

ATLAS SoC - Review

System-on-Chip Workshop - 12th June 2019 – CERN

<https://indico.cern.ch/event/799275/>

R. Kopeliansky, Indiana University, On behalf of the ATLAS collaboration

Many thanks to:

M. Aleksa, C. Amelung, V. Andrei, G. Avolio, F. Carrió Argos, M. Corradi, T. Costa De Paiva, H. Evans, P. Farthouat
B. Gorini, R. Hart, L. Hervas, O. Kepka, F. Lanni, F. Martins, D. Miller, P. Moschovakos, L. Pontecorvo, S. Schlenker, R. Spiwoks,
A. Straessner, W. Vandelli, S. Veneziano, M. Wittgen, C. Yildiz
for the info, help and cooperation in preparing this summary ☺

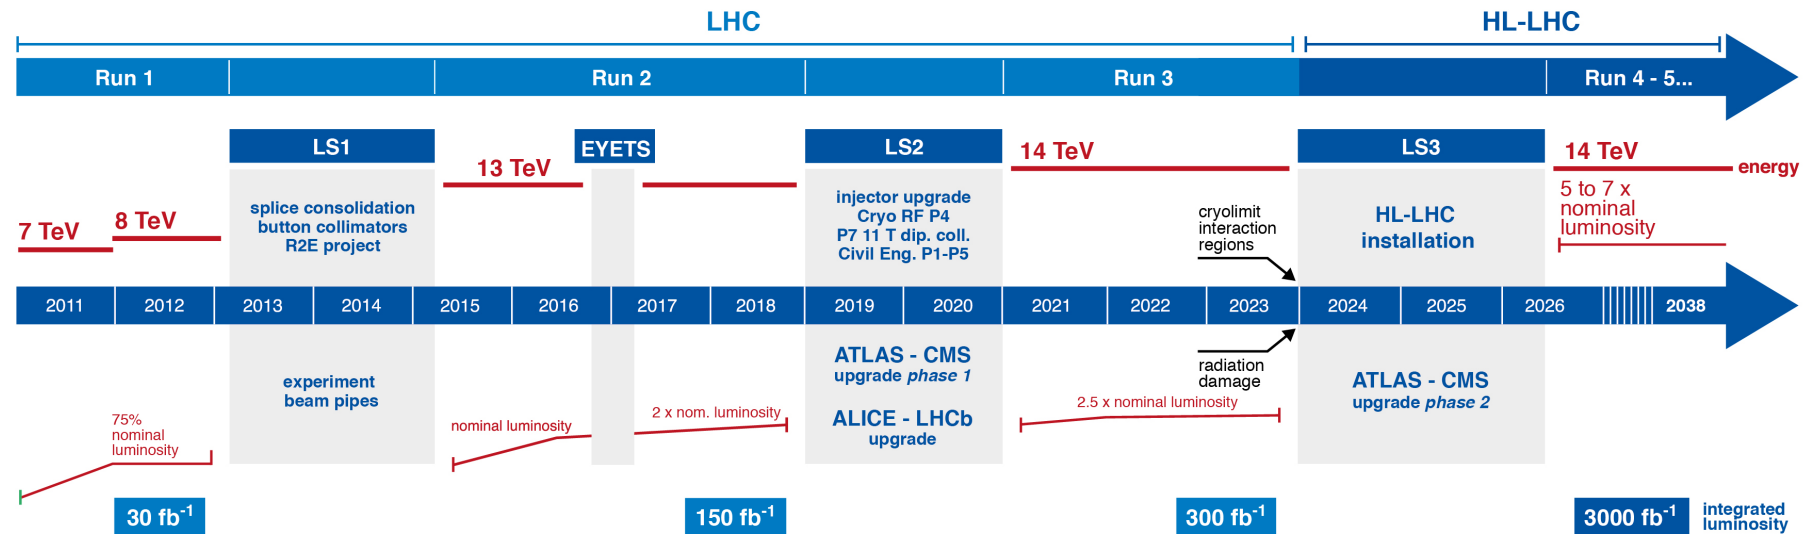
Pixel, LAr, Tile, Muon, TDAQ, TC, UC, ESE



Outline

- ATLAS SOC usage evolution over the upgrade phases
- Associated challenges & coordination of related solutions

