ATLAS SoC usage overview

System-on-Chip Workshop - 8th June 2021 - CERN

https://indico.cern.ch/event/996093/

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

Many thanks to:

C. Gemme, O.Kepka, A. Straessner, T. Hrynova, F. Carrio Argos, F. Martins, O. Solovyanov, M. Corradi , P. Fleischmann,, E. Pasqualucci, K. Polawski, T. Costa De Paiva, M. Ishino, L. Levinson, Y. Okumura, A. Tanaka, R. Vari, A. Annovi, G. Avolio, M. Begel, B. Gorini, E. Hazen, R. Middleton, D. Miller, T. Mkrtchyan, W. Panduro-Vazques, D. Sankey, D. Scannicchio, U. Schafer, W. Vandelli, S. Veneziano, M. Warren, B. Giacobbe, J. Guimaraes de Costa, M. Kocian, K. Korcyl, , F. Lasagni, M. Wittgen, S. Schlenker, S. Haas, R. Spiwoks, F. Lanni

for the info, help and cooperation in preparing this summary ☺

ID, LAr, Tile, Muon, Muon-TDAQ, TDAQ, FW, TC, ESE-TDAQ, UC

Outline

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

LHC / HL-LHC Plan

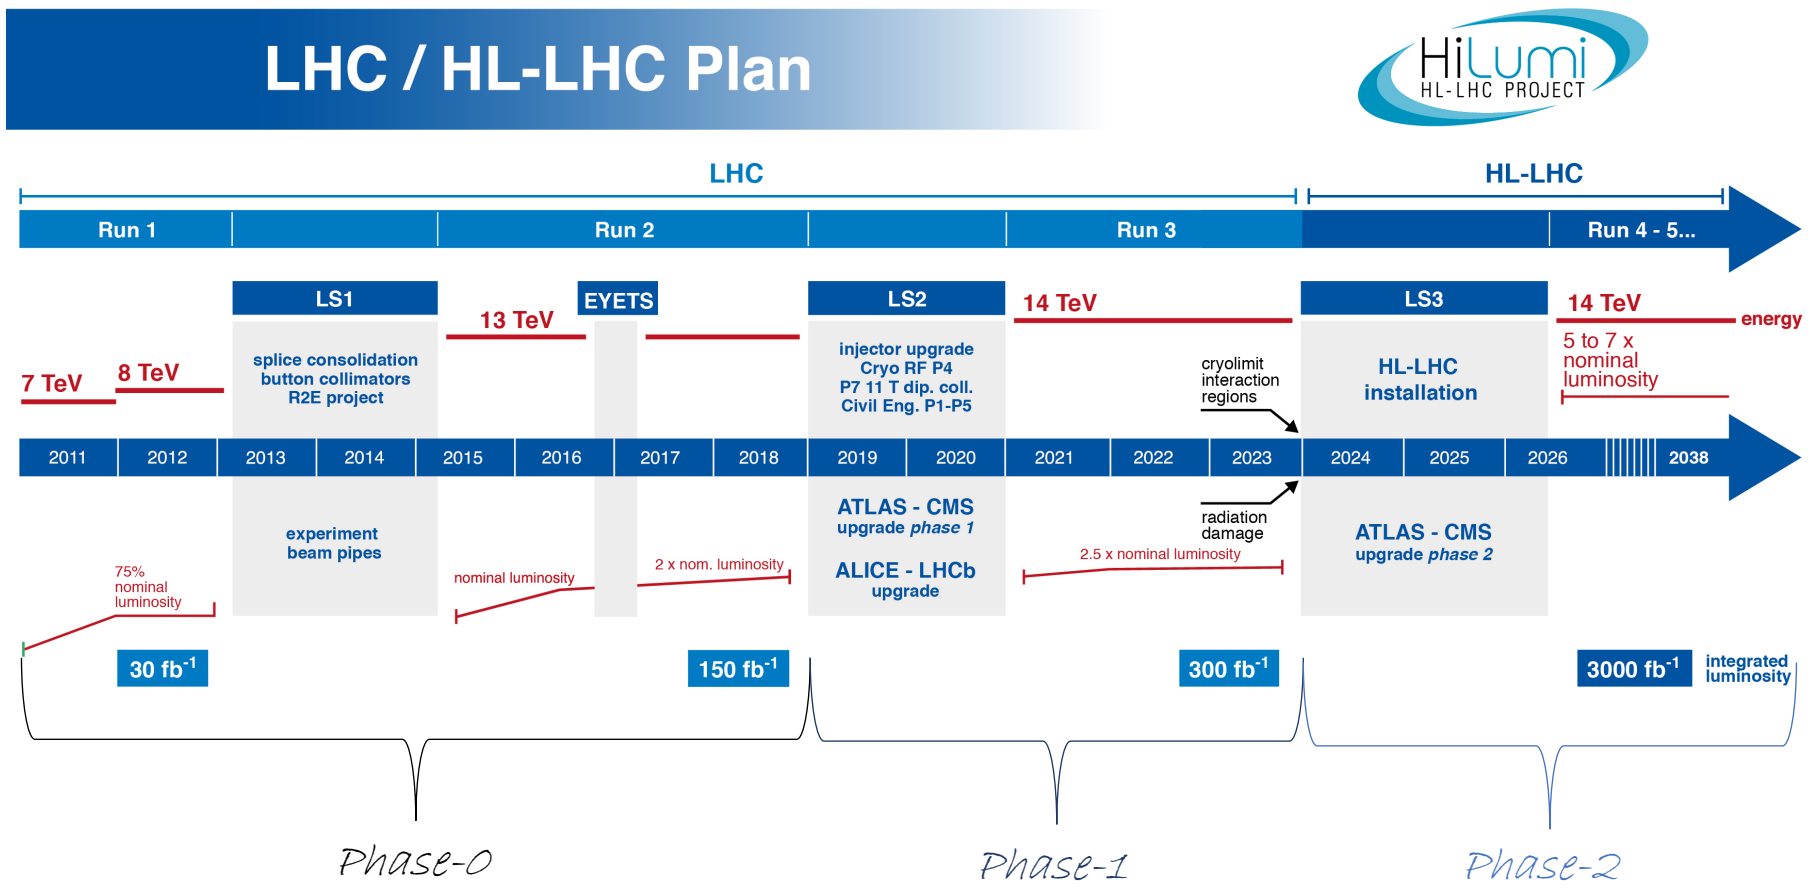

ATLAS SoC usage evolution over the upgrade phases

Evolution: Semi-Spoiler ☺

* As we know today, subject to change

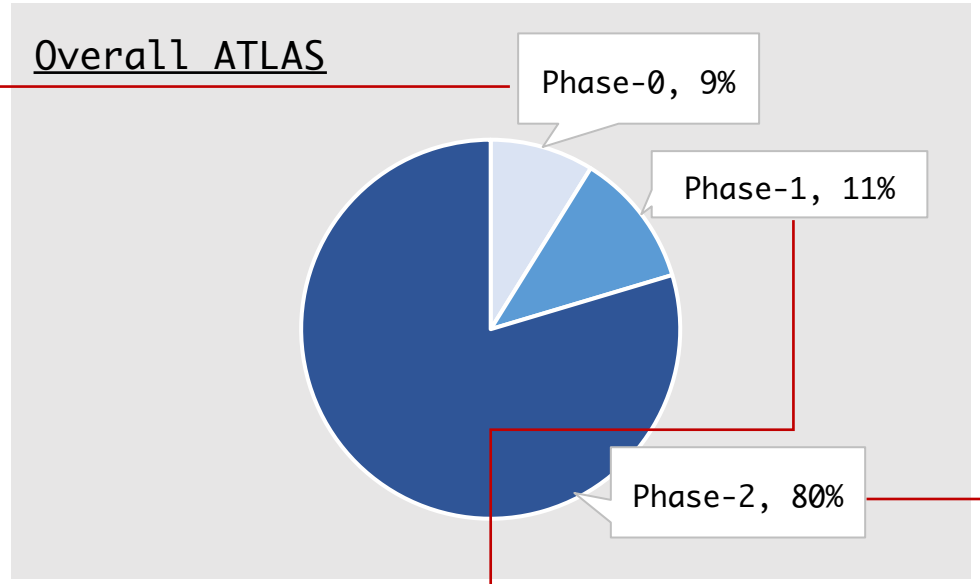

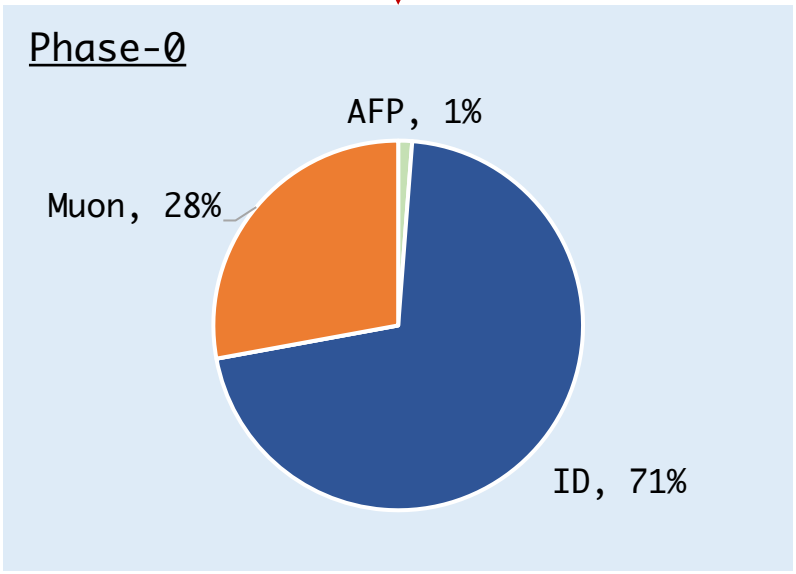

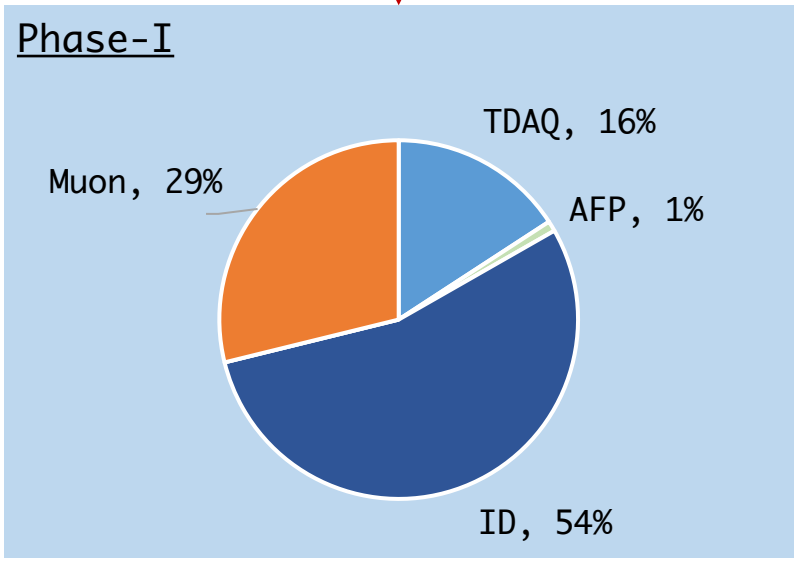

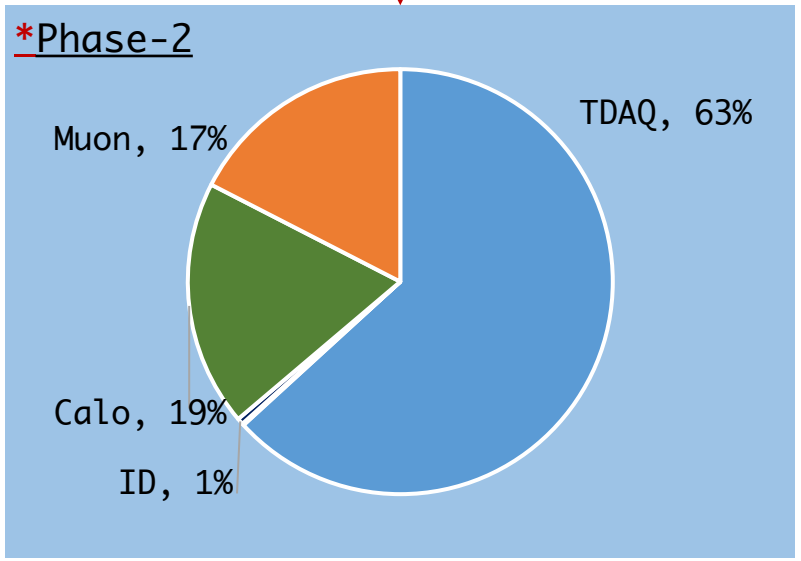

Run-2:

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

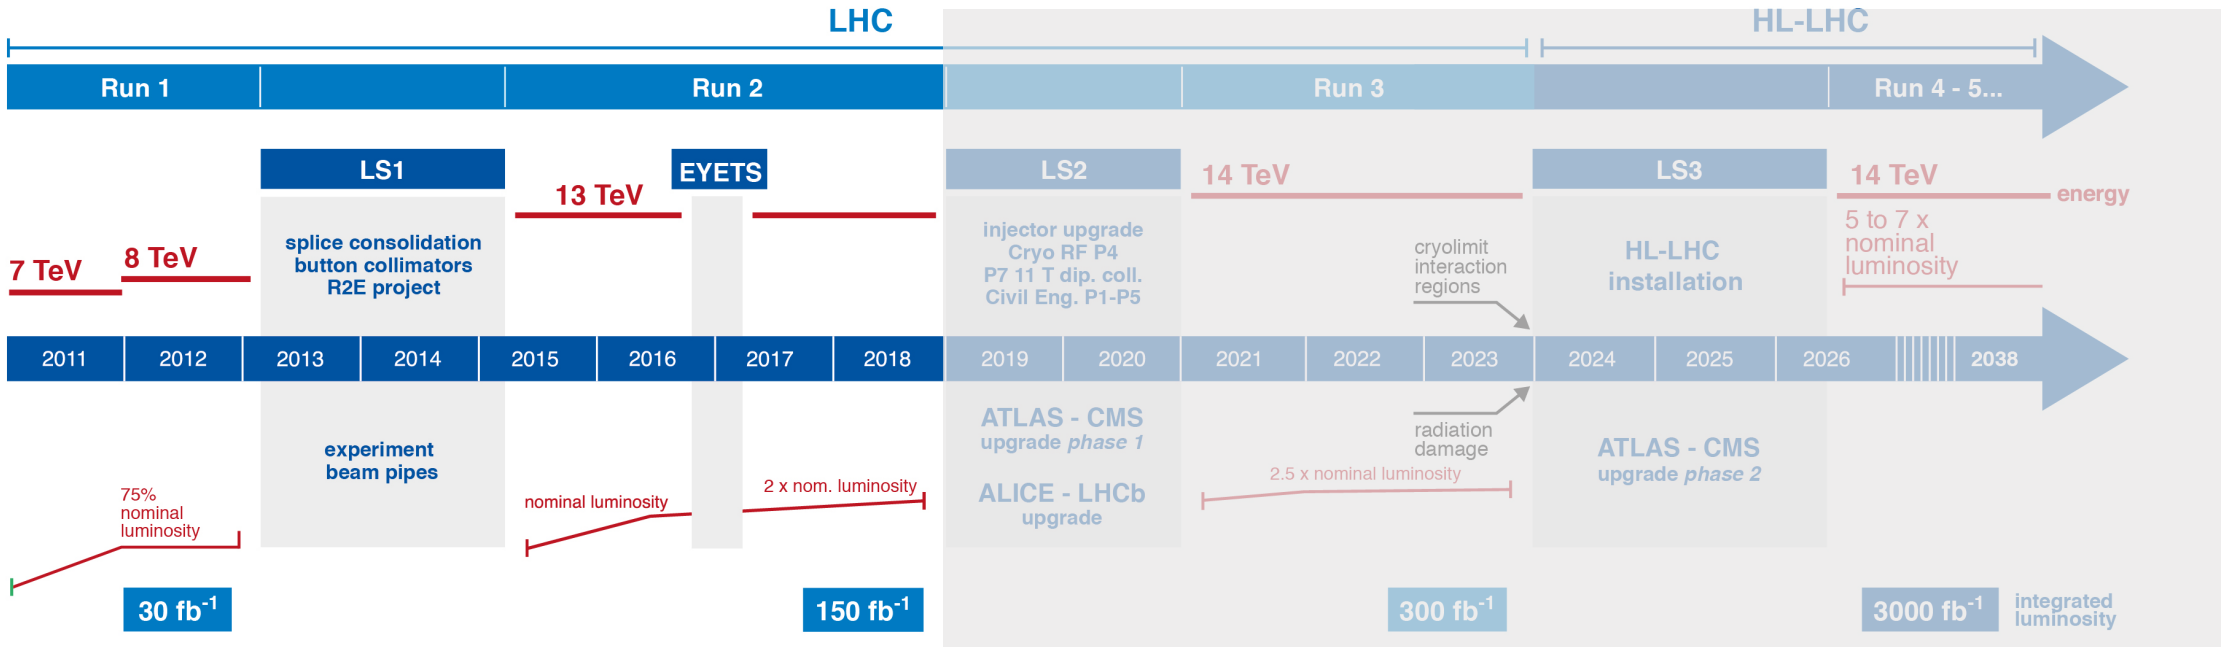

Phase-0 and up to end of Run-2

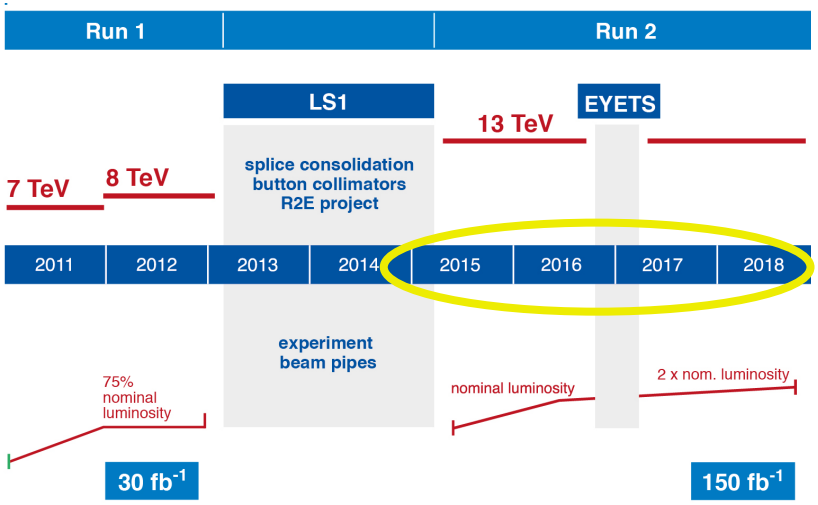

Inner Detector:

