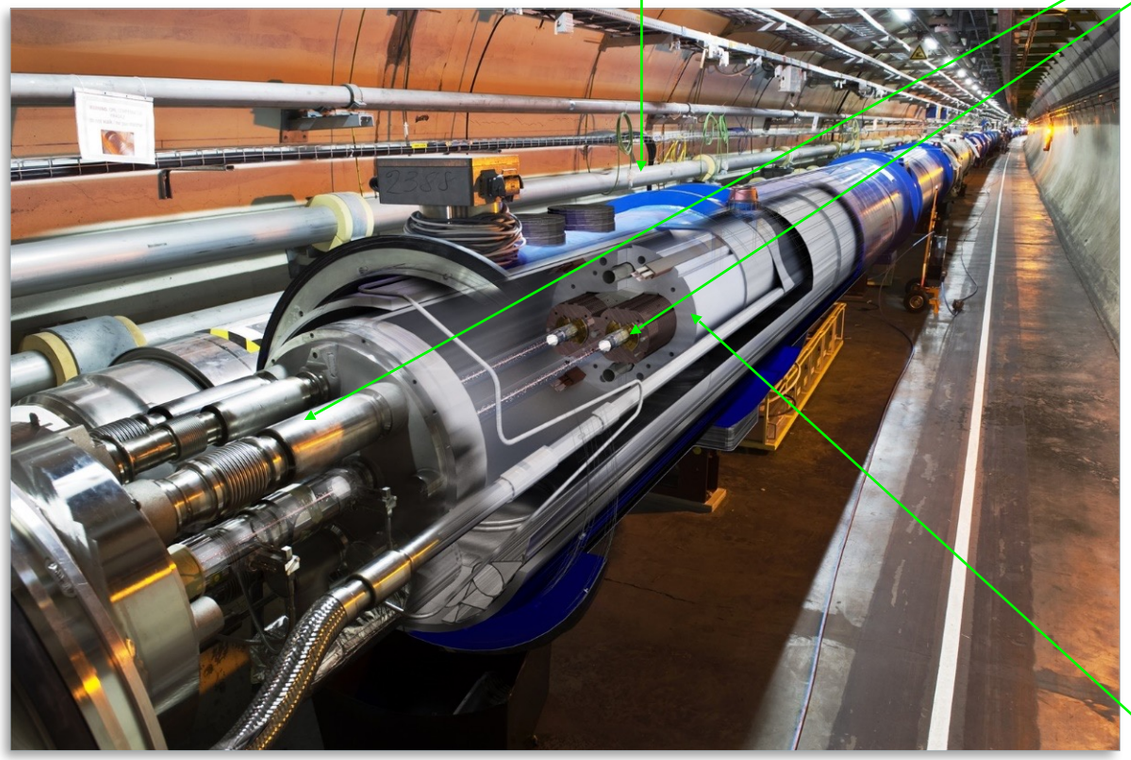

A large highly relational database oriented system based on:

- Oracle 11g database (RAC cluster)
- Set of high-level client Java APIs & Database level client APIs (PL/SQL interfaces)
- Legacy GUIs based on .NET, Oracle Forms & APEX (proprietary Oracle technologies) & new generation GUIs based on BE-CO ACW stack (Spring, HTML5, AngularJS)

Source of information for many (~40) other CERN systems, and serving around 200 users, from all over CERN.

