ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD017

WREN: A VERSATILE WHITE RABBIT EVENT NODE FOR CERN'S TIMING SYSTEM RENOVATION

T. Gingold*, M. Cejp, A. Dujović, E. Gousiou, F. W. Hoguin, F. Irannejad, I. Kozsar, G. Kruk, T. Levens, G. Moscardi, M. Murillo Moya, P. Peronnard, J. Serrano, T. Wlostowski, A. Zeising,
European Organization for Nuclear Research (CERN), Geneva, Switzerland

Abstract

WREN is a versatile White Rabbit (WR) node developed for CERN's event-based timing system renovation. Thousands of WRENs are expected to be deployed across the whole CERN accelerator complex from 2027 onwards. Equipped with dedicated hardware and gateware, WREN integrates synchronisation in both TAI (International Atomic Time) and RF (accelerator Radio Frequency) timing. It can function as a TAI event transmitter and receiver, a Beam Synchronous (RF) transmitter and receiver, and is also capable of FPGA-based time-to-digital conversion and fine-delay generation. WREN is highly adaptable for various timing and trigger distribution systems. It is available in multiple form factors, including PCIe, VME, PXIe, and µTCA. All boards are based on the Zynq UltraScale+ System-on-Chip (SoC), designed using the open-source KiCad tool, and licensed under the CERN Open Hardware License (OHL). The gateware and software are also open source. This paper presents the WREN hardware modules, the gateware architecture, and potential customisations for applications beyond CERN. It also shares insights from the initial pilot deployments at CERN.

INTRODUCTION

International Atomic Time and Accelerator Radio Frequency Timing

CERN's accelerator complex delivers particle beams to a wide range of experiments, each with highly specific requirements in terms of beam properties such as intensity, energy, and size. To meet these demands, beams pass through several accelerators and undergo complex manipulations in the CERN accelerator complex, including particle production, bunching, cooling, steering, acceleration, transfer, and eventual delivery to the experimental points. These processes are orchestrated through precisely timed sequences within each accelerator's operational cycle.

Two timescales are fundamental for these operations and are often used together by different types of equipment: TAI[1] (International Atomic Time, our ordinary everyday time) and RF (accelerator Radio Frequency) timing. TAI provides nanosecond precision when defining Central Timing Events (CTIMs), such as Cycle Start, Beam Injection, Beam Extraction and Cycle End. The RF timing, in contrast, is specific to a given accelerator and derived directly from the radio frequency that acts on the particle beams. It is defined by (a) the Turn clock, or revolution frequency, which is the rate at which particles complete one turn in the accelerator, and (b) the Bunch clock, a unit fraction of the base acceleration frequency. In the Large Hadron Collider (LHC), the Bunch clock (40 MHz) is one tenth of the base acceleration frequency (400 MHz) and represents the maximum frequency at which particle bunches can cross at the experimental points, making it the key reference for synchronising all electronic systems involved in bunch measurements and event detection.

Building on these two timescales, CERN currently operates three distinct timing systems [1] across its accelerator complex:

- **General Machine Timing (GMT):** Based on the TAI timescale, GMT distributes both data and a common notion of time throughout the entire accelerator chain. All accelerator equipment makes use of this system. At the physical layer, GMT relies on RS-485 serial links operating at 500 kbps over copper, with a single transmitter feeding hundreds of receivers.
- **Radio Frequency (RF):** Based on the RF timescale, this system distributes the Turn clock and Bunch clock through optical links with very low jitter, and without any data payload. There are three instances in operation: one for the Super Proton Synchrotron (SPS) and one for each LHC beam. It is used primarily by equipment that depend on knowledge of the arrival time of individual bunches, such as the experiments and injection / extraction kickers, while slower systems (e.g., power converters for dipole magnets) typically do not require it.
- **Beam Synchronous Timing (BST):** Also based on the RF, BST combines a less precise RF reference with GMT data. The GMT data are encoded using the RF clocks into a 160 Mbaud optical stream, allowing both RF clock and data distribution through relatively low-cost electronics. BST is used mainly by beam observation equipment to align their measurements with individual circulating bunches.

White Rabbit Timing (WRT)

The WRT project [1] aims at consolidating these systems (GMT, RF and BST) by using the same architecture and the same infrastructure: a physical White Rabbit network. The creation of such a global, CERN-wide network not only eliminates the need of maintaining the multitude of existing

---

* tristan.gingold@cern.ch

[1] TAI is the basis for UTC, which is used for civil timekeeping all over the Earth's surface and which has leap seconds.

networks, but also opens new possibilities for applications requiring a real-time communication channel and exchange of information across networks, something that is complex today. In addition to a significantly higher throughput (1 Gbps), WR and the WRT network design bring a number of other advantages:

- Automatic link-delay calibration.
- Guaranteed upper bound in message latency.
- Increased availability due to backbone redundancy [2].
- Simplified network and cross-network communication.