- 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
    - program slave (Spartan) FPGA from the SoC PPC440 directly
  - Will run until the end of Run-3
  - *See talk from **Oldřich Kepka** 1st SoC workshop*
  - *See talk from **Matthias Wittgen** 1st Workshop*

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** 1st Workshop*
  - *See talk from **Zijun Xu**, later today*

- 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 Xlinx Artix FPGAs, and Zynq-7000
  - SoC usage:
    - Control & configuration of the FPGAs
  - *See talk from **Matthias Wittgen** 1st Workshop*

Phase-I:

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

Legacy SoCs from Run-2:

- Pixel/IBL ROD
- MDT Alignment
- AFP Readout

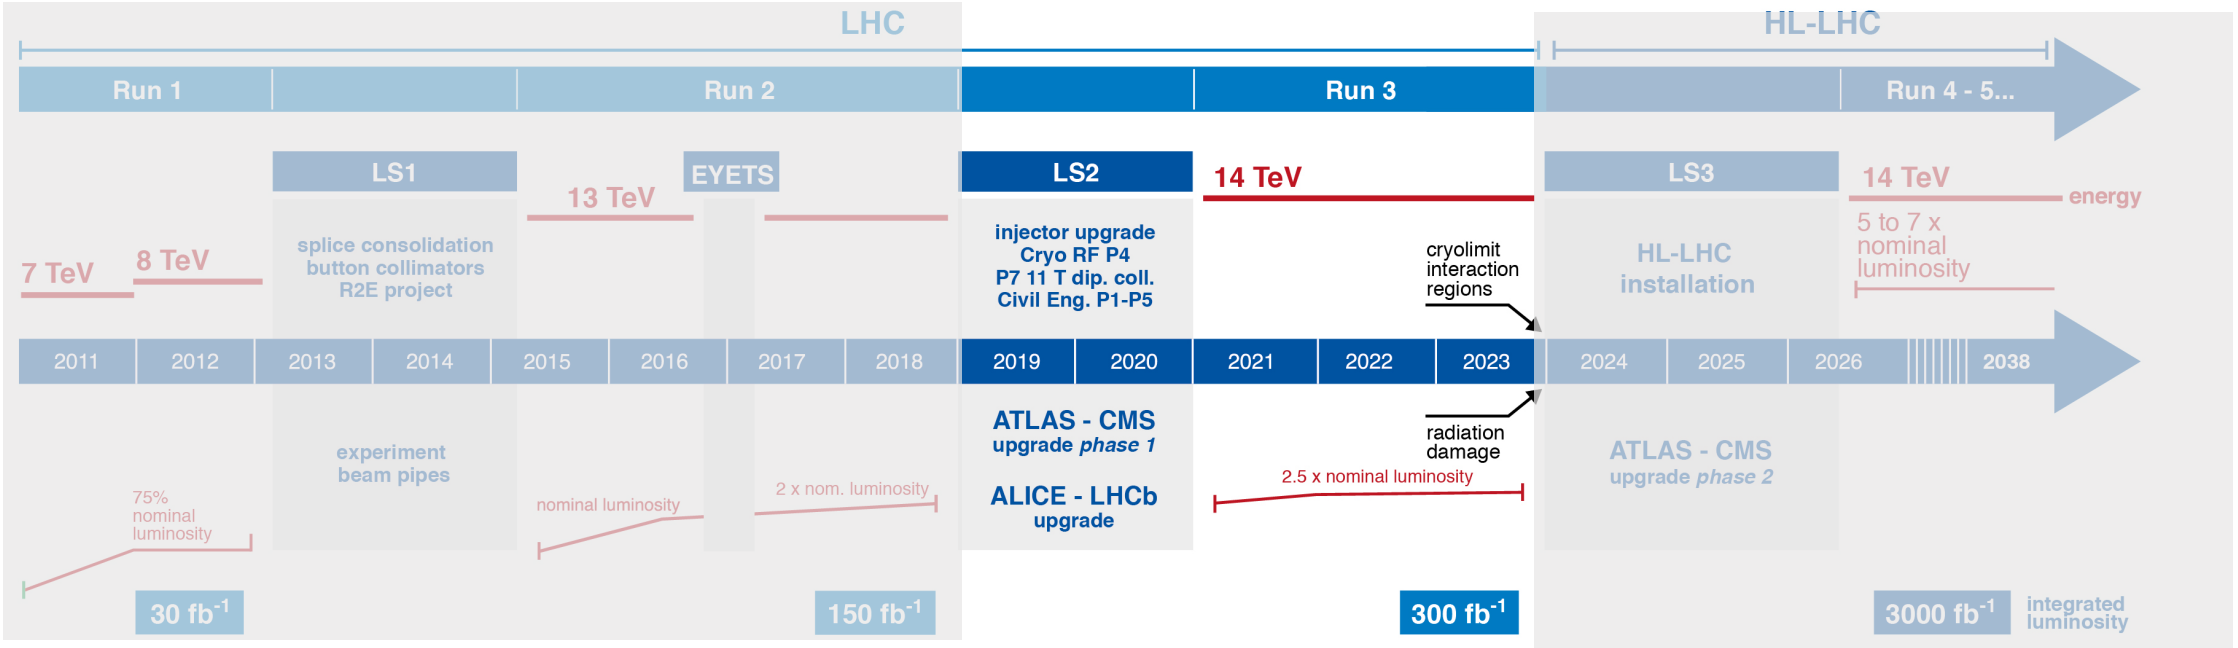

Phase-I upgrade

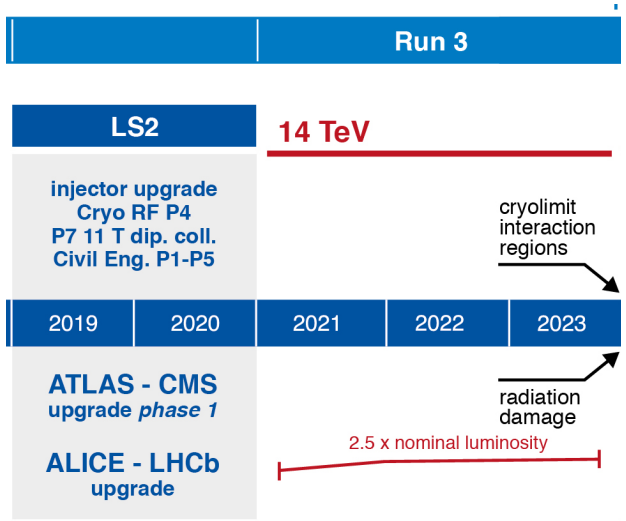

Legacy SoCs from Run-2:
- Pixel/IBL ROD
- MDT Alignment
- AFP Readout

Muon Detector:
- NSW – New Small Wheel
  - The NSW Trigger Processors (TP) are ATCA based
  - Blades with Xilinx US FPGAs for running algorithms and a Zynq-7000
  - SoC usage:
    - HW control and monitoring

Trigger-DAQ (TDAQ):
- L1Calo
  - TREX – Tile Rear Extension
    - VME digital modules
    - Zynq US+ MPSoC on each
    - SoC usage:
      - Configuration, control, monitoring, and communication with the DCS BE
    - *See talk from Tigran Mkrtchyan later today*
  - 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 BE
    - *See talk from Emily Smith later today*
- L1Muon-Central
  - MUCTPI – Muon-to-Central-Trigger-Processor-Interface
    - ATCA blade with Xilinx FPGAs (running algorithms, readout and triggering) + Xilinx Zynq US+
    - SoC usage:
      - Configuration, control and monitoring of the board
      - Running a RunControl application directly on the SoC
    - *See talk from Ralf Spiwoks on Wednesday*