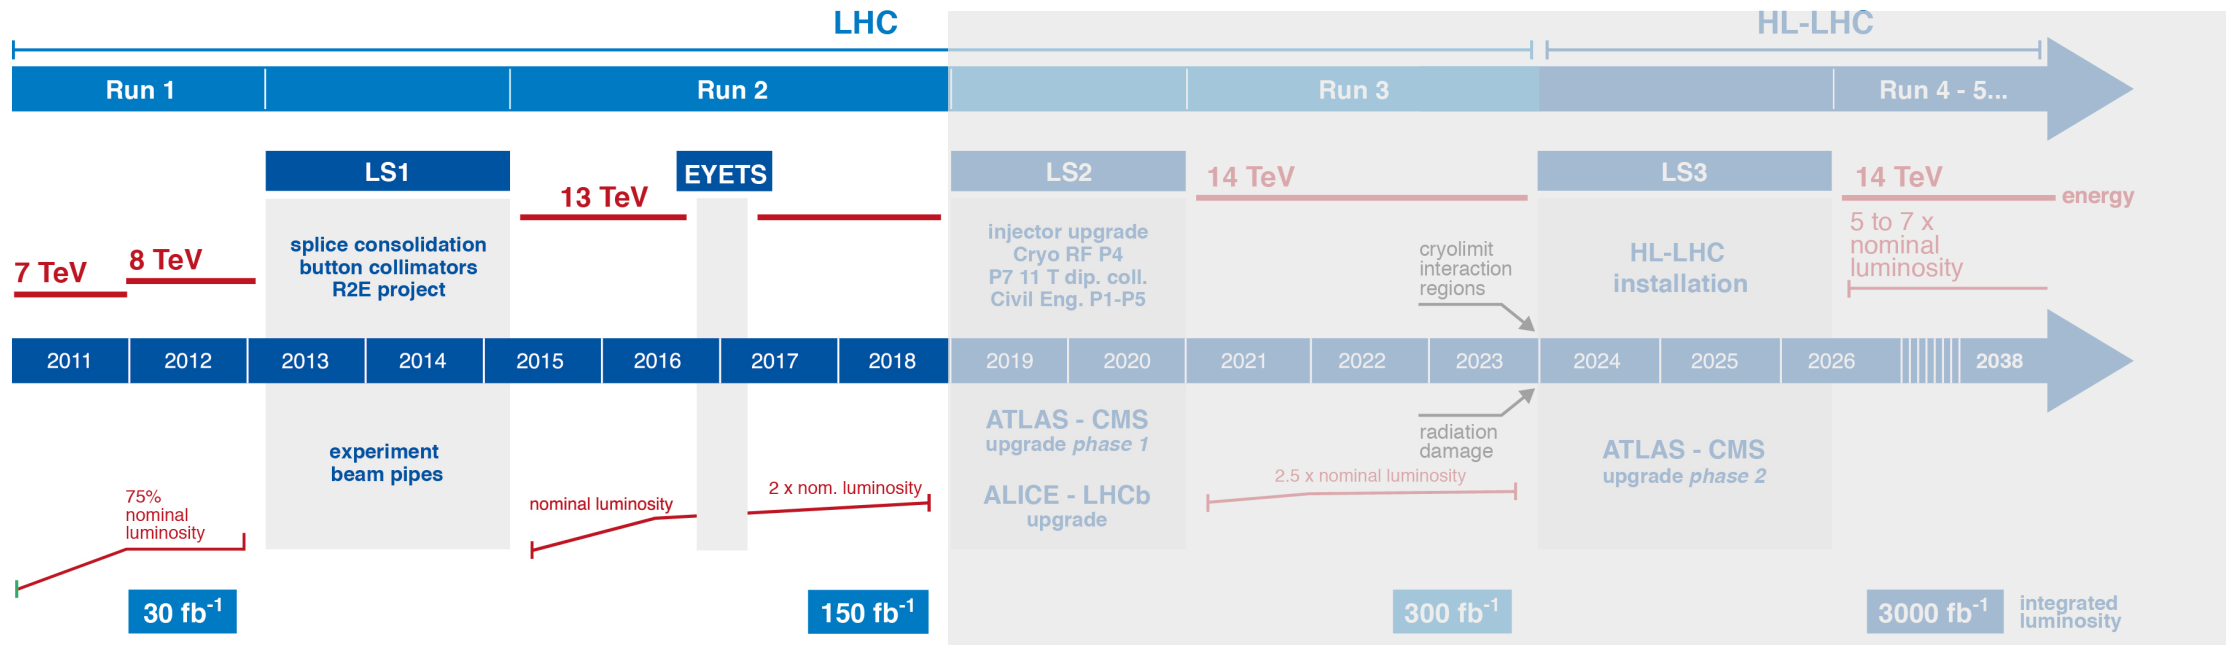
LHC / HL-LHC Plan



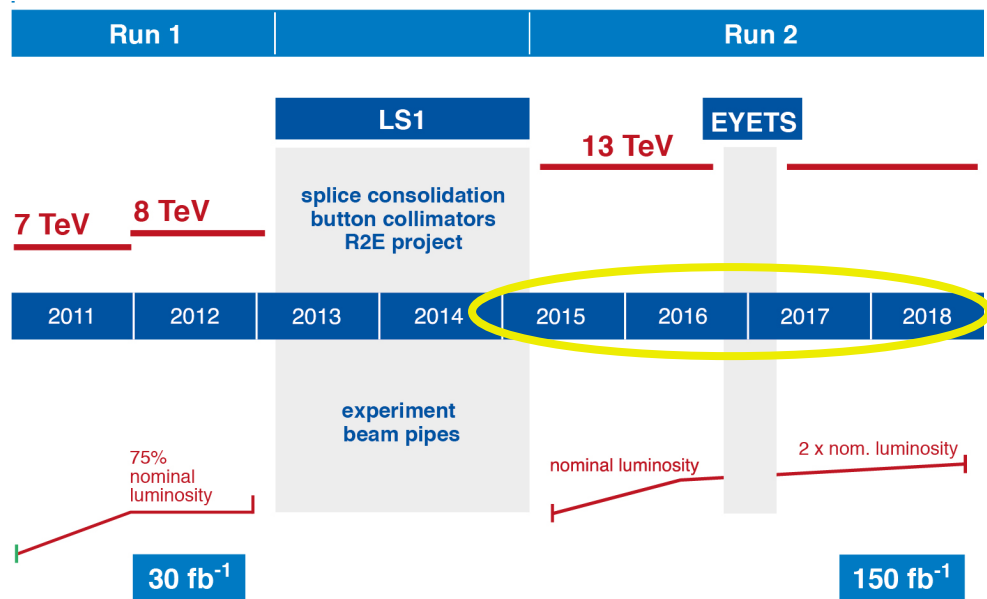
ATLAS SOC usage over the upgrade phases...