The new unified network will comprise nearly 200 WR switches and 1800 nodes, making it the largest WR installation at CERN. To ensure its reliable operation and supervision, a dedicated design and configuration have been implemented, along with advanced diagnostic methodologies and monitoring strategies [2].

WREN

For both GMT and BST, the consolidation relies on the use of a common hardware platform, the WREN board [3], a general-purpose WRT node capable of functioning as a GMT receiver, GMT transmitter, BST receiver, or BST transmitter. Currently, while the WREN can reconstruct RF clocks for BST systems, it does not achieve the precision required for RF systems, as its hardware is limited by cost constraints. In applications demanding higher accuracy, a dedicated WR2RF card will be deployed alongside a WREN.

The WREN offers a layer on top of WR which allows users to send and receive messages, and react to them by a) producing pulses or trains of pulses and/or b) producing interrupts or sequences of interrupts to synchronise software tasks in a host computer, in both TAI and RF time.

The consolidation will cover all accelerators and facilities at CERN. It began in Q2 2025 with pilot installations and will continue during Long Shutdown 3 of the CERN accelerator complex (2026–2028), focusing on the SPS, LHC, and their experimental and test facilities. It will be extended to include the remaining accelerators in Long Shutdown 4 (2034).

WREN HARDWARE

Initially, the use of WR-enabled modular boards, such as SPEC [4] or SVEC [5] paired with FMC-DIO [6], was considered for WREN hardware. However, the high cost of modular boards for the required production volume, along with early SPEC implementations revealing the limitations of soft-core CPUs (e.g., lower CPU frequencies and lack of a DDR controller in soft-core CPUs), highlighted the need for a dedicated solution based on a System on Chip (SoC) with hard CPU cores.

The WREN is therefore implemented as a dedicated hardware design, currently available in VME, PCIe, and PXIe form factors (Fig. 1). While hosted solutions remain the mainstream at CERN, the WREN, with features such as external power support and unused SoC resources, could, if required, be adapted to operate as a stand-alone node in the future.

All WRENs are based on a SoC from the Xilinx Ultrascale+ family and share many identical electronic components. The fundamental design principle is to maximize reusability, only adapting the host-bus bridge and number of I/Os for each form factor.

Regarding the user interface, each WREN provides between 8 and 38 (depending on the form factor) front panel LEMO-00 coaxial connectors that can serve as external inputs or outputs. Some connectors are fixed as inputs (I), others as outputs (O), while a subset is configurable via the host as either input or output (I/O). Inputs can be assigned, through host configuration, as the start, stop, or clock source for a pulser (see following section), whereas outputs can deliver signals generated by a pulser. An on-board micro-coaxial extension connector is also available to interface with the patch panel, which provides 32 outputs in a convenient rack-mounted 1U form factor. Each input includes over/under-voltage protection and selectable termination, while each output includes an enable function and can drive 3 V into 50 Ω loads. All I/Os are single-ended and connected to the High Performance (HP) FPGA bank. Table 1 summarises the main WREN features per form factor.

The BST functionality is implemented entirely in gateware, without the need for dedicated hardware components, a strategic decision aimed at reducing costs during the early stages of the project.

For compatibility with the existing BST system, which exclusively uses the VME form factor, the BST signals (RF clocks and BST Stream) are provided on the WREN in the same way as in the current BST receiver and transmitter [7], via the VME P0 and P2 connectors. In addition, a Rear Transition Module (RTM) in VME form factor has been developed. The BST-RTM connects to the VME P2 connector and mainly provides copper-to-optical conversion and fan-out of the BST Stream. For the PCIe and PXIe form factors, BST functionality is not required but can optionally be enabled through on-board Ultra-Fine Coaxial (UFC) connectors.

WREN will also be available in the µTCA form factor. However, given the low deployment volume at CERN, a dedicated board will not be developed; instead, a generic µTCA carrier, AFCZ [8], combined with a digital I/O FMC mezzanine will be used.

KiCad and Open Hardware

The WREN VME, PCIe, and PXIe boards as well as the patch panel and BST-RTM were designed using the free and open-source tool KiCad [9], demonstrating its capability to handle complex 12-layer designs integrating an 800-pin SoC, multi-Gb lines and over 1500 components.