Layout Service Database - Scope Evolution

Massive growth in Scope, Community, & Support since 2003…

Layout Service Database - Scope Evolution

Massive growth in Scope, Community, & Support since 2003…

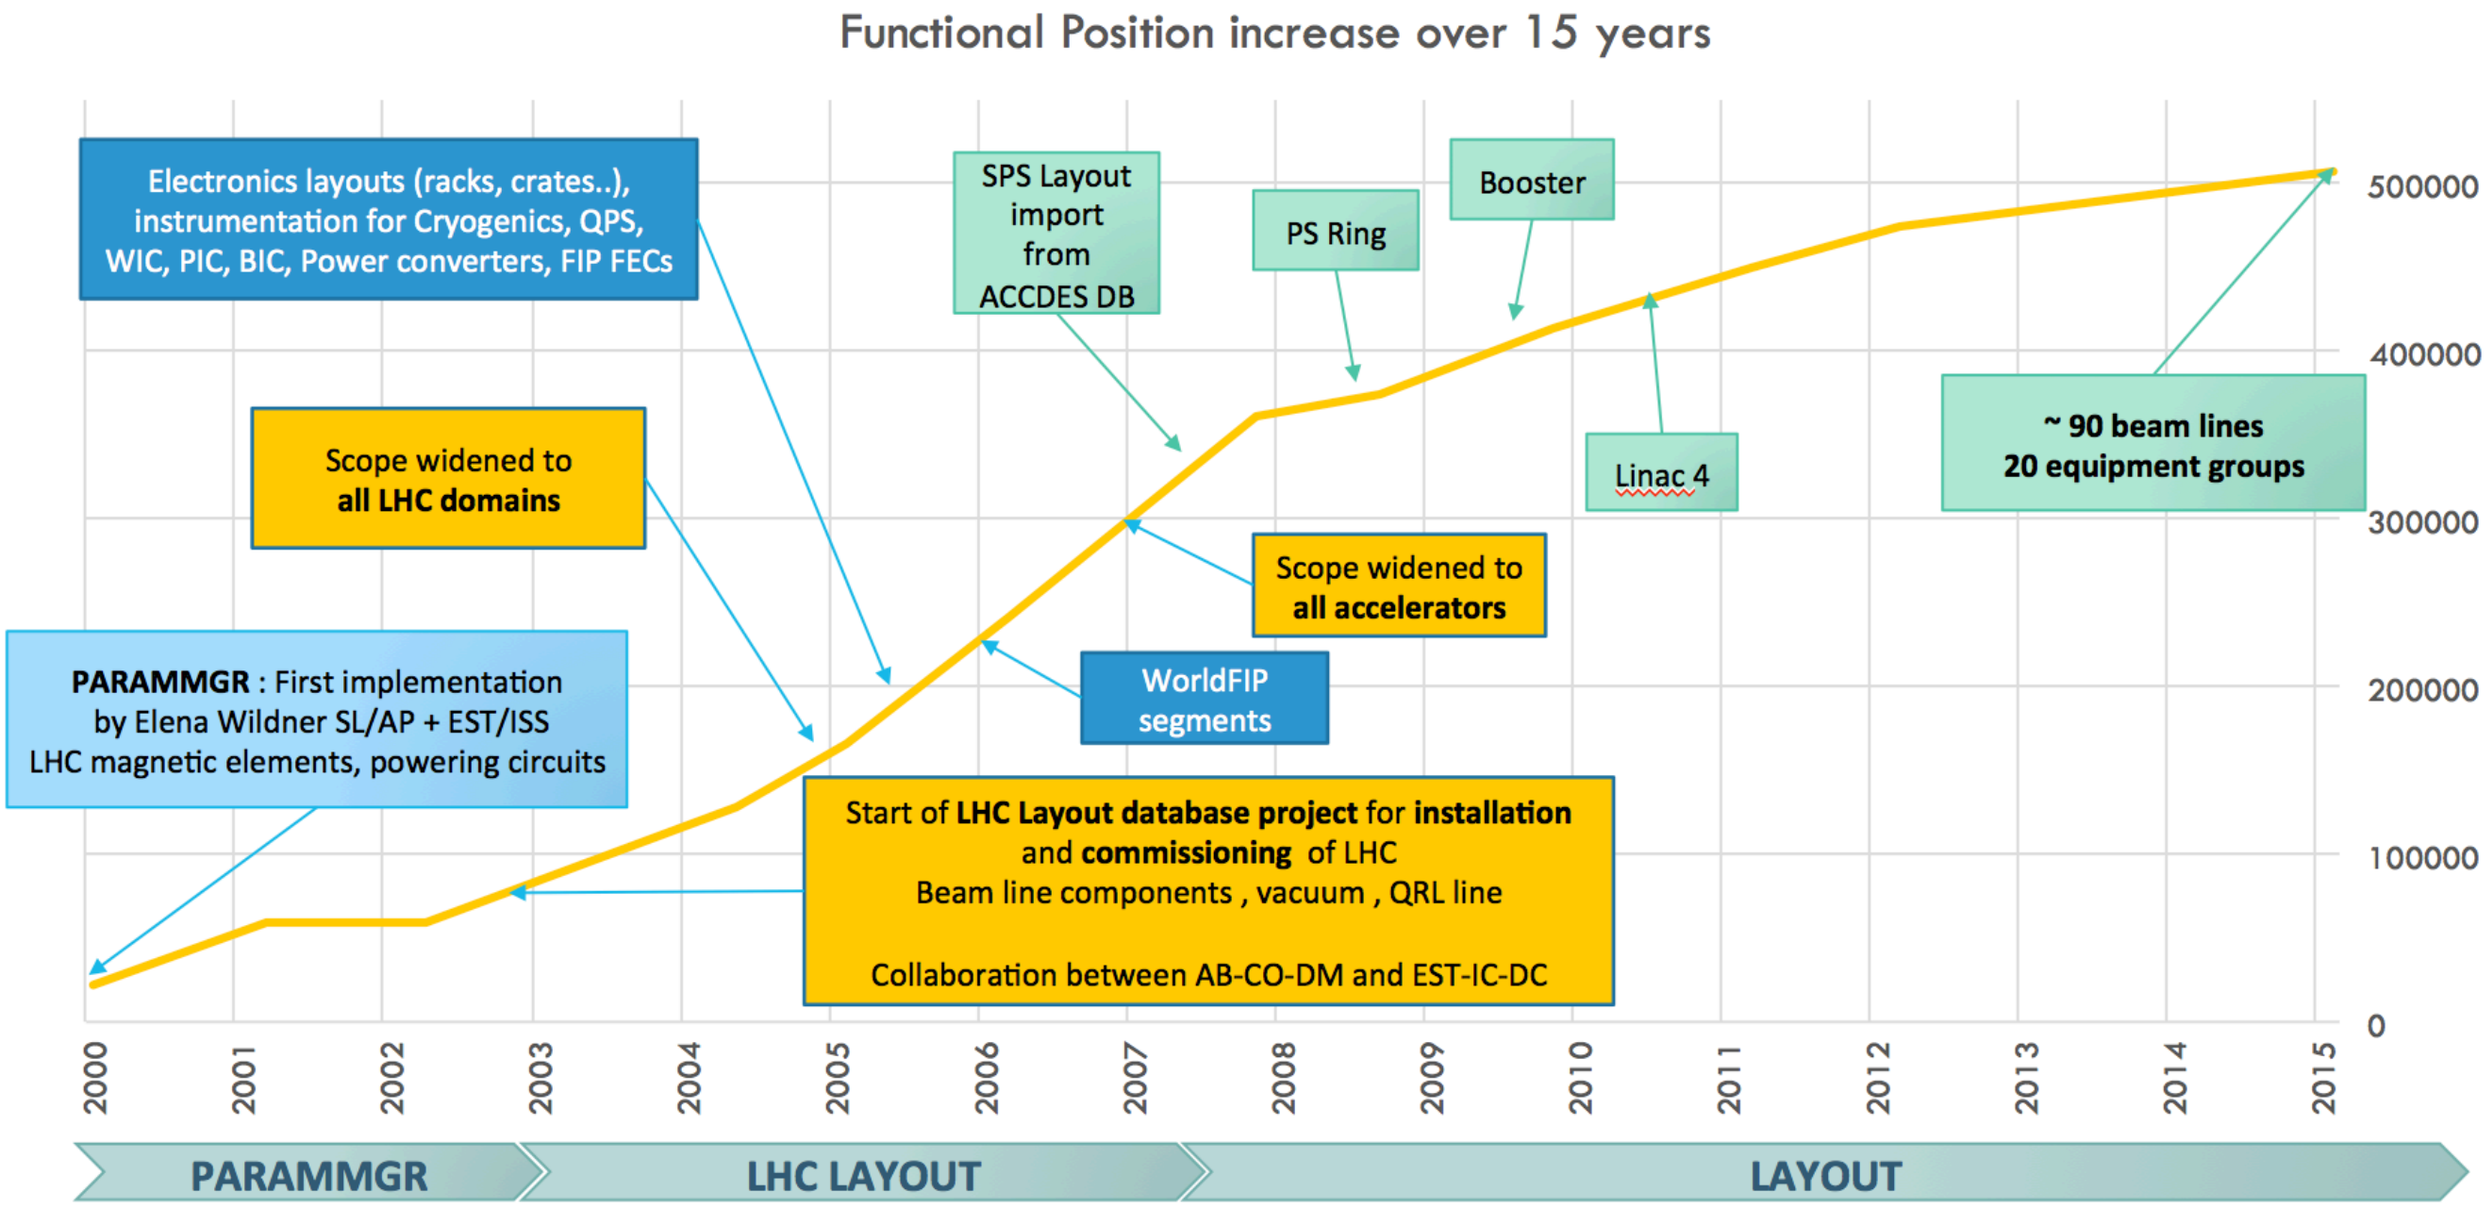

Layout Service Database - Scope Evolution

Massive growth in Scope, Community, & Support since 2003…

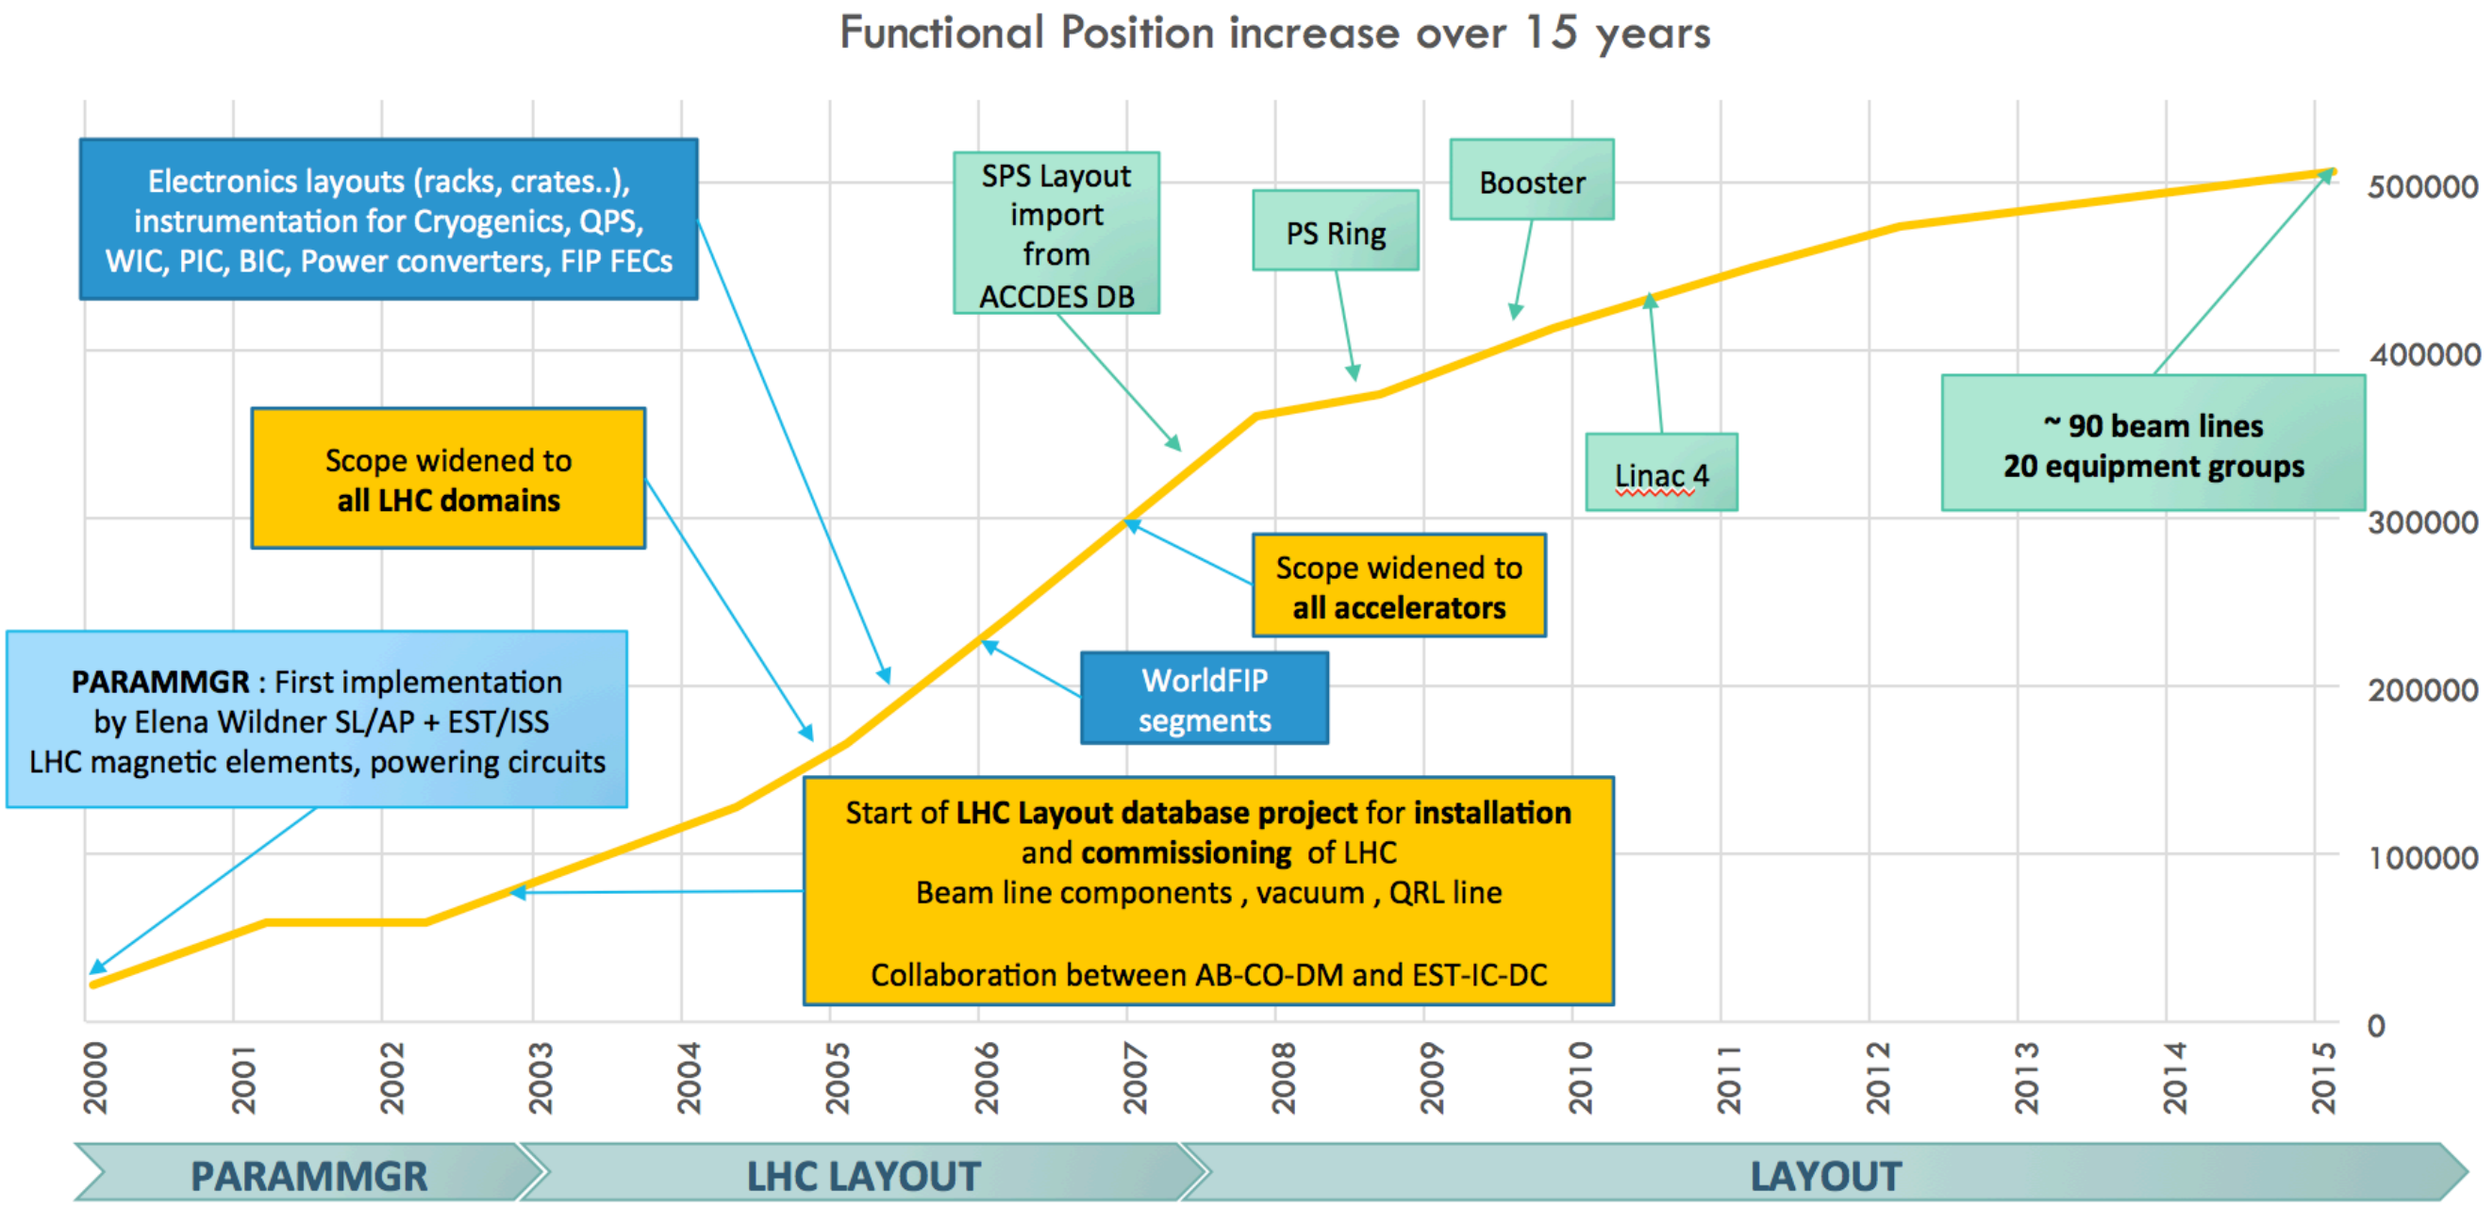

However, more of a domain knowledge and support challenge, rather than a pure technical challenge

Layout Service Database - Activities & Outlook

Layout Service Database - Activities & Outlook

Each accelerator development requires new data entry and updates on a massive scale.

Layout Service Database - Activities & Outlook

Each accelerator development requires new data entry and updates on a massive scale.

New system under development to satisfy new requirements & empower users to manage their own data via suitable tools.

Layout Service Database - Activities & Outlook

Each accelerator development requires new data entry and updates on a massive scale.

New system under development to satisfy new requirements & empower users to manage their own data via suitable tools.

Focus on replacing .NET, Oracle Forms & APEX with more suitable high-quality end-user oriented applications.

Layout Service Database - Activities & Outlook

Each accelerator development requires new data entry and updates on a massive scale.

New system under development to satisfy new requirements & empower users to manage their own data via suitable tools.

Focus on replacing .NET, Oracle Forms & APEX with more suitable high-quality end-user oriented applications.

Overall, the relational model is valid and Oracle satisfies most needs

Layout Service Database - Activities & Outlook

Each accelerator development requires new data entry and updates on a massive scale.

New system under development to satisfy new requirements & empower users to manage their own data via suitable tools.

Focus on replacing .NET, Oracle Forms & APEX with more suitable high-quality end-user oriented applications.