Run-2:

- During LS1 several ATLAS subsystems started deploying SoCs
- SoC main usage:
 - Readout, control & configuration



Phase-0 and up to end of Run-2



Inner Detector:

- **Pixel / IBL**
 - 2015-2018 → New Readout Drivers (RODs)
 - 9U VME modules, containing several FPGAs
 - One master FPGA Xilinx Virtex5 + PPC440
 - SoC usage:
 - Configuration & readback monitoring from the other FPGAs
 - Will run until the end of Run-3
 - *See talk from Oldřich Kepka on Thursday*

Muon Detector:

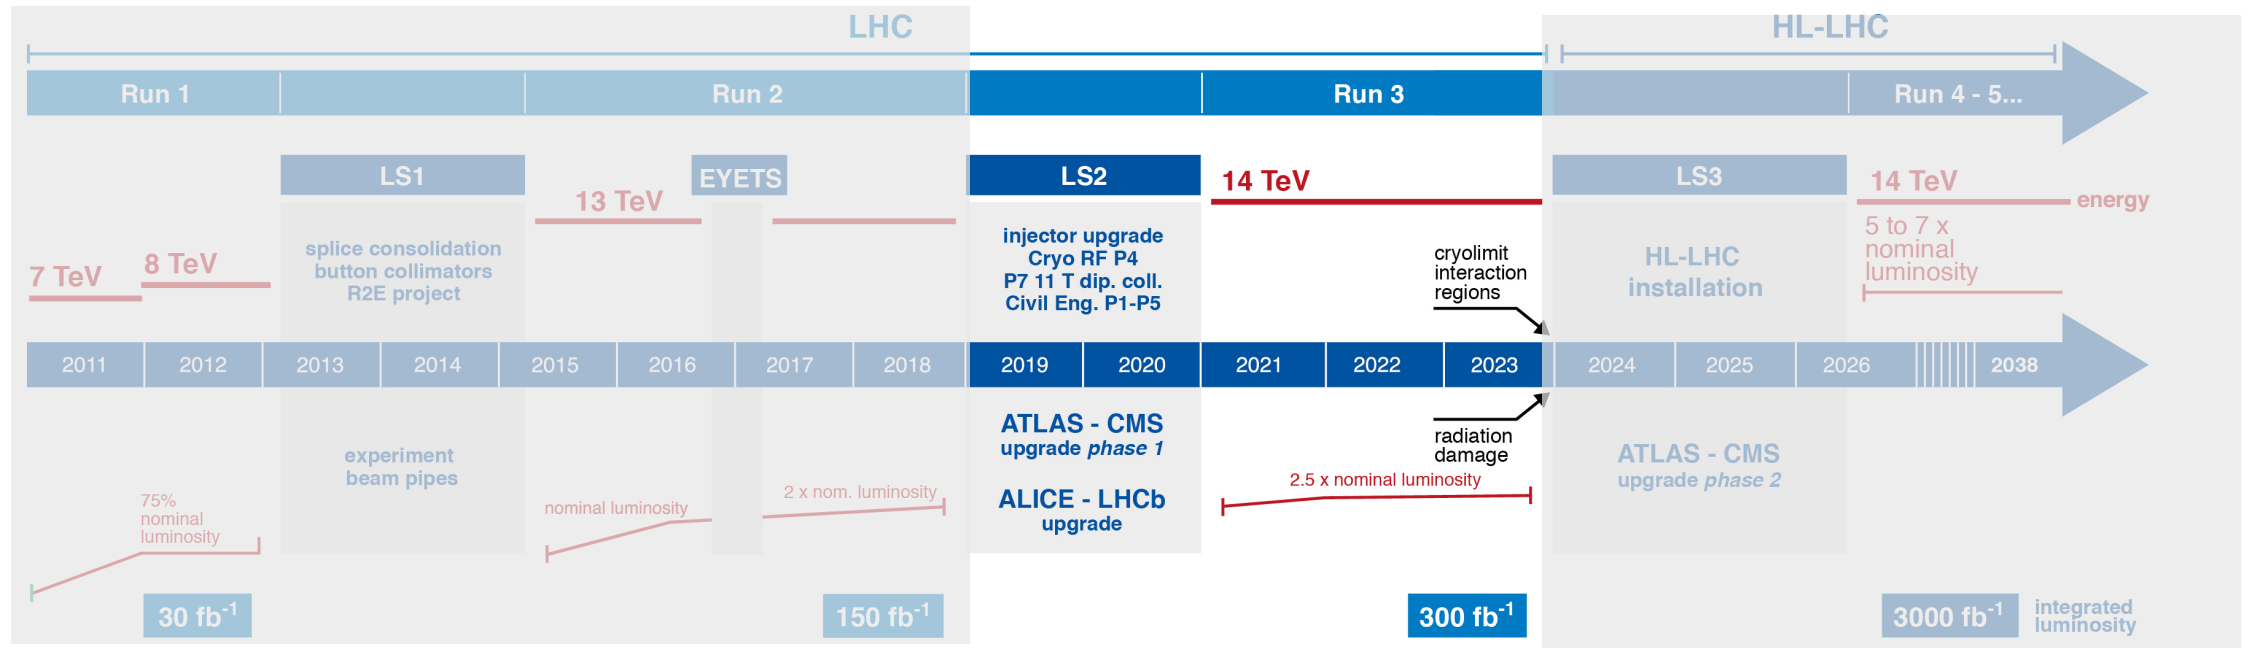
- **CSC** – Cathode Strip Chambers
 - 2015 → New RODs
 - VME to ATCA
 - Modules containing several Xilinx Zynq-7000
 - SoC usage:
 - Readout functionality
 - *See talk from Matthias Wittgen on Thursday*
- **MDT** – Muon Drift Tubes
 - Dec 2016 → Upgraded frame grabbers of the barrel alignment system
 - ‘Florida’ carrier boards, hosting ‘Miami’ SoMs with Xilinx Zynq Z7015
 - SoC usage:
 - Frame grabber functionality

Forward Detector:

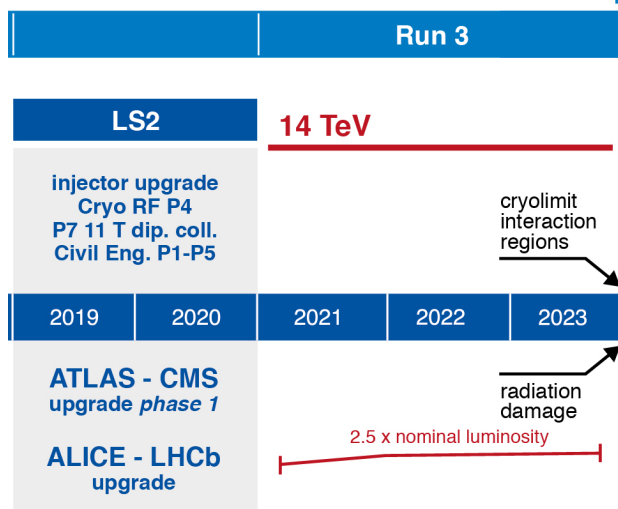
- **AFP** – ATLAS Forward Proton
 - 2016 → 1st installation
 - Readout on Xilinx Artix FPGAs, and Zynq
 - SoC usage:
 - Control & configuration of the FPGAs
 - *See talk from Matthias Wittgen on Thursday*

Phase-I:

- Few trigger systems started to include SoCs in their design
- SoC usage spectrum is being expanded to communicate with the DCS



Phase-1 upgrade



Muon Detector:

- **NSW** – New Small Wheel
 - The NSW Trigger Processors (TP) are ATCA based
 - Blades with Xilinx US FPGAs for running algorithms + Zynq-7000
 - Used for HW control and monitoring

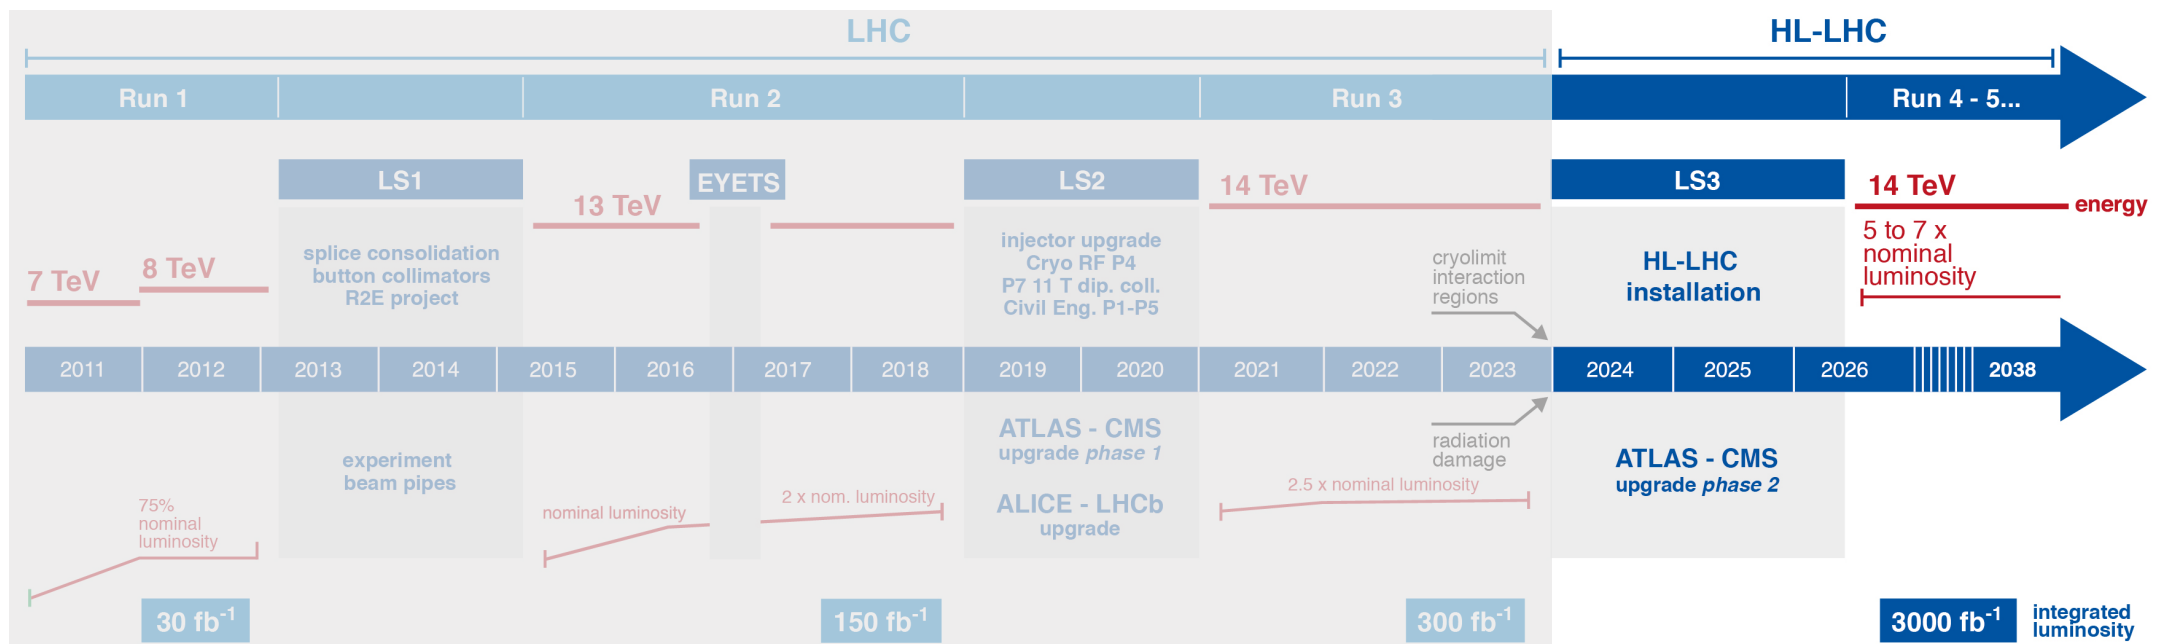
Trigger-DAQ (TDAQ):