All the designs are also licensed under the weakly reciprocal variant (CERN-OHL-W) of the CERN Open Hardware License [10], and are available in the Open Hardware Repository [11].

ISBN: 978-3-95450-255-4 ISSN: 2226-0358 doi: 10.18429/JACoW-ICALEPCS2025-WEPD017

Table 1: Main Hardware Features per Form Factor

WREN Type	**WREN-V (VME)**	**WREN-Ie (PCIe)**	**WREN-Xie (PXIe)**
Host Interface	VME64 and VME32/16	PCIe Gen2 ×4	PCIe Gen2 ×4
Front panel I/Os	Modular, with 2 options: • Single: 6× I, 8×IO • Double: 6× I, 8×IO, 24×O	2× I, 6×IO	2× I, 6×IO
Backplane I/Os	32×IO on P2	N/A	• 10/100 MHz distribution • DSTARA/B/C as on typical NI Timing board
Patch panel I/Os	µ-coax with 32×O	µ-coax with 32×O	µ-coax with 32×O
BST	P0/P2	UFC (optional)	UFC (optional)

Figure 1: WREN-Ie and WREN-V hardware.

WREN GATEWARE/FIRMWARE

The implementation leverages both gateware on the FPGA (HDL) and firmware (embedded software) on the hard CPUs of the SoC. Time-critical functions, such as RF reconstruction and precise trigger generation, are implemented in gateware, while less time-critical tasks, including frame parsing, configuration, and diagnostics, are executed in bare-metal code on the Arm Cortex-R5F. The Arm Cortex-A53 cores of the SoC are currently unused to allow flexibility for future developments. A stand-alone WREN in the future could employ these cores to run an operating system, eliminating the need for a host.

As a GMT transmitter, the WREN primarily receives data from the host in the form of events [1], each with its TAI due timestamp in the future, which it encapsulates using a custom protocol into WR frames for broadcast. Frame transmission can also be triggered by an external input on a LEMO connector. Currently, an event's due time is defined in the TAI timescale, but this could be extended to support events defined in the RF timescale (see Section "Beyond CERN"). The transmitter is mostly implemented in firmware.

The architecture of the WREN as a GMT receiver is more complex. As shown in Figure 2, the WR PTP Core (WRPC) handles the WR protocol frames and provides TAI synchronisation. A filter separates RF-type traffic, which is processed in gateware by the BST core, from other traffic, which is handled by the Network Interface Controller (NIC).

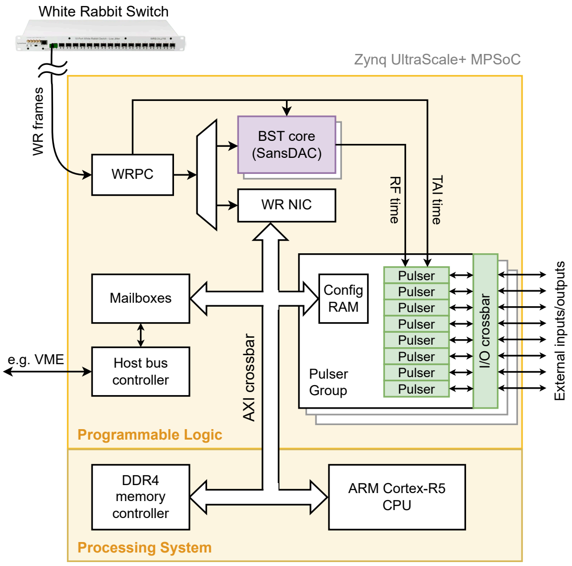

Figure 2: Architecture of the WREN as GMT receiver.

Pulser

The pulser is an elaborate counter and a core component of the WREN receiver. Developed in gateware it is essential for executing the main WREN functions: digital pulse and pulse-train generation, as well as host system interrupts or sequences of interrupts. In all form factors, the WREN provides 32 pulsers, organized into four groups of eight. The digital output of a pulser can be routed to an external connector, the interrupt controller or internally to another pulser unit in the same group. Additionally, outputs can be combined through logic gates (AND, NAND, OR, NOR) before being sent to a physical connector. In this way, chains and combinations of pulsers can be created, providing users with a wide range of possibilities for pulse and interrupt generation.

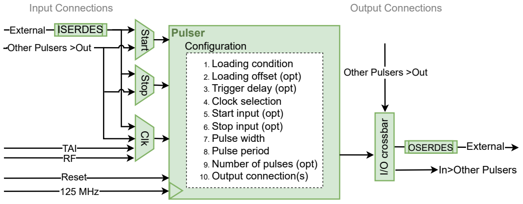

Figure 3: Pulser configuration and connections.