Overall, the relational model is valid and Oracle satisfies most needs

One new key requirement is to manage past, present and parallel future Layout versions. Need for data relations with temporal validity - answer questions like "*What was the Layout in Run 1?*" or "*What will be the Layout in Run 3?*"

Layout Service Database - Activities & Outlook

Each accelerator development requires new data entry and updates on a massive scale.

New system under development to satisfy new requirements & empower users to manage their own data via suitable tools.

Focus on replacing .NET, Oracle Forms & APEX with more suitable high-quality end-user oriented applications.

Overall, the relational model is valid and Oracle satisfies most needs

One new key requirement is to manage past, present and parallel future Layout versions. Need for data relations with temporal validity - answer questions like "*What was the Layout in Run 1?*" or "*What will be the Layout in Run 3?*"

→Currently implemented in-house, excludes usages of Foreign Keys for such cases. Something out of the box would be very welcome to reduce complexity and maintenance. Oracle 12c starts to provide something, however it seems not a complete solution.

Controls Configuration Service Database

Controls Configuration Service Database

The Controls Configuration Service helps bind all Control system layers together by providing complete and coherent configurations.

Controls Configuration Service Database

The Controls Configuration Service helps bind all Control system layers together by providing complete and coherent configurations.

A large highly relational database oriented system based on:

- Oracle 11g database (2-node RAC cluster)
- Set of high-level client Java APIs & Database level client APIs (PL/SQL interfaces)
- Legacy GUIs based on APEX (proprietary Oracle technology) & new generation GUIs based on BE-CO ACW stack (Spring, HTML5, AngularJS)

Controls Configuration Service Database

The Controls Configuration Service helps bind all Control system layers together by providing complete and coherent configurations.