*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, DCS interface

*As we know today, subject to change

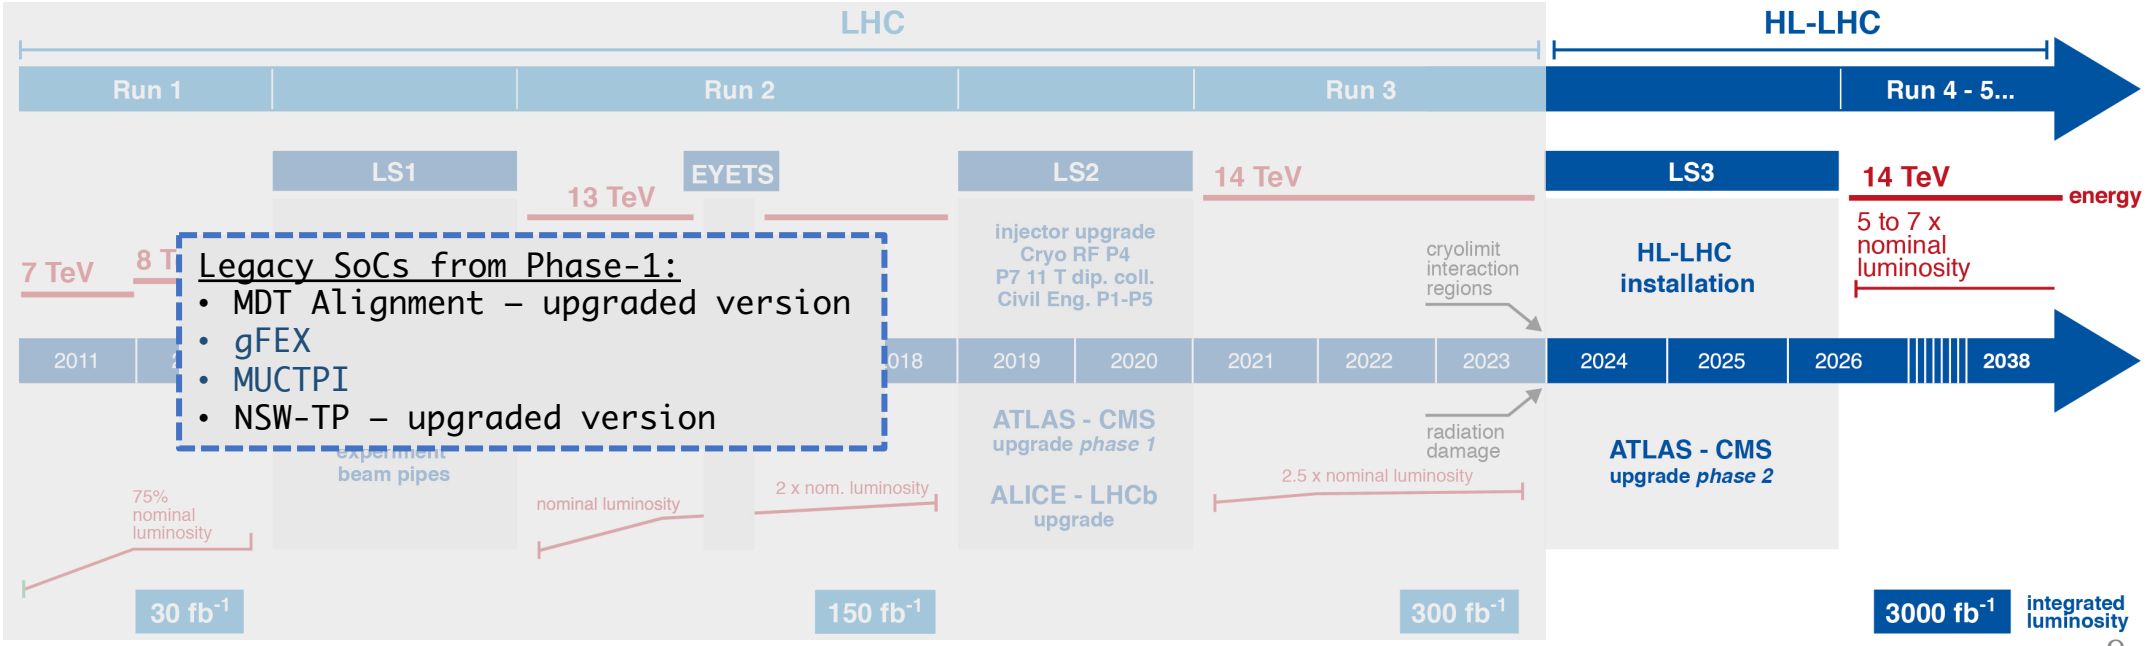

Phase-II upgrade

Non ATCA systems

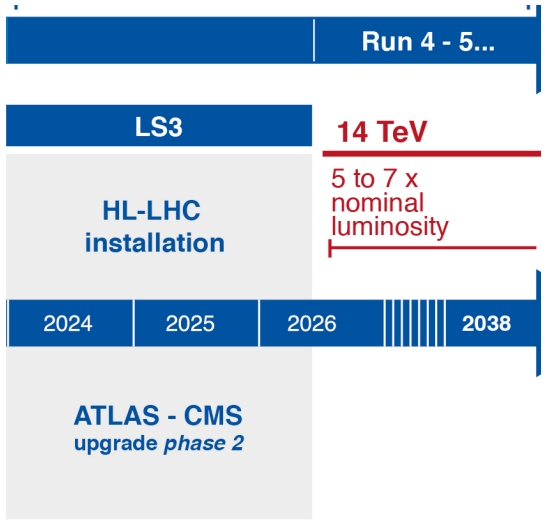

*** NOT FINAL ***
Ongoing design discussions
Final SoC model yet to be chosen

Muon detector:
- Thin Gap Chambers (TGC)
  - **Charge Monitoring system:**
    - Charge Monitoring Boards (CMB): VME modules with Trenz SoM with Xilinx Zynq-7000
    - SoC usage:
      - control and monitoring of the CMB operations
  - **JATHub**
    - VME module with Xilinx Zynq-7000
    - SoC usage:
      - Configuration, control and monitoring of detector front-end FPGAs
      - Managed remotely via a host-PC
      - *See talk from **Aoto Tanaka** on Wednesday*

Legacy SoCs from Phase-1:
- MDT Alignment – upgraded version

DCS interface – EMCI/EMP

EMCI - Embedded Monitoring & Control Interface
EMP - Embedded Monitoring Processor

- EMCI to front-end detector via e-links:
  - **Inner-Tracker (ITk)**
  - **Tile – Cs calibration**
  - **Thin Gap Chambers (TGC) - charge monitoring**
  - **High-Granularity-Timing Detector**
- EMCI to EMP via optical links
- Xilinx Zynq US+ on the EMP running OPC-UA and interfacing into the DCS back-end
- Control & monitoring services
- *See talk from **Paris Moschovakos** on Tuesday*

Phase-II upgrade

ATCA systems

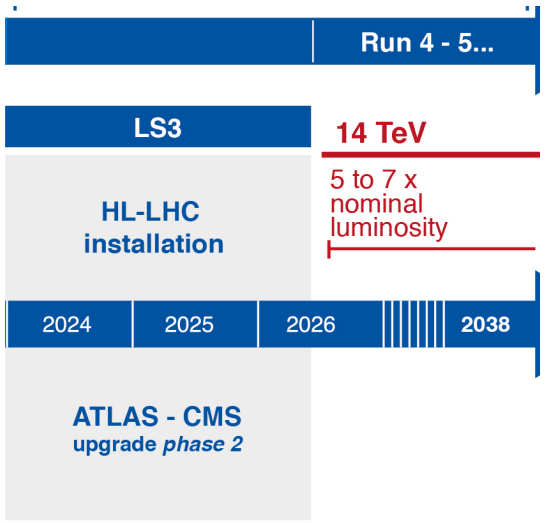

Legacy SoCs from Phase-1:
- gFEX
- MUCTPI
- NSW-TP – upgraded version

Common Phase-II SoC usage:
- Interfacing the other onboard components
- Monitoring, control, interfacing with the DCS BE

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

LAr:
- 2 new ATCA-based systems:
- **LASP** – LAr Signal Processor
  - ATCA blade with 2 Intel Stratix-10 FPGAs
  - RTMs with Xilinx Zynq US+
  - *See talk from **John Hobbs & Dean Schamberger** on Wednesday*
- **LATS** – LAr Timing System
  - ATCA blades with Enclustra SoM MARS XU3 including Xilinx Zynq US+

Tile:
- **Tile PPr** – PreProcessor
  - ATCA blades hosting TileCoM mezzanine with Xilinx Zynq US+
  - *See talk from **Mpho Gift Gololo** later today*
- **Zybo Z7** as HV control module with Xilinx Zynq SoC

TDAQ:
- All trigger systems electronics will be based on ATCA technology:
  - **L0 Calorimeter trigger**
  - **L0 Muon trigger**
    - MDT Trigger Processor - *See talk from **Dan Gastler** on Wednesday*
  - **Central trigger**
  - **Global trigger**
  - **Hardware Track Trigger**
- ATCA blades with FPGAs for processing algorithms and SoC