- L1Calo
 - **TREX** – Tile Rear Extension
 - VME digital modules
 - Zynq MPSoC on each
 - SoC usage:
 - Slow control, monitoring, and communication with DCS
 - *See talk from Víctor Andrei on Thursday*
 - **gFEX** – global Feature Extractor
 - ATCA blade with US+ Virtex FPGAs and Zynq US+MPSoC
 - FPGAs running the processing algorithms
 - SoC usage:
 - Coordinate the FPGAs operation, control & monitoring of the blade as well as communicating with DCS
 - *See talk from David Miller on Thursday*
- L1Muon
 - **MUCTPI** – Muon-to-Central-Trigger-Processor-Interface
 - ATCA blade with Xilinx FPGAs (running algorithms, readout and triggering) + Zynq
 - SoC usage:
 - Configuration, control and monitoring of the board
 - Running a RunControl application directly on the SoC
 - *See talk from Ralf Spiwoks today*


***Phase-II:**

- Massive increase in systems integrating SOCs in their designs
- Effort is coordinated with respect to previous phases
- SOC usage – similar functionalities as phase-I:
 - Control, monitoring, configuration of onboard FPGAs and interfacing with DCS

** As we know today, subject to change*



Phase-II upgrade



Common Phase-II SoC usage:

- Interfacing the other onboard components
- Monitoring, control & DCS

LAr:

- 2 new ATCA-based systems:
- LASP – LAr Signal Processor
 - ATCA blade with Intel Stratix-10 SX SoC
 - RTMs with Xilinx Zynq
- LATS – LAr Timing System
 - ATCA blades with controller FPGA, potentially a SoC

Tile:

- Tile PPr – PreProcessor
 - ATCA blades hosting Tile CoM-SoC mezzanine
- Digilent Zybo Z7 as HV control module using Zynq SoC
- *See talk from Fernando Carrio Argos on Thursday*

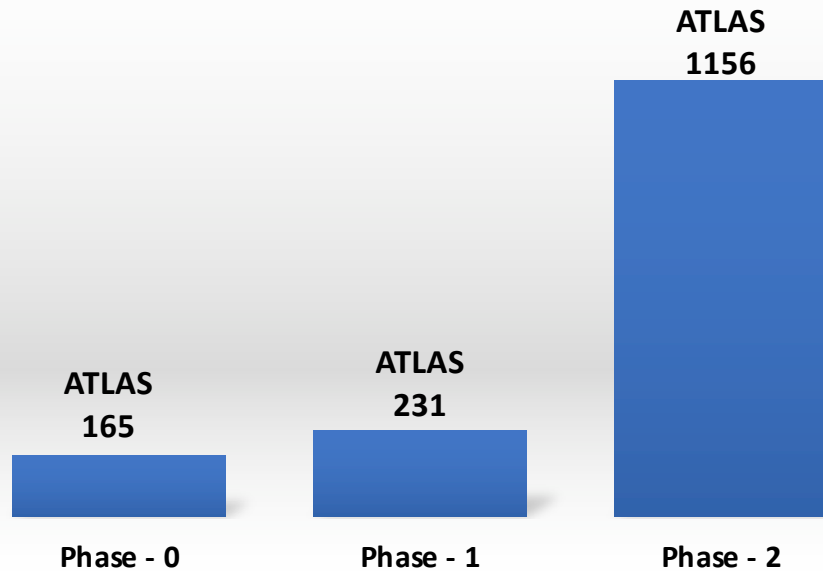
TDAQ:

- All trigger systems electronics will be based on ATCA technology:
 - **L0 Calorimeter trigger**
 - **L0 Muon trigger**
 - **Central trigger**
 - **Hardware Track Trigger**
- ATCA blades with FPGAs for processing algorithms and SoC