A large highly relational database oriented system based on:

- Oracle 11g database (2-node RAC cluster)
- Set of high-level client Java APIs & Database level client APIs (PL/SQL interfaces)
- Legacy GUIs based on APEX (proprietary Oracle technology) & new generation GUIs based on BE-CO ACW stack (Spring, HTML5, AngularJS)

Mission-critical serving around 400 users, mainly from the A&T Sector

Controls Configuration Service Database

The Controls Configuration Service helps bind all Control system layers together by providing complete and coherent configurations.

A large highly relational database oriented system based on:

- Oracle 11g database (2-node RAC cluster)
- Set of high-level client Java APIs & Database level client APIs (PL/SQL interfaces)
- Legacy GUIs based on APEX (proprietary Oracle technology) & new generation GUIs based on BE-CO ACW stack (Spring, HTML5, AngularJS)

Mission-critical serving around 400 users, mainly from the A&T Sector

One of the first Oracle Database systems at CERN

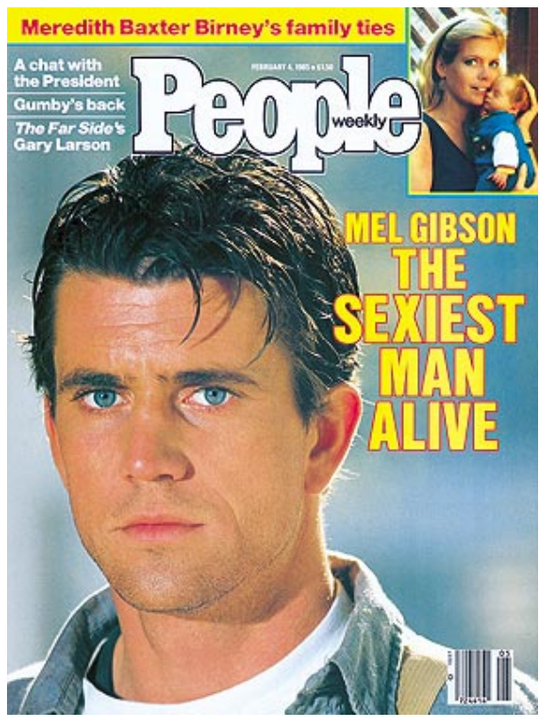

Controls Configuration Service Database - Domains & Stakeholders

high-level applications

- **LSA** — classes, devices migrations
- **CALS** — classes, devices
- **LASER*** — alarms definition, data-management
- **DIAMON*** — metrics, agents, data-management
- **GM** — meta-model, treatment-codes
- **CCM** — OP Configs, applications, menus configuration

middleware

- **CMW** — servers, proxies
- **FESA** — class meta-model, deployment, instantiation, promotion
- **FGC** — meta-model, equip configuration, history
- **Timing** — GMT/HL events, networks, Timing XMLs generation
- **OASIS** — connections, events, stats, viewer menus
- **RBAC** — roles, access rules, tokens history

low-level components

- **Hardware** — computers, hw types, crates, modules, cables
- **FEC Startup** — DSC programs, reboot, FEC families
- **White Rabbit** — switches, calibration data

generic aspects

- **Device property-model** — HW, Virtual, SL, FESA, GM
- **Device migrations** — classes, devices

external systems

- **WinCC/Unicos*** — Unicos applications conf
- **TE-MPE** — PIC, WIC, BIS, SMP

DB access · GUI Tools · Processing · Java API · read-only · read-write

Controls Configuration Service Database - Activities & Outlook

Controls Configuration Service Database - Activities & Outlook

Huge levels of Technical Debt - steady but sure service-wide consolidation effort since late 2014 (first time in 29 years!) - foreseen to continue into LS2

Controls Configuration Service Database - Activities & Outlook

Huge levels of Technical Debt - steady but sure service-wide consolidation effort since late 2014 (first time in 29 years!) - foreseen to continue into LS2

Focus on eradicating APEX and providing integrated user experience via more suitable high-quality end-user oriented applications and APIs

Controls Configuration Service Database - Activities & Outlook

Huge levels of Technical Debt - steady but sure service-wide consolidation effort since late 2014 (first time in 29 years!) - foreseen to continue into LS2

Focus on eradicating APEX and providing integrated user experience via more suitable high-quality end-user oriented applications and APIs

Overall, the relational model is valid and Oracle satisfies most needs

Controls Configuration Service Database - Activities & Outlook

Huge levels of Technical Debt - steady but sure service-wide consolidation effort since late 2014 (first time in 29 years!) - foreseen to continue into LS2

Focus on eradicating APEX and providing integrated user experience via more suitable high-quality end-user oriented applications and APIs

Overall, the relational model is valid and Oracle satisfies most needs

Strong need for Enhanced Auditing / Tracing - answer questions like "*who changed what, when, and how (from which tools)?*"

→Should we continue to rely on our in-house "Commons4Oracle" trigger-based solutions - or can we get more out of the box from Oracle / 12c?

Controls Configuration Service Database - Activities & Outlook

Huge levels of Technical Debt - steady but sure service-wide consolidation effort since late 2014 (first time in 29 years!) - foreseen to continue into LS2

Focus on eradicating APEX and providing integrated user experience via more suitable high-quality end-user oriented applications and APIs

Overall, the relational model is valid and Oracle satisfies most needs

Strong need for Enhanced Auditing / Tracing - answer questions like "*who changed what, when, and how (from which tools)?*"

→Should we continue to rely on our in-house "Commons4Oracle" trigger-based solutions - or can we get more out of the box from Oracle / 12c?

Increasingly difficult to schedule downtime for interventions - needed for operations, and needed for access and interventions.

→Can we move towards zero downtime? Fast switching stand-by service?

LSA Accelerator Settings Management Database

LSA Accelerator Settings Management Database

The LSA Accelerator Settings Management System is used to modify & drive settings to 10,000's of accelerator devices to control particle beams.

LSA Accelerator Settings Management Database

The LSA Accelerator Settings Management System is used to modify & drive settings to 10,000's of accelerator devices to control particle beams.

A large highly relational database oriented system based on:

- Java Servers with complex business logic
- Oracle 11g database (2-node RAC cluster)
- Java client Java APIs & Java GUIs

LSA Accelerator Settings Management Database

The LSA Accelerator Settings Management System is used to modify & drive settings to 10,000's of accelerator devices to control particle beams.

A large highly relational database oriented system based on:

- Java Servers with complex business logic
- Oracle 11g database (2-node RAC cluster)
- Java client Java APIs & Java GUIs

Mission-critical used on-line by Operators and Experts to control the entire Accelerator complex.

LSA Accelerator Settings Management Database

The LSA Accelerator Settings Management System is used to modify & drive settings to 10,000's of accelerator devices to control particle beams.

A large highly relational database oriented system based on:

- Java Servers with complex business logic
- Oracle 11g database (2-node RAC cluster)
- Java client Java APIs & Java GUIs

Mission-critical used on-line by Operators and Experts to control the entire Accelerator complex.

<Acc Operators>

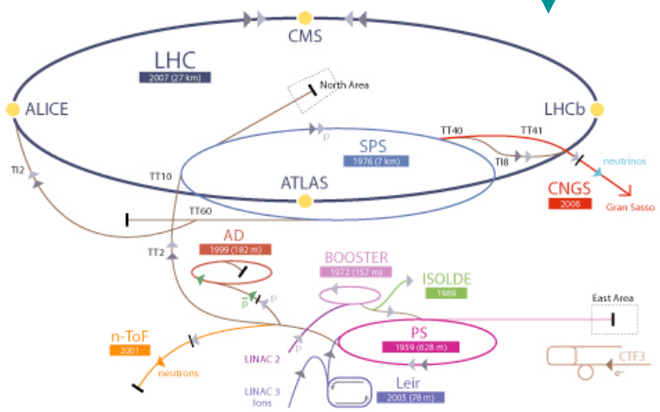

<Accelerators>

LSA Accelerator Settings Management Database

The LSA Accelerator Settings Management System is used to modify & drive settings to 10,000's of accelerator devices to control particle beams.

A large highly relational database oriented system based on:

- Java Servers with complex business logic
- Oracle 11g database (2-node RAC cluster)
- Java client Java APIs & Java GUIs

Mission-critical used on-line by Operators and Experts to control the entire Accelerator complex.

100's Billions of active & historical settings.

Massive data skew in many dimensions.