SoC usage evolution in ATLAS

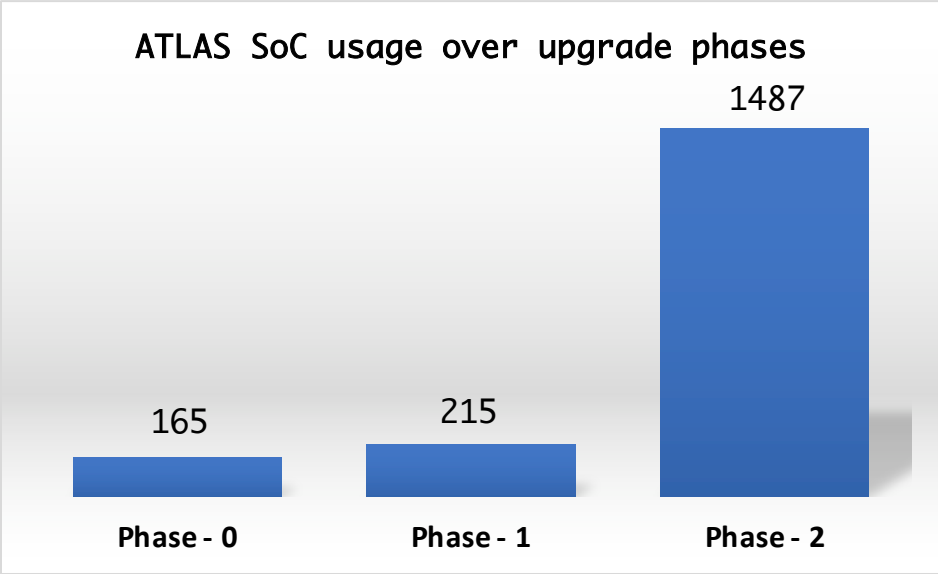

Main usages:

- Slow control, monitoring
- Other onboard FPGAs configuration
- Running algorithms
- Readout
- DCS functionalities