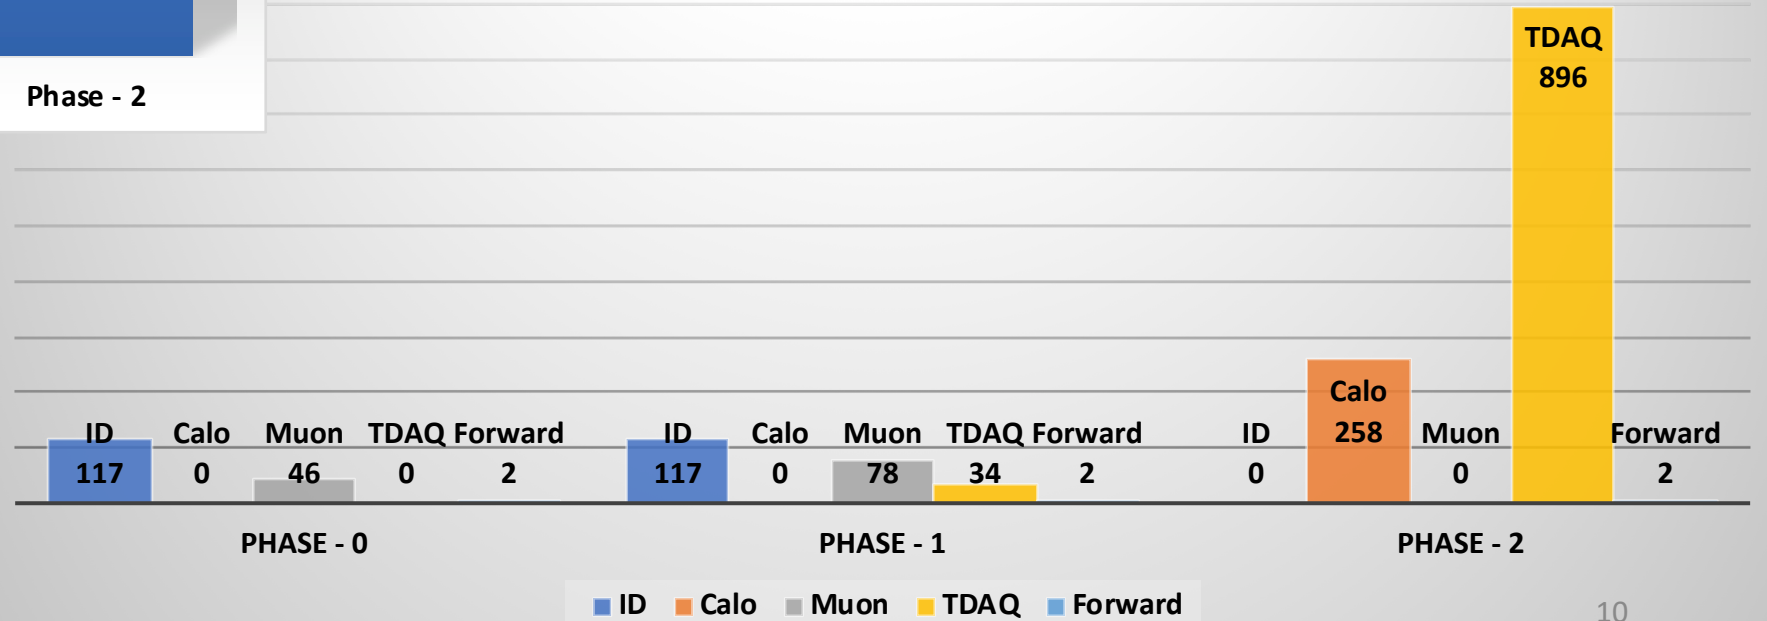
***** NOT FINAL *****
Ongoing design discussions
Final SoC model yet to be chosen