<Acc Operators>

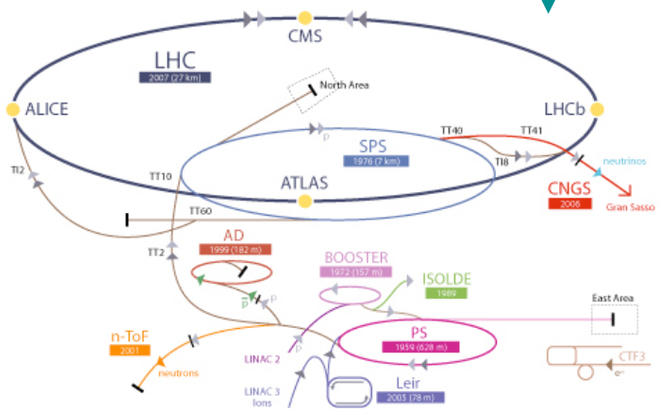

<Accelerators>

LSA Accelerator Settings Management Database

The LSA Accelerator Settings Management System is used to modify & drive settings to 10,000's of accelerator devices to control particle beams.

A large highly relational database oriented system based on:

- Java Servers with complex business logic
- Oracle 11g database (2-node RAC cluster)
- Java client Java APIs & Java GUIs

Mission-critical used on-line by Operators and Experts to control the entire Accelerator complex.

100's Billions of active & historical settings.

Massive data skew in many dimensions.

Need to maintain a full history of all changes & be able to quickly revert to prior settings.

<Acc Operators>

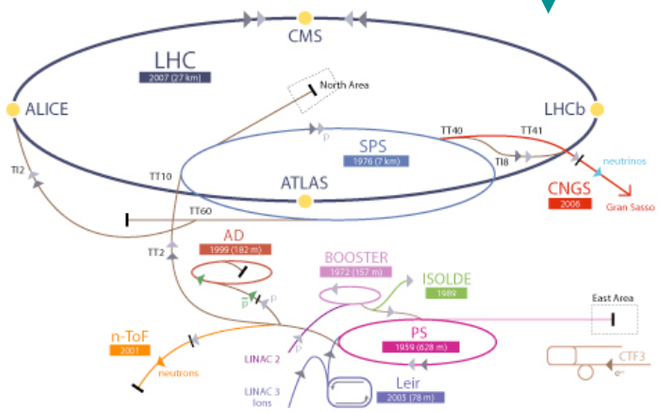

<Accelerators>

Accelerator Settings Management - Activities & Outlook

Accelerator Settings Management - Activities & Outlook

Increasing number of active & historical settings together with data skew have led to poor performance in recent years. Temporary measures in place, but operators have high expectations for response times.

Accelerator Settings Management - Activities & Outlook

Increasing number of active & historical settings together with data skew have led to poor performance in recent years. Temporary measures in place, but operators have high expectations for response times.

Luminosity levelling for HL-LHC will generate more settings changes via on-line feedback systems.

Accelerator Settings Management - Activities & Outlook

Increasing number of active & historical settings together with data skew have led to poor performance in recent years. Temporary measures in place, but operators have high expectations for response times.

Luminosity levelling for HL-LHC will generate more settings changes via on-line feedback systems.

Apart from an expected increase in settings (currently ~150GB → ~1TB max for the future?), no other new requirements are on the horizon.

Accelerator Settings Management - Activities & Outlook

Increasing number of active & historical settings together with data skew have led to poor performance in recent years. Temporary measures in place, but operators have high expectations for response times.

Luminosity levelling for HL-LHC will generate more settings changes via on-line feedback systems.

Apart from an expected increase in settings (currently ~150GB → ~1TB max for the future?), no other new requirements are on the horizon.

Overall, the relational model is valid and Oracle satisfies *most* needs

Accelerator Settings Management - Activities & Outlook

Increasing number of active & historical settings together with data skew have led to poor performance in recent years. Temporary measures in place, but operators have high expectations for response times.

Luminosity levelling for HL-LHC will generate more settings changes via on-line feedback systems.

Apart from an expected increase in settings (currently ~150GB → ~1TB max for the future?), no other new requirements are on the horizon.

Overall, the relational model is valid and Oracle satisfies *most* needs

Concerning response time performance:

→Should we rely on hardware improvements for response times? In-memory features? New 12c partitioning schemes? Alternative technology for settings archiving?

Accelerator Settings Management - Activities & Outlook

Increasing number of active & historical settings together with data skew have led to poor performance in recent years. Temporary measures in place, but operators have high expectations for response times.

Luminosity levelling for HL-LHC will generate more settings changes via on-line feedback systems.

Apart from an expected increase in settings (currently ~150GB → ~1TB max for the future?), no other new requirements are on the horizon.

Overall, the relational model is valid and Oracle satisfies *most* needs

Concerning response time performance:
→Should we rely on hardware improvements for response times? In-memory features? New 12c partitioning schemes? Alternative technology for settings archiving?

Extremely difficult to schedule downtime for interventions outside of a YETS.
→Can we move towards zero downtime? Fast switching stand-by service?

Logging Service Databases

Logging Service Databases

The Logging Service stores accelerator beam & equipment data on-line, to be kept beyond the lifetime of the LHC (>20y). Used to analyse and improve behaviour of accelerators & their sub-systems over time.

Logging Service Databases

The Logging Service stores accelerator beam & equipment data on-line, to be kept beyond the lifetime of the LHC (>20y). Used to analyse and improve behaviour of accelerators & their sub-systems over time.

A very large time series database oriented system based on:

- Oracle 11g database (2-node RAC cluster)
- Java and PL/SQL data logging processes and APIs
- Java data extraction API & GUI (TIMBER)

Logging Service Databases

The Logging Service stores accelerator beam & equipment data on-line, to be kept beyond the lifetime of the LHC (>20y). Used to analyse and improve behaviour of accelerators & their sub-systems over time.

A very large time series database oriented system based on:

- Oracle 11g database (2-node RAC cluster)
- Java and PL/SQL data logging processes and APIs
- Java data extraction API & GUI (TIMBER)

Close to 1PB stored for > 1.5 million signals

Mission-critical serving > 1,000 users from all over CERN

Logging Service Databases

The Logging Service stores accelerator beam & equipment data on-line, to be kept beyond the lifetime of the LHC (>20y). Used to analyse and improve behaviour of accelerators & their sub-systems over time.

A very large time series database oriented system based on:

- Oracle 11g database (2-node RAC cluster)
- Java and PL/SQL data logging processes and APIs
- Java data extraction API & GUI (TIMBER)

Close to 1PB stored for > 1.5 million signals

Mission-critical serving > 1,000 users from all over CERN

Designed in 2001 based on LEP experience, in production since 2003. Original Estimated data rate: 1TB/year…

Logging Service Databases - 6 years ago...

Logging Service Databases - 6 years ago...