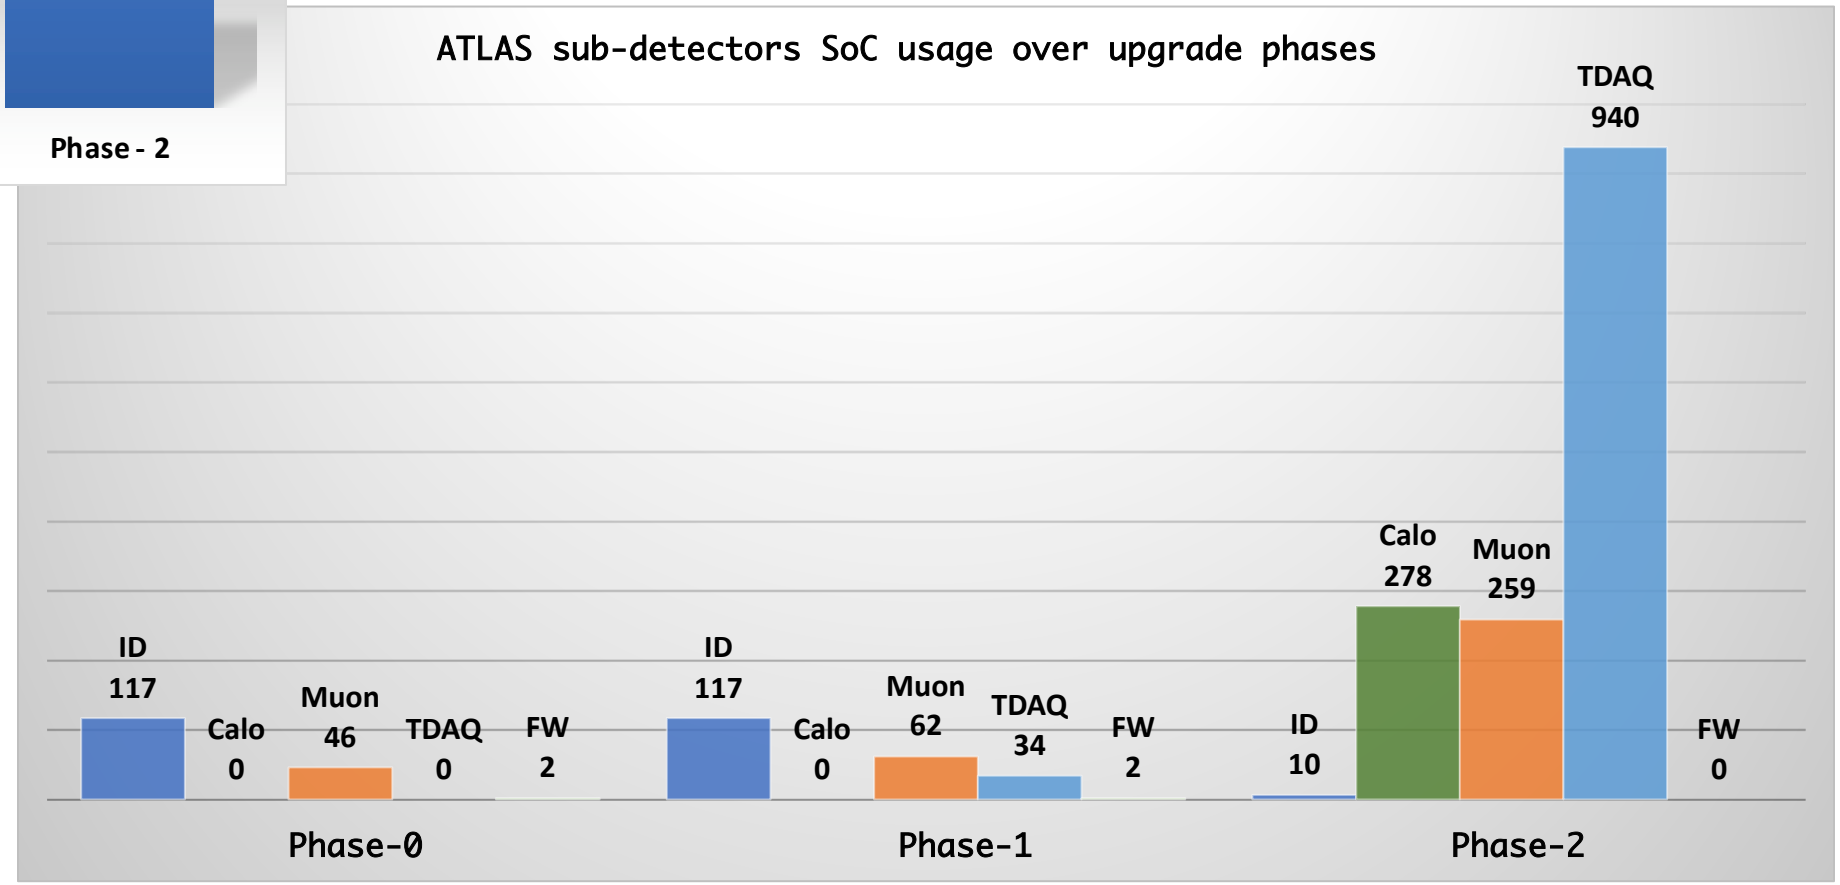

*Phase-2: As we know today, subject to change

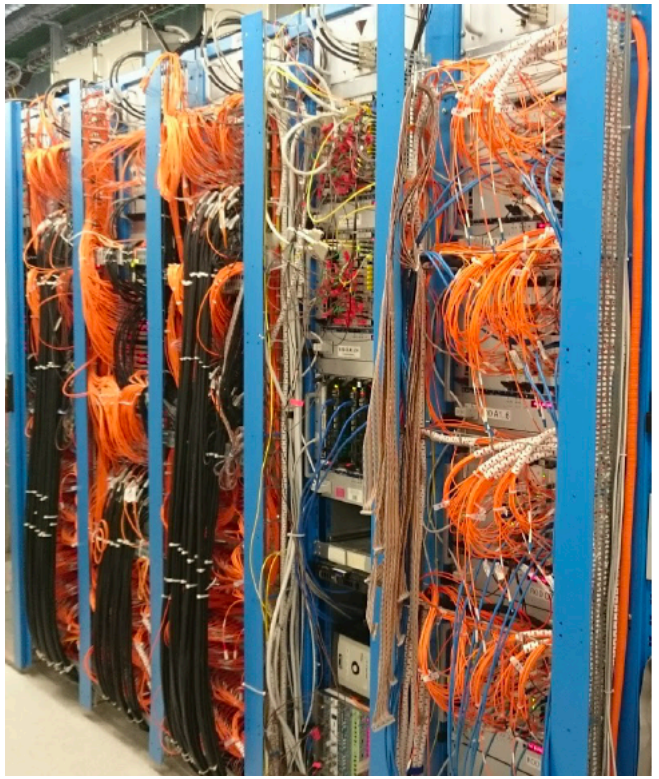

©W. Vandelli

Associated challenges & coordination of related solutions

SoC usage evolution in ATLAS

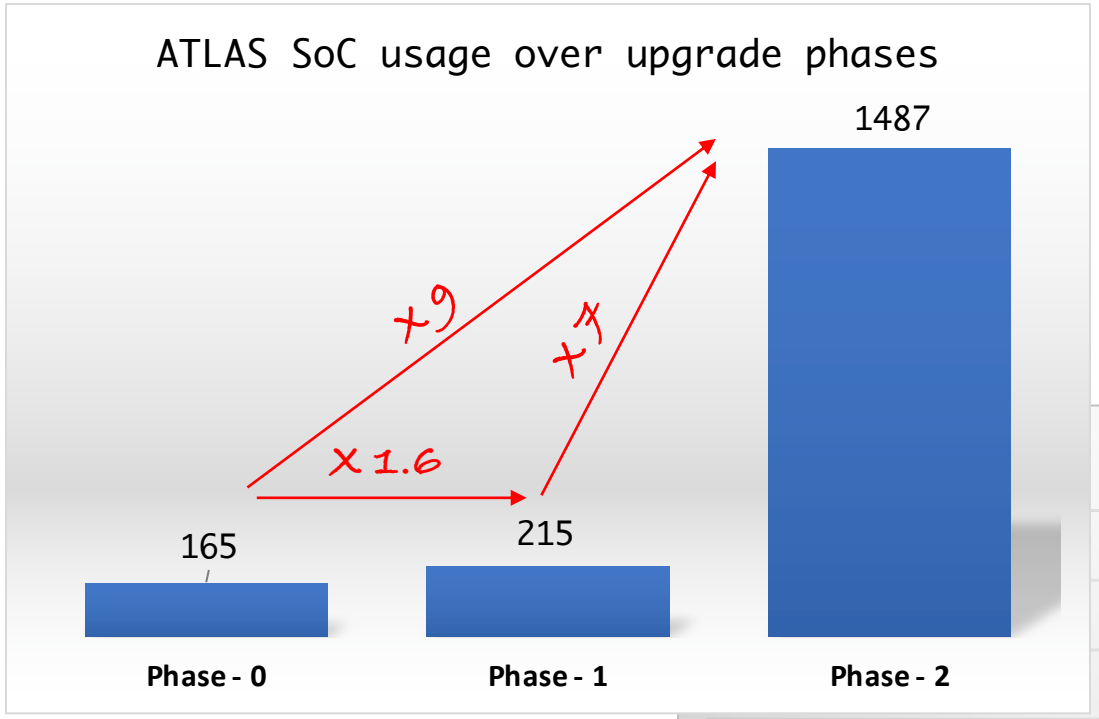

Phase-II TDAQ has the highest (planned) SoC usage in ATLAS

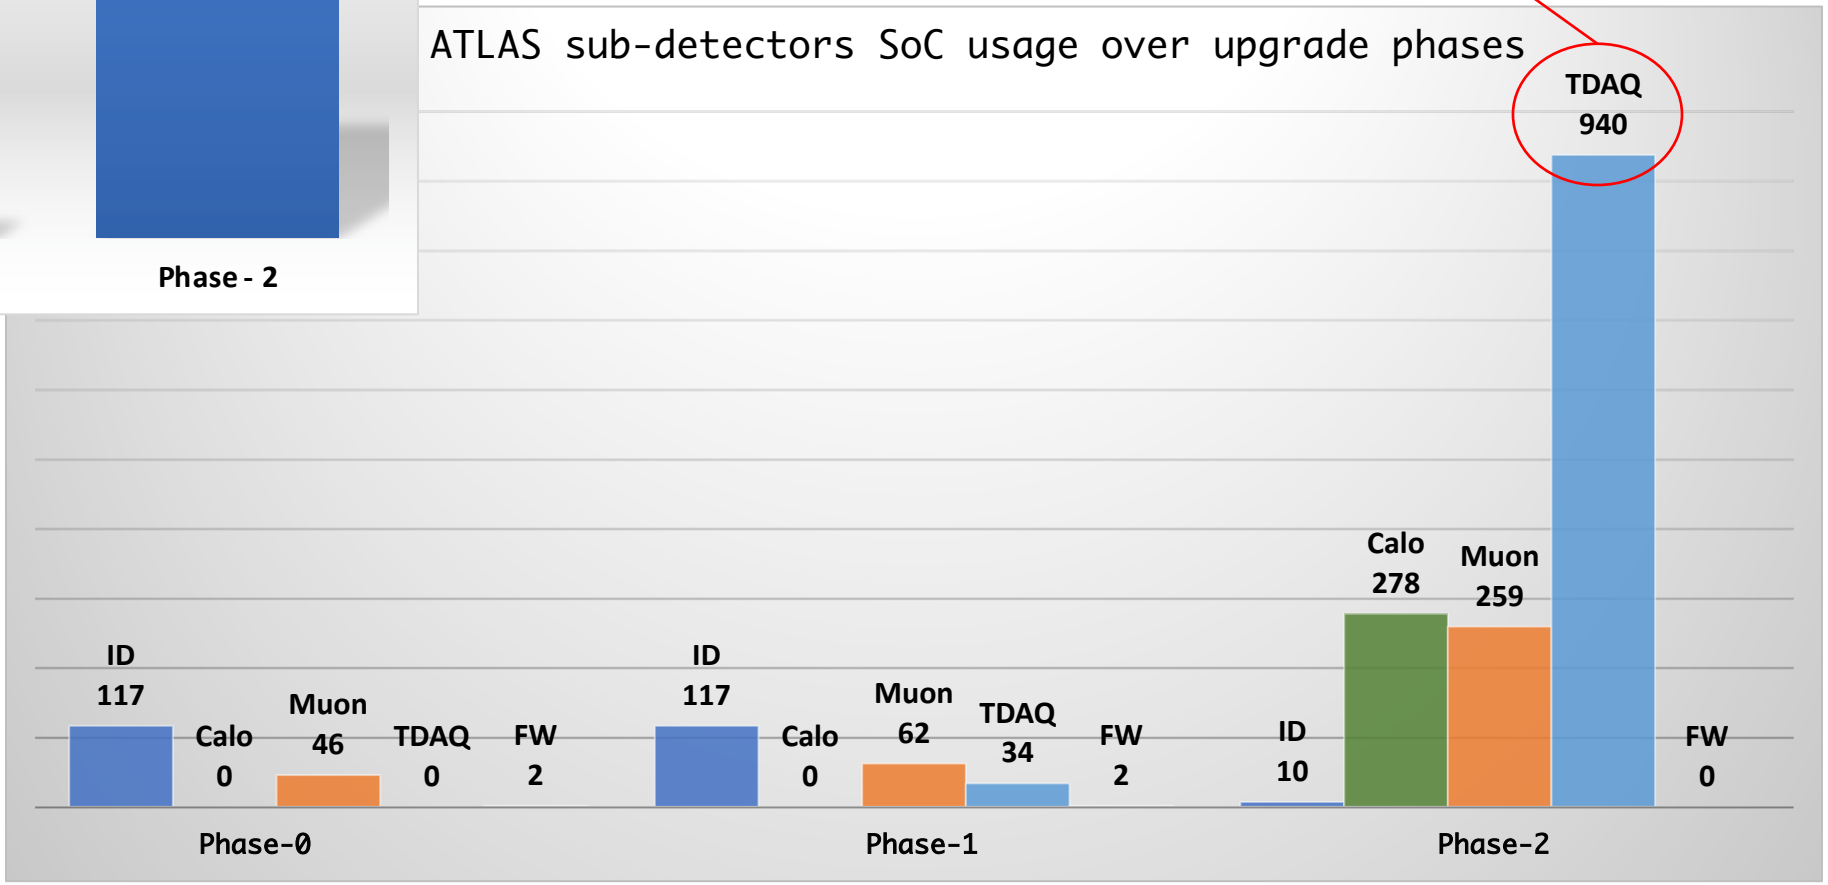

*Phase-2: As we know today, subject to change

ATLAS Phase-II TDAQ SoC usage & related coordination

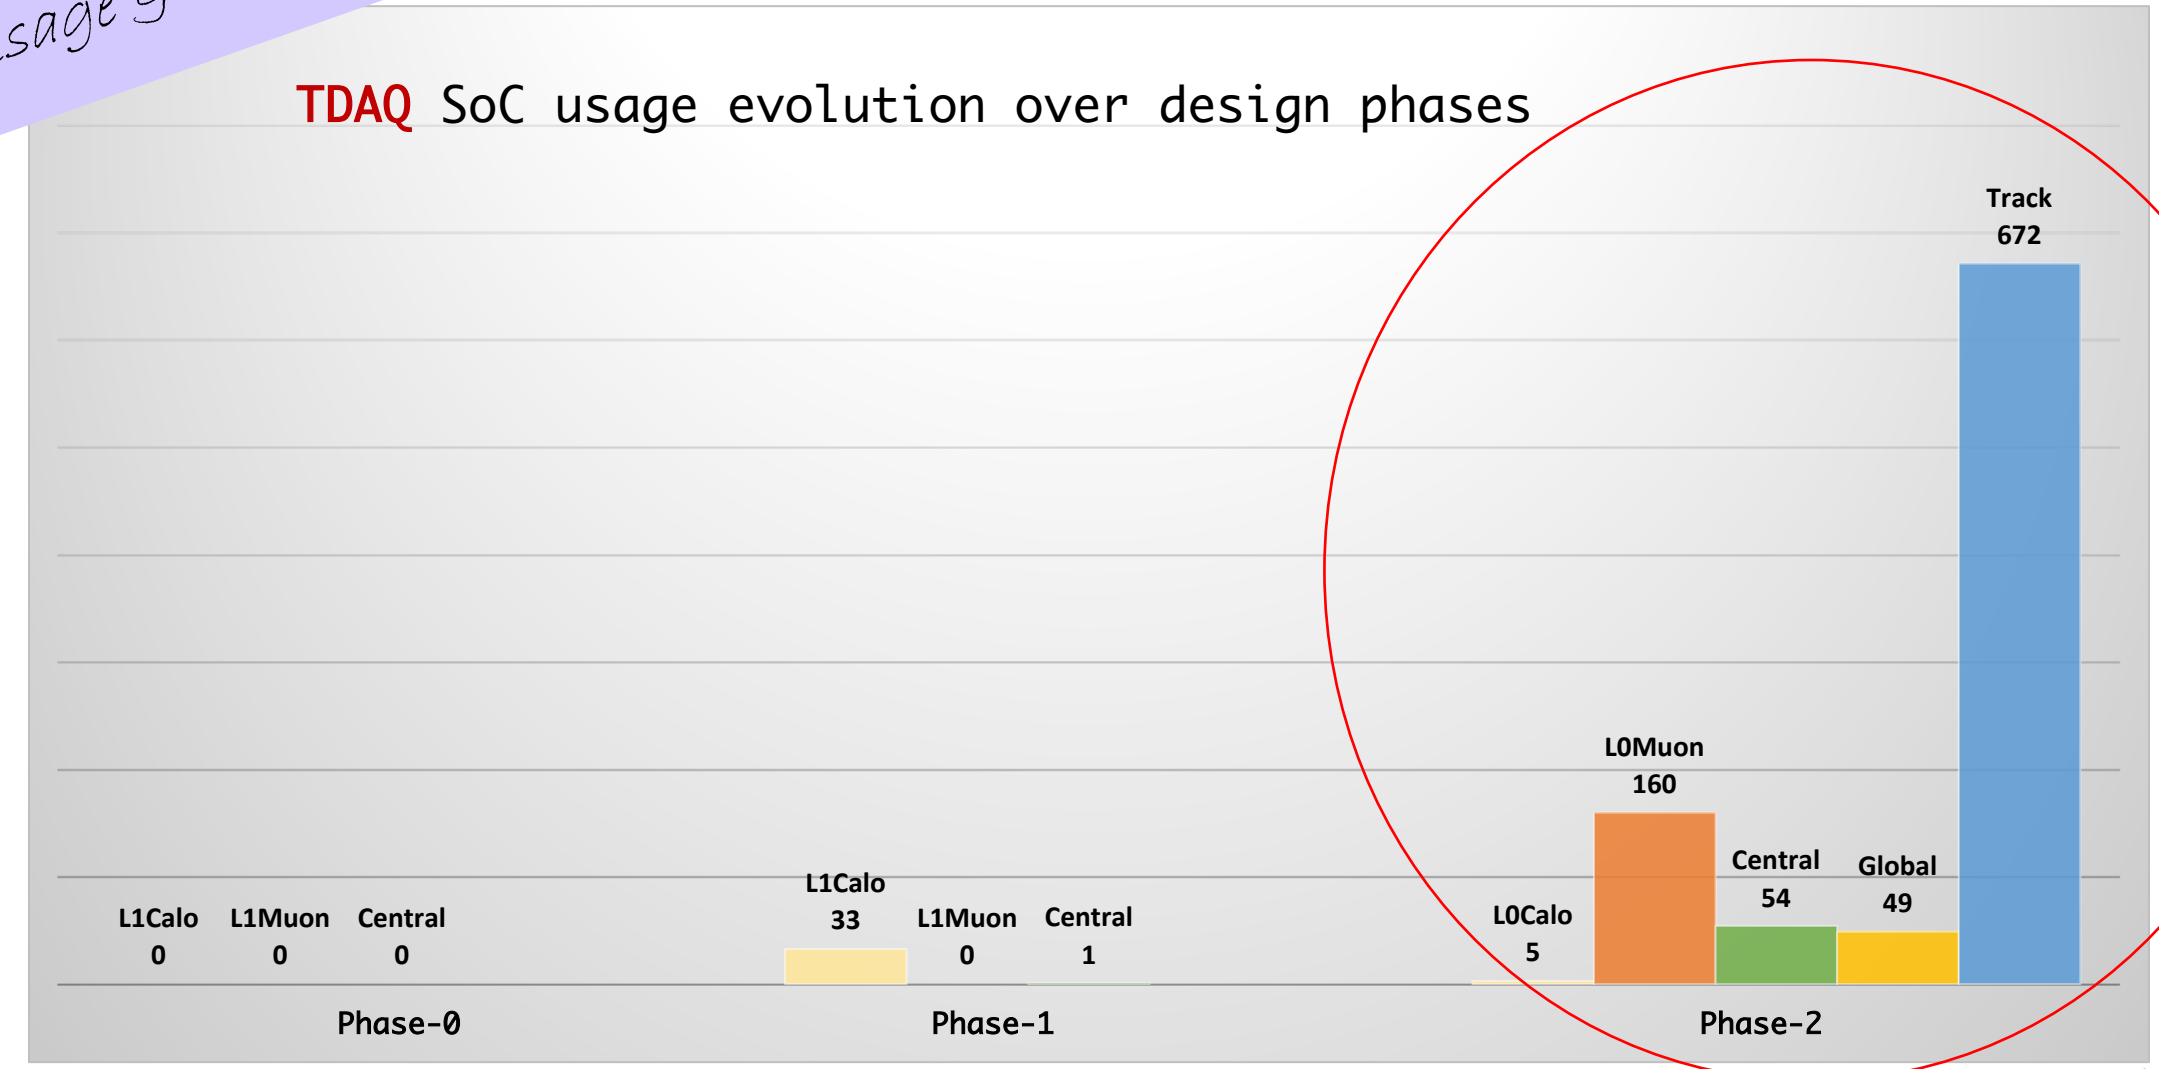

*Phase-2: As we know today, subject to change

Phase-II TDAQ – challenges & coordination

- ~900 ATCA blades with one SoC on each

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:

- Dedicated discussions involving our **systems experts, TDAQ Phase-II & external-services coordinators**, allowing both **design, integration & commissioning concerns** to be raised

- SoC-survey - a questionnaire prepared by several system engineers & TDAQ coordinators (see slide 18)
  - *Provides a better understanding on the different systems requirements & wish-list*
  - *Attempt to spot commonality*

- Define SoC-user requirements in an official document (see slide 19-23)
  - Reflecting the **common systems needs** and establish **a uniform baseline**