SOC usage evolution in ATLAS

ATLAS SoC usage over upgrade phases

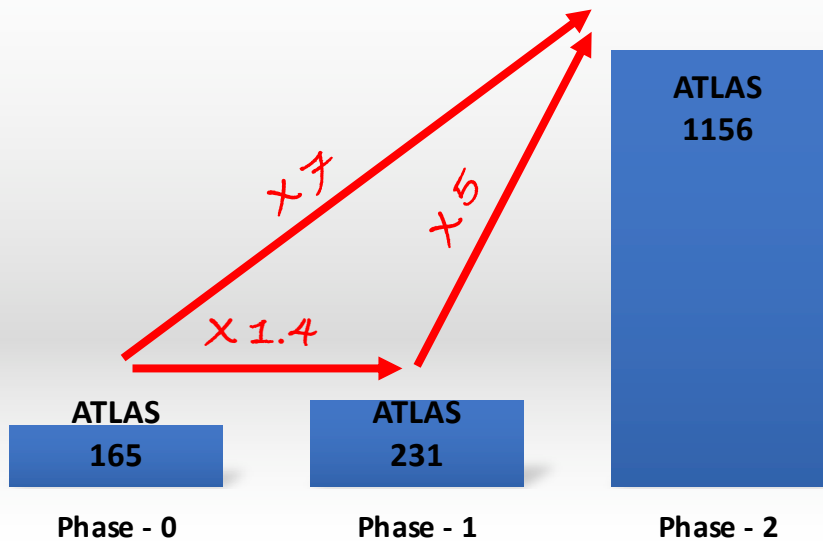


ATLAS detectors SoC usage over upgrade phases




SOC usage evolution in ATLAS

ATLAS SoC usage over upgrade phases



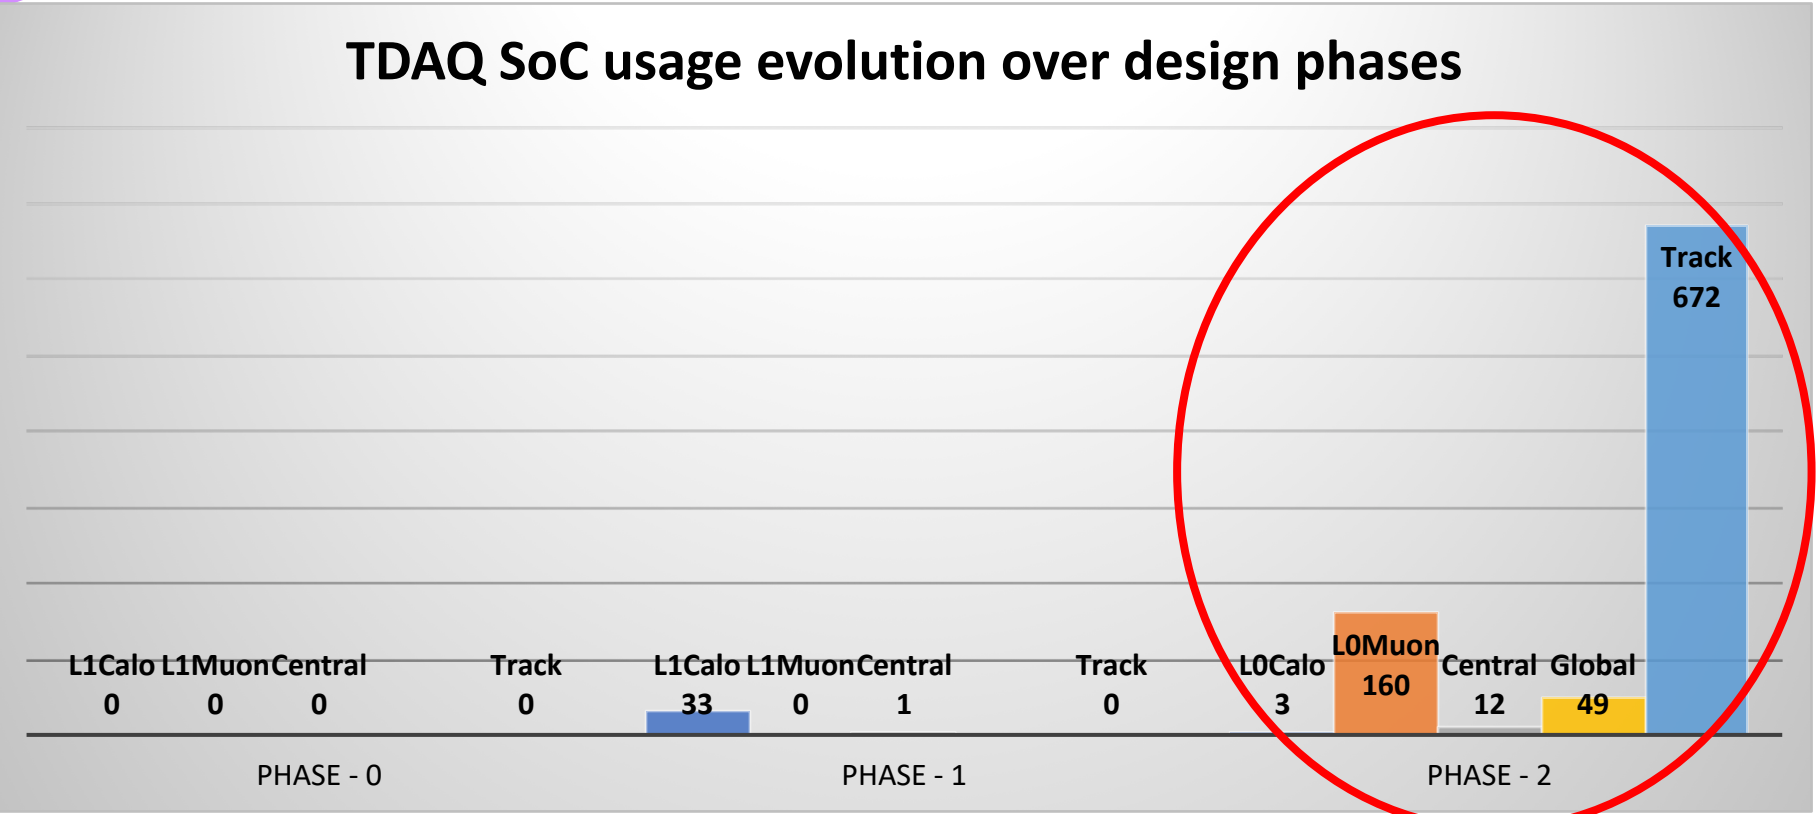
Phase-II TDAQ has the highest (planned) SoC usage in ATLAS

ATLAS detectors SoC usage over upgrade phases



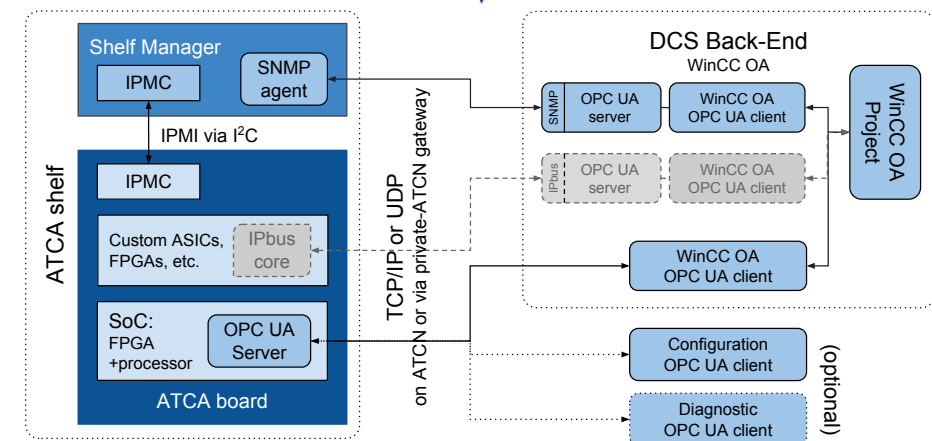
ATLAS Phase-II TDAQ
SoC usage & related coordination

TDAQ SoC usage evolution over design phases



Phase-II TDAQ - SoC functionality

- ~900 ATCA blades with one SoC on each (exception MDT TP - *See talk from Dan Gastelr on Thursday*)
- General ATCA-blade architecture will include both IPMC & SoC, each master of a different I2Cbus:
 - **IPMC I2Cbus** - All critical onboard components following the ATCA compliance.
 - **SoC I2Cbus** - Optical transceivers and any other non-critical info to be monitored.
- Two interfaces to the DCS backend: IPMC & SoC
 - The SoC allows **more flexibility in the number of monitored parameters** (OPC UA Server)
 - OPC-UA server built using **quasar** (Quick OPC-UA server generation framework)
See more on Friday morning: S. Schlenker's talk + Tutorial from P. Nikiel
- Control and monitoring functionalities implemented with dedicated hardware connections
 - **Specific per board**, with many commonalities (see later)
 - Take advantage of **programmable logic for interfacing FPGA(SoC)-to-FPGA**
- Baseline plan for **online software** to **provide a common communication library**
 - **Minimize software dependencies** in the SoC domain
- **Realtime processing** is being considered by few systems, e.g. ATLAS Global trigger-processors considering implementation of topological real-time algorithms