From: Database Futures Workshop, 2011-06-06

LHC Logging Service - Logged records

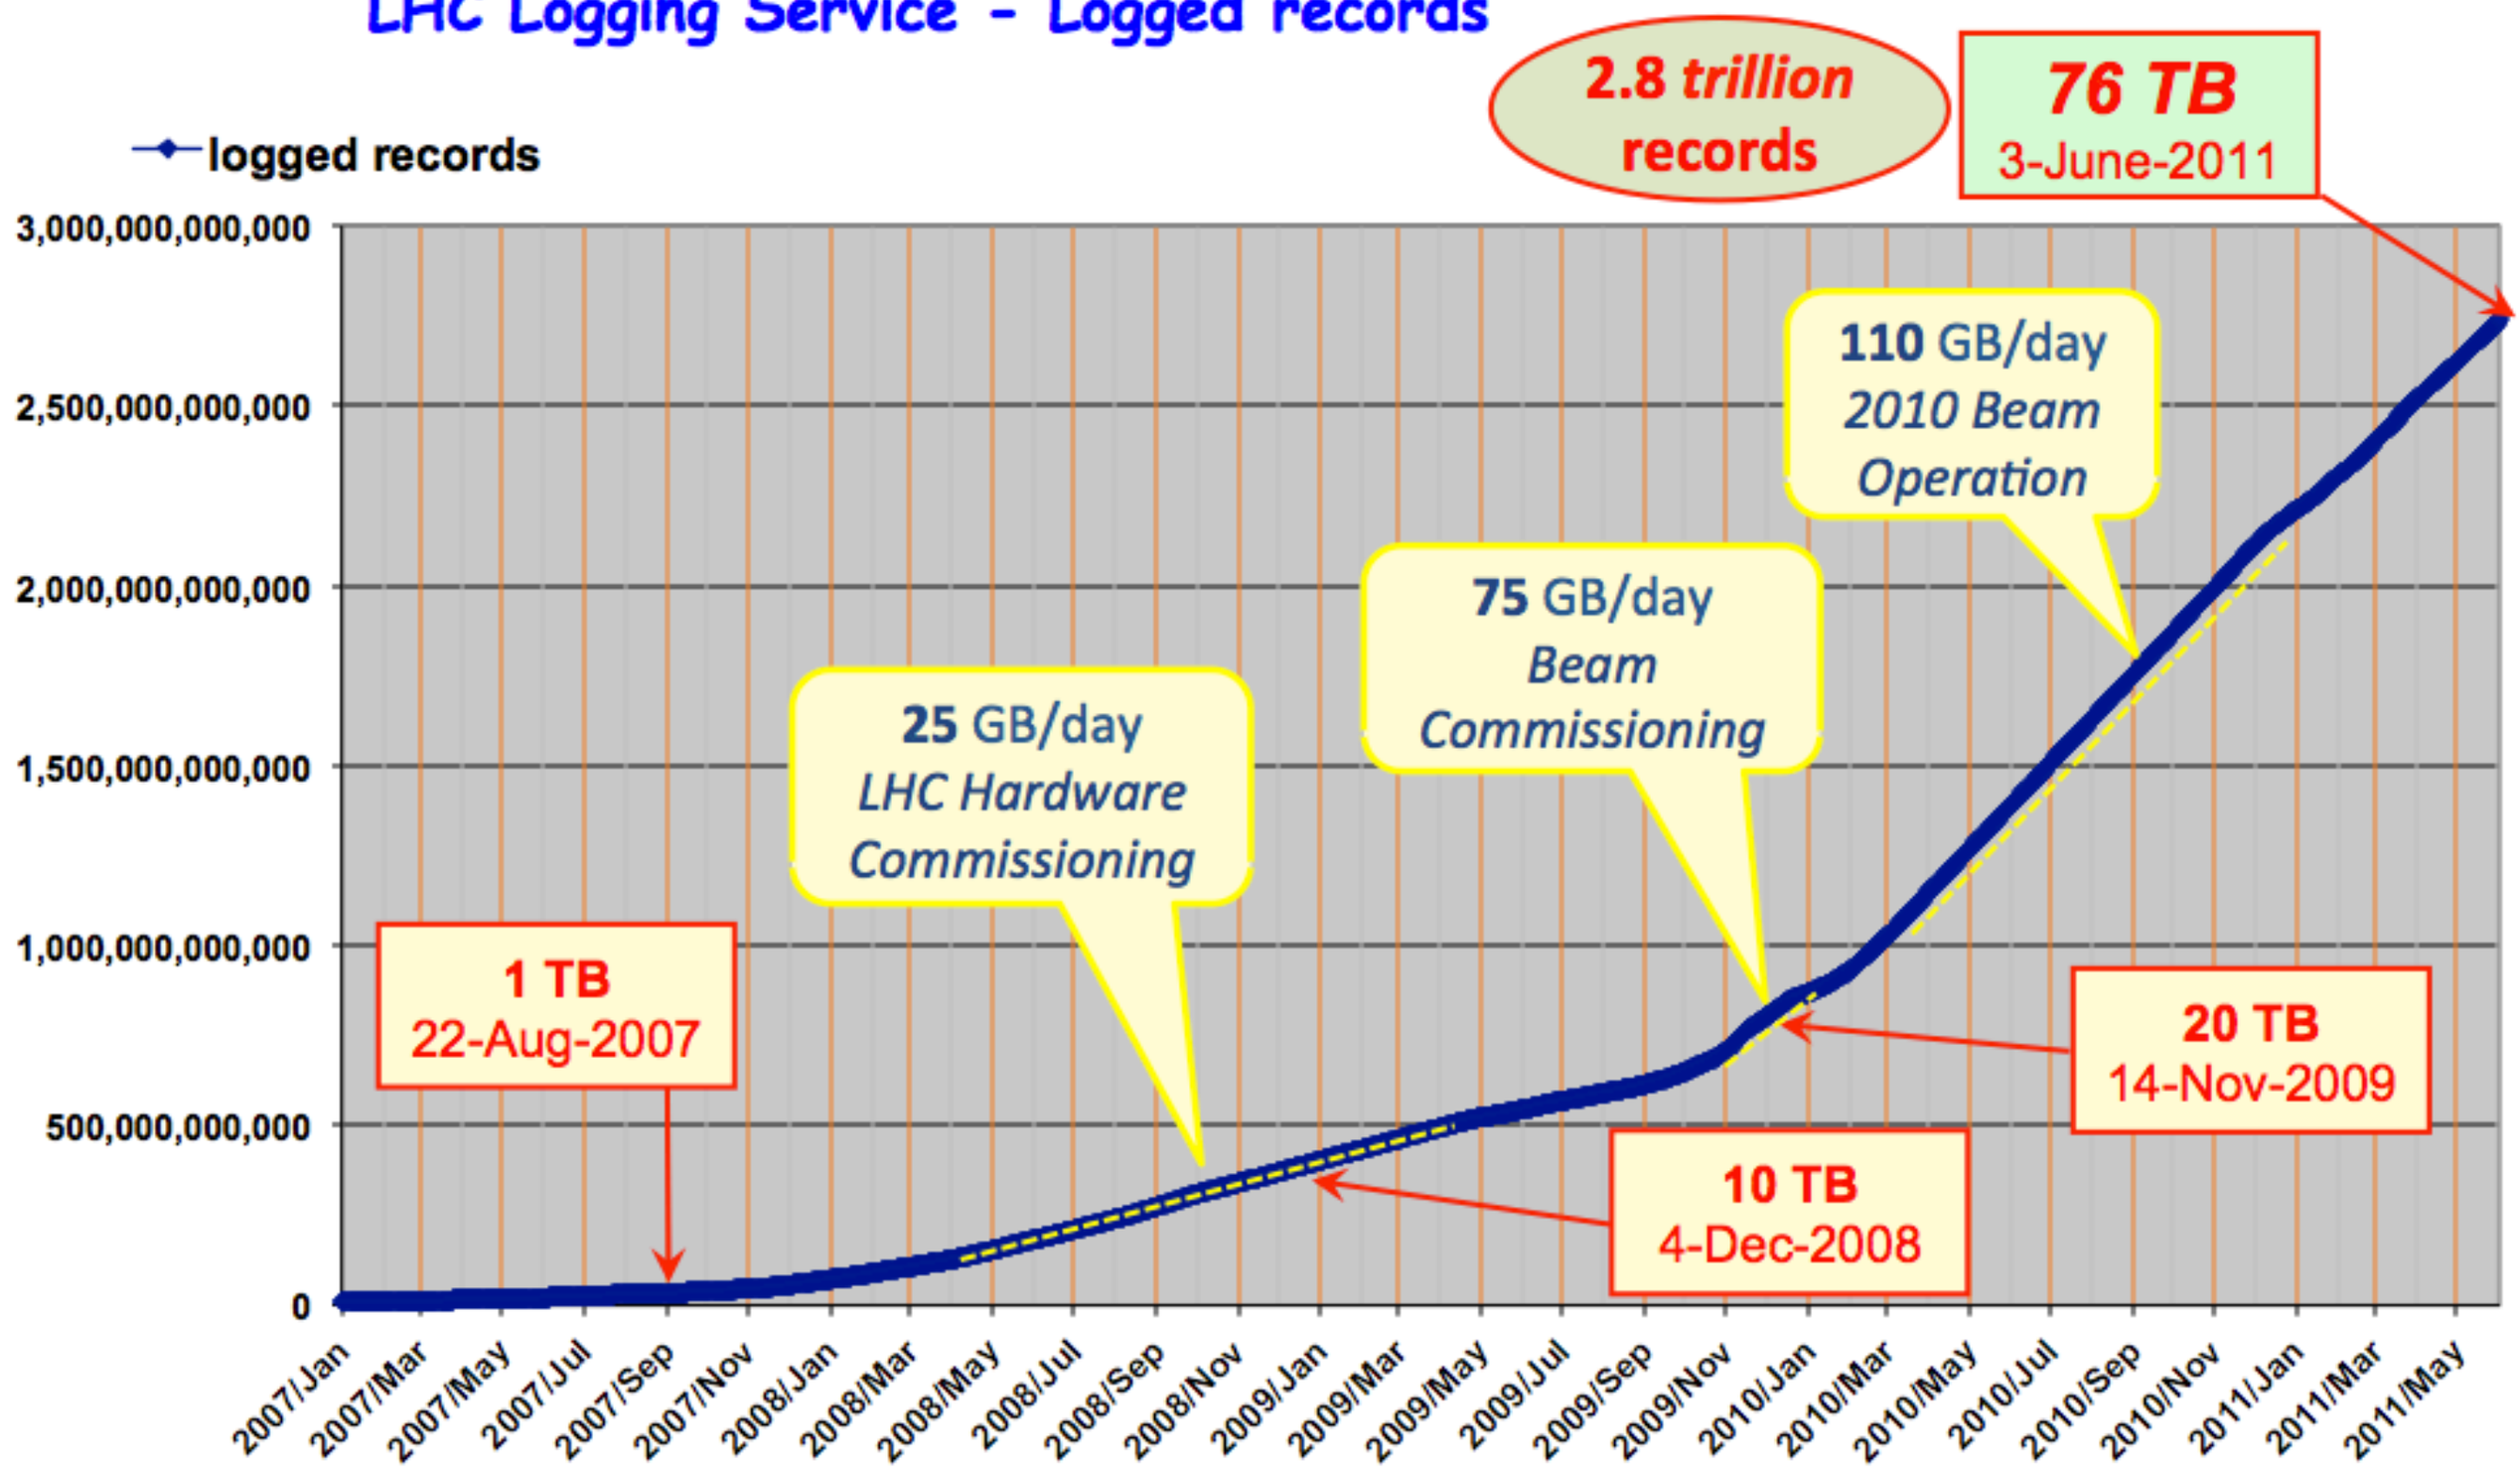

Logging Service Databases - Activities & Outlook

Logging Service Databases - Activities & Outlook

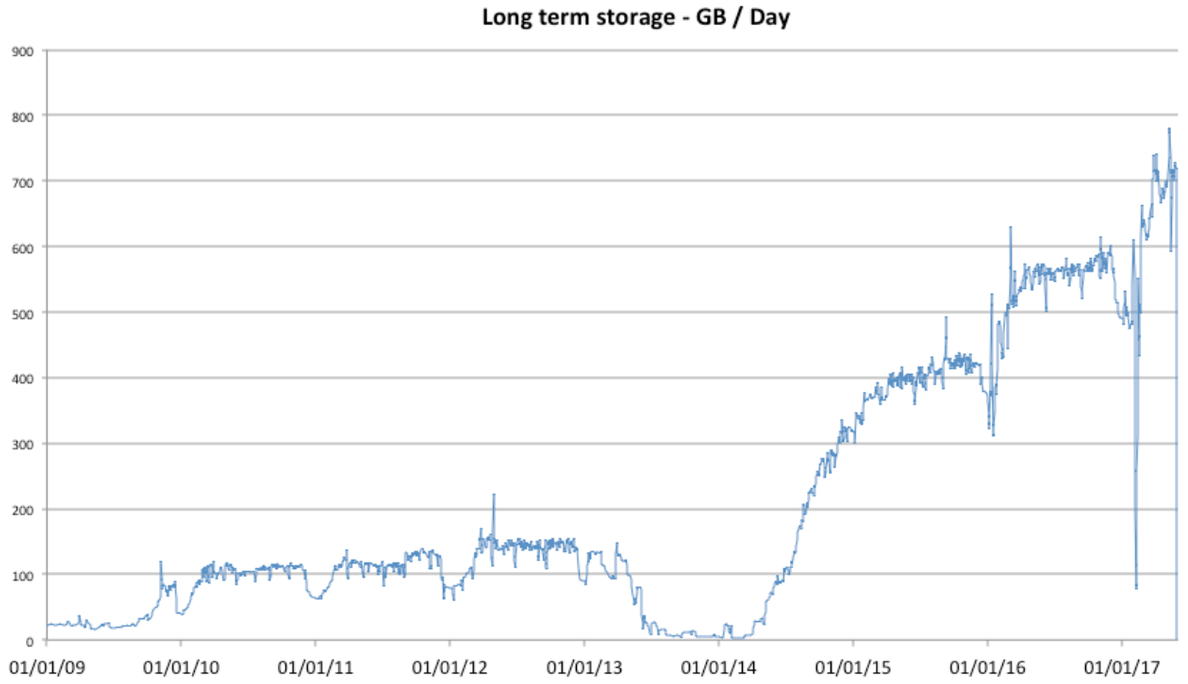