  - **Cross-referenced** in other official requirement documents (systems design-reports, networking, DCS, etc…)

- Dedicated test-rig
  - 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**

Phase-II TDAQ – SoC on ATCA blades

- 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.

- Control and monitoring functionalities implemented with dedicated hardware connections
  - **Specific per board**, with many commonalities
  - Take advantage of **programmable logic for interfacing FPGA(SoC)-to-FPGA**

- External services:
  - Centrally managed by ATLAS TDAQ sysadmin
    - More info in the next slides
    - *See talk from **Marc Dobson & Diana Scannicchio** on Friday*
  - Interface of ATCA blades into the DCS backend:
    - IPMC (via Shelf-Manager) & SoC
    - The SoC allows **flexibility in the number of monitored parameters** (OPC UA Server)
    - OPC-UA server built using **quasar** (Quick OPC-UA server generation framework)
    - *See talk from **París Moschovakos** on Tuesday*
  - SoC- ATLAS DAQ interface
    - Baseline plan for online software to provide a common communication library
    - **Minimize software dependencies** in the SoC domain
    - *See talk from **Andrei Kazarov** on Friday*

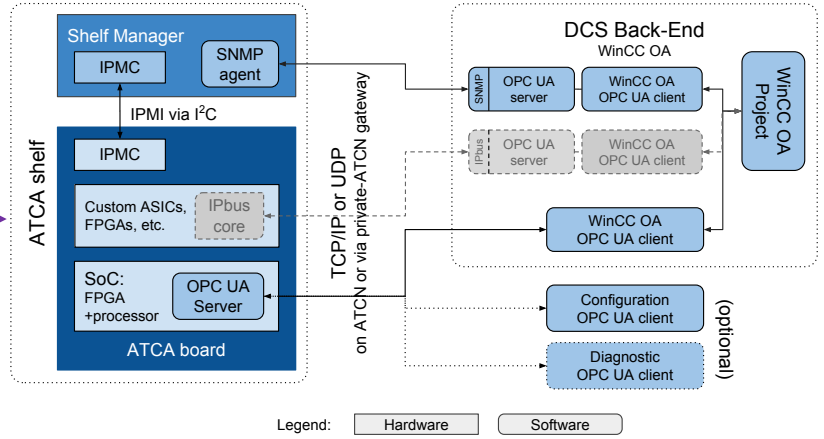

©S. Schlenker

Phase-II TDAQ – SoC survey

Choice of a chip:

- *Usage - Mainly for monitoring & control, with the programmable logic used to interface with the other onboard components*
- *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*

→ *Most systems indicated the **Xilinx Zynq US+ ARM 64bits** as a baseline example*

Mezzanine – SoM (to ensure upgradability):

- *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 doesn't imply re-design of the blade*

Choice of an OS:

- *Initial-preference - Linux-based OS: most mentioned CentOS*
- *Wish-list - ssh-access, sudo-rights*

→ *Full agreement on Linux-based OS - Cent OS main candidate*

Network connection – ATLAS Technical Control Network (ATCN) / Isolation:

- Either direct connection to the ATCN (ATLAS Technical Control Network) or Isolation

→ *The choice of either will be taken by subsystems based on use cases, while respecting external restrictions (sysAdmin,TC,etc…)*

TDAQ SoC user requirements document

- Survey and discussion conclusions were compiled into a '**SoC Requirements document**' for **TDAQ** subsystems – **APPROVED** Mar20

Main issues addressed:

- **Hardware**
  - Recommendations:
    - Common **vendor** - *for common tooling and support*
    - **Ethernet** connectivity - *1 + 1 spare*
  - Requirements:
    - Assigned **MAC** address - *either via CERN or vendor*
    - Minimal **memory** - *to ensure ability to run the chosen common OS*
    - **SoM** - *to ensure upgradability*
    - **HW/SW compatible with network isolation** - *For non upgradable SoCs*

Ensuring SoC & related HW dependencies will allow proper operation of the device and it's interface with the infrastructure

- **Functional requirements**
  - *non critical* HW components (**following ATCA compliance**)

Ensuring the SoC implementation on ATCA blades will not interfere with the shelf-management https://cdsweb.cern.ch/record/1159877?ln=en

- **System integration and software** requirements & recommendations
  - **ATCN** compliance
  - **Common OS**
  - **External services** requirements: *DCS, RunCtrl, etc...*

Ensuring the SoC related software dependencies will allow proper operation

- The TDAQ document served as baseline for the final **ATLAS SoC requirements document** – **APPROVED** Jul20

Remark 2.1: *TDAQ SoC user requirements document*

Being the largest consumer of SoC devices in ATLAS, TDAQ has compiled a SoC user requirements document that is serving as baseline for the ATLAS document. As such, some of the more detailed requirements & recommendations have been generalized in order to reflect the overall usage in ATLAS. Further detailed can be found in [5].

ATLAS SoC user Requirements document

ATLAS Phase-II Upgrade Project

SoC for ATLAS subsystems: Requirements Document for HL-LHC

Abstract

This document describes the high-level functional and performance requirements of the SoC for all the deliverables of ATLAS Phase-II Upgrade Project (UPR) and their interfaces. While the conceptual design and implementation of the ATLAS detector upgrades is described in the individual sub-detectors TDRs, this document provides the UPR systems and sub-systems with a common framework for requirements capture and the specification of interfaces, allowing detailed designs to proceed concurrently.

ATLAS-SoC Interface Requirements Document		
ATLAS Doc:	ATU-GE-ES-0005	
EDMS Id:	2373932	
EDMS Url:	https://edms.cern.ch/document/2373932	
Version:	1.2	
Created:	May 7, 2020	
Last modified:	July 24, 2020	
Prepared by:	**Checked by:**	**Approved by:**
Revital Kopeliansky Wainer Vandelli	ATLAS Upgrade Steering Committee	ATLAS Upgrade Steering Committee

*** highlights in the next slides ***

ACKNOWLEDGEMENT

We acknowledge the help and contributions of the ATLAS Upgrade community, the ATLAS TDAQ Phase-II community and the ATLAS TDAQ system administration team. This document is the result of fruitful discussions with and valuable input from many members of these communities.

SoC-user requirements – Mezzanine (SoM)

- 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 2.3: *SoC lifetime requirements*

Upgradability of the SoC OS is required for maintaining a secure interface into the network during its operation lifetime. In order to ensure long term functionality, systems including SoC in their hardware should make sure the device is integrated as a SoM.
In case the implementation of a SoM is not technically feasible, the SoC will not be replaceable and therefore may become outdated and unsupported during the ATLAS Phase-II lifetime. In this scenario the device will have to be isolated from the ATLAS control network. While different isolation scenarios are possible, it must be possible to operate the system under the most constraining one.
In this case the system would become accessible exclusively through a gateway machine and it shall be ensured that neither software nor hardware limitations (e.g. missing rack space for additional components) would prevent this. Moreover, the appropriate subsystem 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, Open Platform Communications Unified Architecture (OPC UA) server for DCS, etc.).

SoC-user requirements - Functions

Requirement 2.4: *Monitoring of non-critical components*

All non-critical components (e.g. optical transceivers, LED status registers, extra temperature sensors), as defined by the ATCA compliance [4], shall be monitored through a SoC. The suggested implementation includes a dedicated bus mastered by the SoC to which the components are connected. Equivalent schemes implementing a communication link between the SoC and another devices with control of the monitoring components can be envisaged as long as they are compliant with the ATCA specification [4].

Remark 2.7: *ATCA compliance - critical components*

The critical components according to the ATCA compliance are specified under Section 3.9.3.1 on page 3-197 of PICMG 3.0 R3.0_withErrata001 [4] and are detailed over REQ 3.735-3.743. These include at least one temperature sensor per blade and it is also recommended to incorporate power-supply sensors.

Requirement 2.5: *ATCA blades onboard critical components list*

The baseline design of the different subsystems ATCA blades is rather common in terms of the basic chosen onboard components. The following list of common components specifies the ones that are considered critical according to both the ATCA compliance [4] and the blades safe functionality, and as such should be included in the IPMC bus, rather than on the SoC bus:

- Power-management chips
- DCDC converters
- FPGA power-modules
- At least one temperature-sensor per FPGA
- At least one onboard temperature sensor, near regions of high-power density in the PCB

Remark 2.8: *Critical components monitoring while limited resources*

In case the total of number of critical components to be monitored exceeds the input, output or processing capabilities of the IPMC, a careful decision needs to be made with regards to the frequency of the monitoring, while giving priority to power devices.

SoC-user requirements – OS & connectivity

- 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:
  - *ATLAS-community, CERN IT, CERN Security*

Requirement 2.6: *Direct ATCN connectivity*

Direct connection of computing devices to the ATCN requires compliance with different rules and practices.

- the Operating System must be approved by CERN
- it must be possible to regularly update the system within reasonable time
  - severe security flaws may require prompt patching, compatibly with the operation needs

Requirement 2.7: *The SoC Operating System*

All ATLAS sub-systems implementing SoCs devices (with an aim to be connected directly to the ATCN) should utilize a common OS. The choice of the OS to be deployed on the SoCs falls under the responsibility of the ATLAS Phase-II Upgrade Coordinator, after consulting with the sub-detector Upgrade Project Leaders, TDAQ System Administrators, CERN IT and ATLAS management at large.

Remark 2.9: *Operating System Definition*

A centrally managed, common operating system will include:

- kernel version and common driver set
- system tools and libraries (e.g. init system, compiler, ...)
- user-space libraries and software not distributed with CERN or ATLAS releases

Summary

- **Increase in SoC usage** across ATLAS systems over the phases of design
  - *9 times more than from Run2 (Phase-0)*
- 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***
- Formulation of a **TDAQ SoC-user requirements,** generalized and converted into **an** official **ATLAS SoC-user requirements**
  - **Cross referenced to other** official ATLAS technical documentation
    - *Subsystems design reports*
    - *ATLAS DCS user requirements document*
    - *SoC-DAQ interface* requirements document
- **Continues discussions with CERN IT, LHC experiments and LHC departments** (1st workshop, and on the coming Friday session)

Many thanks again *to:

ID - C. Gemme, O.Kepka

Lar - A. Straessner, T. Hrynova

Tile - F. Carrio Argos, F. Martins, O. Solovyanov

Muon - M. Corradi , P. Fleischmann,, E. Pasqualucci, K. Polawski

Muon/TDAQ - T. Costa De Paiva, M. Ishino, L. Levinson, Y. Okumura, A. Tanaka, R. Vari

TDAQ – A. Annovi, G. Avolio, M. Begel, B. Gorini, E. Hazen, R. Middleton, D. Miller, T. Mkrtchyan, W. Panduro-Vazques, D. Sankey, D. Scannicchio, U. Schafer, W. Vandelli, S. Veneziano, M. Warren,

ESE-TDAQ, - S. Haas, R. Spiwoks

FW – B. Giacobbe, J. Guimaraes de Costa, M. Kocian, K. Korcyl, F. Lasagni, M. Wittgen,

TC - S. Schlenker

UC - F. Lanni

for the info, help and cooperation in preparing this summary ☺

*Apologies in case I forgot someone

ID , LAr , Tile , Muon , Muon-TDAQ , TDAQ , ESE-TDAQ , FW , TC , UC