Phase-II TDAQ – challenges & coordination

Two main challenges:

- **SoC connection to the ATLAS Control Network (ATCN)**
 - **Long-term maintenance & support**
- Overcoming the challenges will most likely rely on commonality across-systems

Related actions:

- **Define SoC-user requirements** in an official document
- **Dedicated discussions** involving our **systems experts & TDAQ Phase-II coordinators**, allowing both **design, integration & commissioning concerns** to be raised
- **SoC-survey** - a questionnaire prepared by several system engineers & TDAQ coordinators
 - *Provides a better understanding on the different systems requirements & wish-list*
 - *Attempt to spot commonality*
- **Dedicated test-rig** – Currently under construction
 - ATCA-related R&D studies, for testing & evaluation of common-related proposed solutions (IPMC, SoC flavors, DCS tools, etc...)
 - Mimicking ATCA environment in ATLAS counting room (USA15), for testing SoC OS management by sysAdmin

SoC-user requirements – *Mezzanine*

- Concerns & motivation:
 - *Phase-II hardware lifetime (>10 years) might require a replaceable support*
 - *Software lifecycle is usually shorter (few years)*
 - *Dependency on the vendor's software support*
- Implications:
 - *OS freezing and most-likely network isolation*

Requirement 2.x: Soc integration on ATCA blades

ATCA blades should integrate SoC devices through mezzanines in order to guarantee upgradeability during the HL-LHC lifetime.

In case the implementation of a mezzanine is not possible, the device may have to be isolated from the control network during its lifetime. It must be possible to operate the system under the most constraining isolation scenario, in which case the system would become accessible exclusively through a gateway machine. Hence it shall be ensured that neither software nor hardware limitations (e.g. missing rack space for additional components) would prevent this. The system will be responsible for the development and maintenance of any additional software layer required in order to interface with common software tools (online-software libraries, OPC-UA server for DCS, etc...).

Updated version is currently under approval process

SoC-user requirements – *Common OS*

- Motivation for common-OS:
 - Mostly *long-term maintenance - OS patching & upgrading regularly*
 - *Allowing central OS support within ATLAS or from CERN will also enable direct connection of the device to the ATLAS Control Network (ATCN)*
- Converging on a common OS while considering both internal & external parties:
 - *TDAQ-community within ATLAS, CERN IT, CERN Security*

Requirement 2.7: Common Operating System

All sub-systems implementing SoCs devices shall utilize a common OS and BSP.

Requirement 2.8: Choice of the operating system

The choice of the OS to be deployed on each SoCs falls under the responsibility of the TDAQ Phase-II UPR Technical Coordinator, after consultation with the UPR sub-systems, TDAQ SysAdmins, CERN IT, and ATLAS management at large.

SoC Survey - Summary - (1/2)

- **Choice of a chip:**

- *Usage* - Mainly for monitoring & control, with the programmable logic used to interface with the other onboard components (*exception* - Realtime processing - Global & MDT TP)
- *No. of GPIOs* - max 150 (*exception*: Global 300)
- *No. of MGTs* - less than 10 (*exception* - Global needs 72)
- *limitations/wish-list* - RAM up to 4GB (*exception*: Global - 16 GB), 2 Ethernet ports, SD-card
- *preferred computing architecture* - ARM 64-bit SoC

→ **The input shows similar requirements among most systems.** A distinction has been mainly spotted in systems that consider to implement realtime processing in SoC devices

- **Choice of a mezzanine:**

- *Limitations* - small form-factor, 12V powering of the mezzanine with on-board I/O power supplies
- *Important peripherals* - Supports both local-flash & SD-card, All spectrum of I/O interfaces: UART, I2C, SPI, GPIO, JTAG, AXI

→ **Most systems are not objecting to deploy a mezzanine** as long as the minimal requirements are set and **the size does not imply re-design of the blade**

SOC Survey - Summary - (2/2)

- Choice of an OS:

- *Initial-preference* - Linux-based OS: most mentioned CentOS
- *Wish-list* - ssh-access, sudo-rights (*exception*: Realtime-OS for processing - Global, MDT TP)

→ **Converging on common OS seems feasible: Full agreement on Linux-based OS.**

- Considering **Realtime processing and regular-processing** requests → expect at least **two flavors**

- Connection preference to the ATCN :

→ **The choice of either direct or isolated connection should be taken by subsystems based on use cases** (e.g. Realtime) and needs

- Corresponding implications and **requirements have to be taken into account** (see slides 15 & 16)

Summary

- **Increase in SoC usage** across ATLAS systems over the phases of design
- Highest usage identified in **Phase-II upgrade of TDAQ**
- **Main challenges** identified:
 - *SoC connection to the ATLAS Control Network (ATCN)*
 - *Long-term maintenance & support*
- **TDAQ have taken** the initiative to find **common solutions**
 - Attempt to overcome the challenges by **deploying a mezzanine & converging on common SoC-OS seems feasible**
- **Open discussion with CERN IT, LHC experiments** and LHC departments (this workshop)
- **To avoid a sharp increase in required ATCN connectivity, a solution is being investigated to connect the SoC on the individual blades to an in-shelf switch.** If a direct connection to the ATCN is necessary, a solution has **to be discussed with ATLAS TC**
- **ATLAS management considering the possibility to adopt the TDAQ approach as ATLAS standard for all upgrade projects**