Logging Service Databases - Activities & Outlook

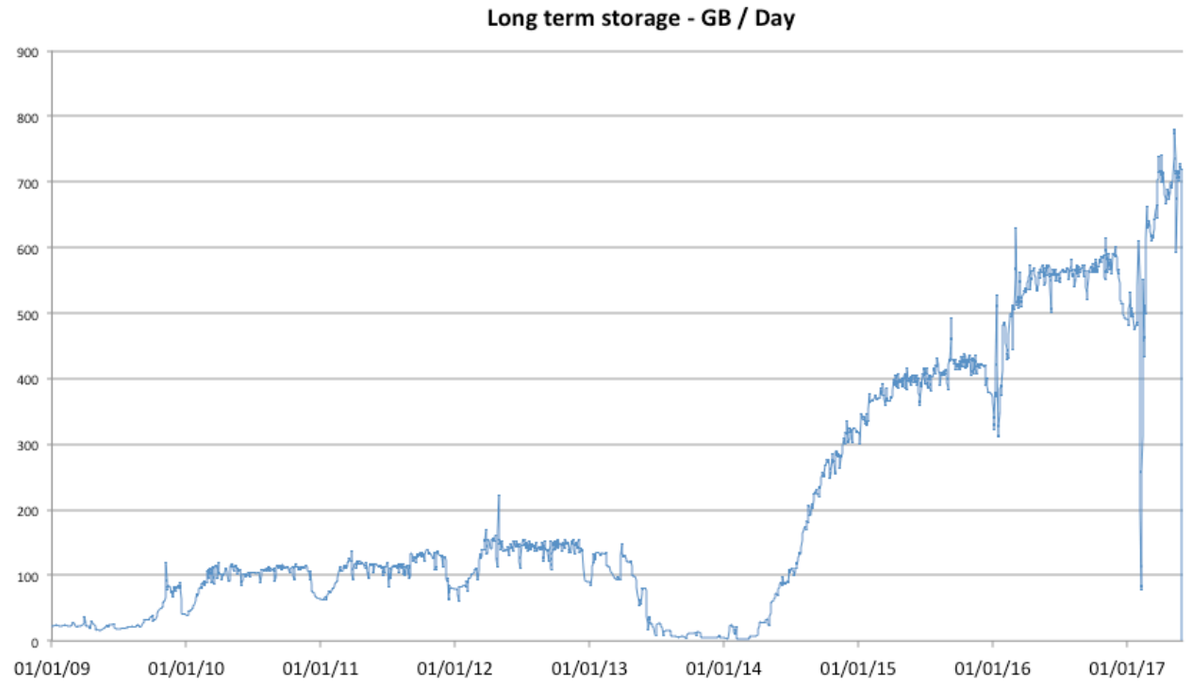

In 2016, after 13 years in service and with a throughput of >1.5 TB / day - acknowledged that the service cannot satisfy new demands to analyse huge data sets in an efficient manner.

Logging Service Databases - Activities & Outlook

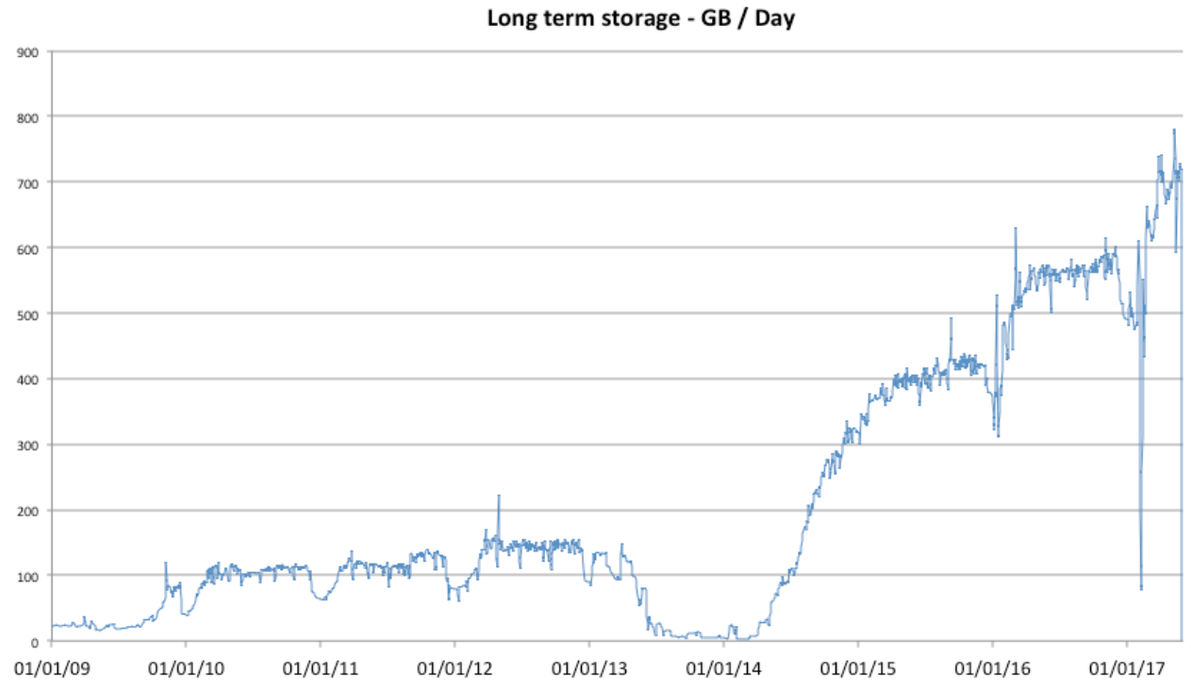

In 2016, after 13 years in service and with a throughput of >1.5 TB / day - acknowledged that the service cannot satisfy new demands to analyse huge data sets in an efficient manner.

A new system "NXCALS" is under development based on "Big Data" technologies and in collaboration with IT (see the talk from Nikolay Tsvetkov at 16:00).

→Current "CALS" system is foreseen to be turned off during 2019 (LS2).

Logging Service Databases - Activities & Outlook

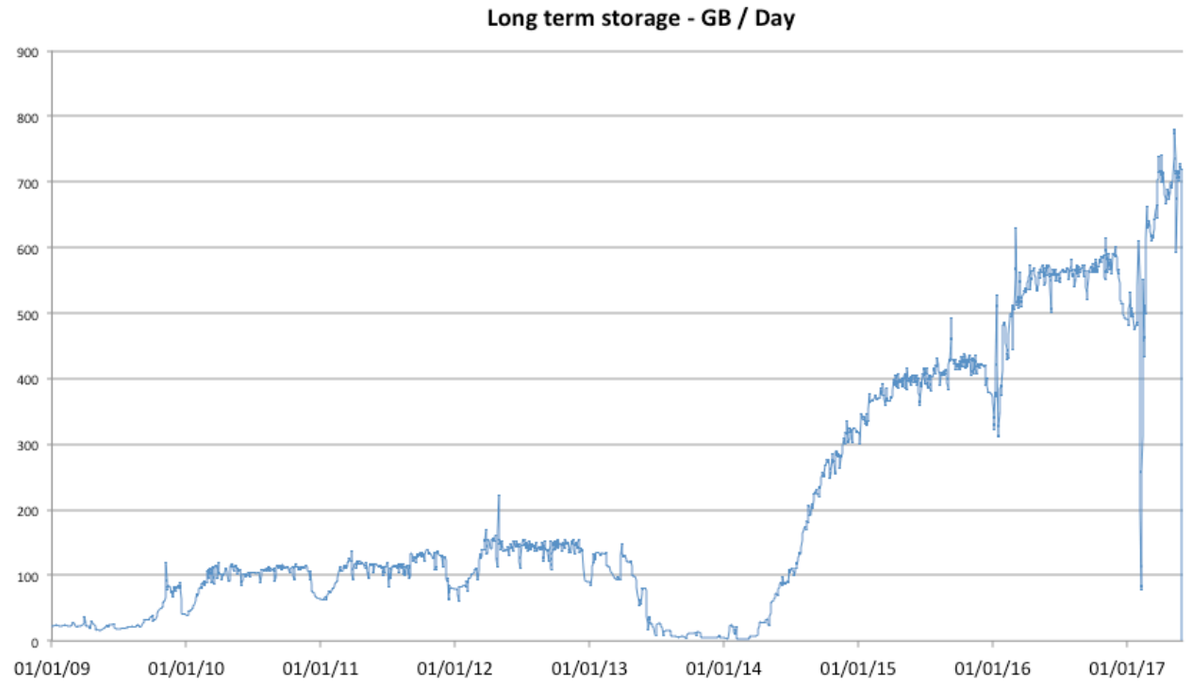

In 2016, after 13 years in service and with a throughput of >1.5 TB / day - acknowledged that the service cannot satisfy new demands to analyse huge data sets in an efficient manner.

A new system "NXCALS" is under development based on "Big Data" technologies and in collaboration with IT (see the talk from Nikolay Tsvetkov at 16:00).

→Current "CALS" system is foreseen to be turned off during 2019 (LS2).

Opportunities to collaborate with other CERN teams and streamline solutions:

- Serhiy Boychenko: "Next generation for Post Mortem event storage and analysis" at 17:00
- Piotr Golonka: "Next generation Archiver" at 17:20

General: Collaborations & Open Source

General: Collaborations & Open Source

Increasing wish to be able to:

- Collaborate with other institutes
- Open source our systems (or sub-systems)

General: Collaborations & Open Source

Increasing wish to be able to:

- Collaborate with other institutes
- Open source our systems (or sub-systems)

For many cases, Oracle is a bottleneck to this

General: Collaborations & Open Source

Increasing wish to be able to:

- Collaborate with other institutes
- Open source our systems (or sub-systems)

For many cases, Oracle is a bottleneck to this

Trade off:

- Efficiently develop solutions for CERN in the short-term
- Establish longer term potentially rewarding collaborations for the community

General: Collaborations & Open Source

Increasing wish to be able to:

- Collaborate with other institutes
- Open source our systems (or sub-systems)

For many cases, Oracle is a bottleneck to this

Trade off:

- Efficiently develop solutions for CERN in the short-term
- Establish longer term potentially rewarding collaborations for the community

Foresee to start putting our hands on some other relational database solutions in the coming years

Summary

Summary

Accelerator Databases are well established and new requirements are uncommon and slow to emerge

Summary

Accelerator Databases are well established and new requirements are uncommon and slow to emerge

NXCALS will replace existing CALS Logging system in 2019 (16 years in production)

- Data volumes and analysis workloads will surely increase
- Expect horizontally scalable architecture to pay off
- IT should foresee a gradual increase in cluster size (space, power, cooling)

Summary

Accelerator Databases are well established and new requirements are uncommon and slow to emerge

NXCALS will replace existing CALS Logging system in 2019 (16 years in production)

- Data volumes and analysis workloads will surely increase
- Expect horizontally scalable architecture to pay off
- IT should foresee a gradual increase in cluster size (space, power, cooling)

No other *significant* changes are foreseen:

Summary

Accelerator Databases are well established and new requirements are uncommon and slow to emerge

NXCALS will replace existing CALS Logging system in 2019 (16 years in production)

- Data volumes and analysis workloads will surely increase
- Expect horizontally scalable architecture to pay off
- IT should foresee a gradual increase in cluster size (space, power, cooling)

No other *significant* changes are foreseen:

Remain with Oracle: most cost-effective platform in terms of functionality & performance, also considering existing investments and expertise

LS2: we will move to 12cR2 - start profiting from new feature set in Run 3

Summary

Accelerator Databases are well established and new requirements are uncommon and slow to emerge

NXCALS will replace existing CALS Logging system in 2019 (16 years in production)

- Data volumes and analysis workloads will surely increase
- Expect horizontally scalable architecture to pay off
- IT should foresee a gradual increase in cluster size (space, power, cooling)

No other *significant* changes are foreseen:

Remain with Oracle: most cost-effective platform in terms of functionality & performance, also considering existing investments and expertise

LS2: we will move to 12cR2 - start profiting from new feature set in Run 3

Start getting experience with other relational database technologies for smaller projects: to facilitate collaborations & open sourcing

Summary

Accelerator Databases are well established and new requirements are uncommon and slow to emerge

NXCALS will replace existing CALS Logging system in 2019 (16 years in production)

- Data volumes and analysis workloads will surely increase
- Expect horizontally scalable architecture to pay off
- IT should foresee a gradual increase in cluster size (space, power, cooling)

No other *significant* changes are foreseen:

Remain with Oracle: most cost-effective platform in terms of functionality & performance, also considering existing investments and expertise

LS2: we will move to 12cR2 - start profiting from new feature set in Run 3

Start getting experience with other relational database technologies for smaller projects: to facilitate collaborations & open sourcing

General: anticipate the need for better performance (response times), easier auditing, and aim for even higher availability

Summary

Questions?