The behaviour of a pulser is determined by its configuration (Fig. 3), provided by the host, in combination with external inputs or chained pulser outputs. The pulser output is generated:

- **If** the pulser conditions are satisfied by the parameters of an incoming event frame. A comparator, implemented in firmware, evaluates the received frame data against the pre-configured conditions of each pulser.
- **When** the matching event's due timestamp is reached, with optional offsets defined in the pulser configuration.

As mentioned earlier, although events are defined only in the TAI timescale, the ability to clock the pulser with RF clocks makes it possible to generate RF-synchronous outputs, for example, a pulse after beam injection, three turns later, at the fourth bunch.

Pulser Clocking

Because of the large number of clocks involved (RF, TAI, external inputs, and outputs from other pulsers), it was not feasible to implement multiple clock domains for the pulsers. Therefore, from the start of the design, the decision was made to use a single 125 MHz clock with a sub-period value. On the outputs, OSERDES modules (high speed parallel-to-serial converters) are used to convert the period and sub-period values into 1 ns granularity pulses, while on the inputs, ISERDES modules (high speed serial-to-parallel converters) are used to translate signals into period and sub-period values.

This 1 ns granularity represents a compromise between requirements and hardware capabilities. It allows us to meet the precision needs for generating pulses based on TAI and RF clocks. However, it introduces too much jitter for generating the BST signals, which is why transceivers are used in that case (see next section).

This solution is also future-proof for integrating additional IP cores. The FPGA currently uses only about 50% of its available resources, leaving ample capacity for additional IP cores. By operating entirely within a single clock domain, the pulser IP is already optimized for seamless integration with other cores. In future designs, next-generation timing functionality could be implemented directly as an IP core within applications.

Time-to-Digital Converter and Fine Delay

In addition to the timely execution of actions (pulses and interrupts), a key feature embedded in the WREN is its ability to precisely time-stamp input signals via its Time-to-Digital Converter (TDC) functionality. Leveraging FPGA transceivers, it achieves a resolution of 1 ns. The pulser also supports fine-delay adjustments with the same resolution.

For applications requiring such precision, the WREN could potentially replace WR Trigger Distribution (WRTD) [12], with only minor adaptations. At present, WREN is configured to operate either as a transmitter or as a receiver, but not concurrently. Trigger distribution would necessitate simultaneous support for both roles, a feature that is technically feasible though not yet implemented.

Diagnostics

The WREN records all incoming frames as well as the timestamps of input and output pulses. Diagnostic data is stored in a 2-gigabyte DDR4 memory. The gateware logs input and output timestamps, while the firmware logs the frames. This data can be gradually retrieved by the host.

BEAM SYNCHRONOUS TIMING (BST)

WREN as BST Receiver: RF Reconstruction

At CERN, the RF frequency excursion during acceleration is generally too large to be reliably used as an encoding clock for an FPGA transceiver. The Phase-Locked Loop (PLL) inside the transceiver would lose lock during the frequency ramp. This limitation motivated the development of a WR-based solution for frequency distribution. RF transfers over WR have already been demonstrated with the SPS Low Level RF upgrade in 2021 [13]. The nodes in that system are WR2RF boards, which provide high precision clocking but are costly. In contrast, the WREN solution is an FPGA-only implementation without a Digital-to-Analog Converter (DAC), hence the corresponding HDL code is named *Sans-DAC* (i.e. "without DAC" in French). The SansDAC core (Fig. 4) leverages FPGA resources to directly transmit digital output samples through the GTH transceiver, using raw encoding. This generates a digital pattern corresponding to the Bunch clock signal. The transceiver output is then physically looped back into the FPGA, where the recovered RF clock is made available to the rest of the WREN core. The Bunch and derived Turn clock are output on the VME backplane P0 and P2 connectors, as well as on optional on-board UFC connectors for the PCIe/ PXIe form factors. Additionally, they can be used as start, stop, or clock signals in the pulsers. The WREN can generate up to four RF clocks from two different Frequency Tuning Word (FTW) sources.

Frequency Tuning Words are broadcast into the WRT network as Radio over Ethernet (RoE) frames [14] and are processed by WREN's SansDAC module, as shown in Figure 4. The Numerically Controlled Oscillator (NCO), identical to the one used in the WR2RF, accumulates phase at each WR clock cycle to recover the Bunch clock and derive the corresponding Turn clock. RF time is thus defined by the deterministic accumulation of FTWs within the NCO. Each RF epoch corresponds to the TAI reference at the moment of an NCO reset, when phase accumulation restarts, ensuring

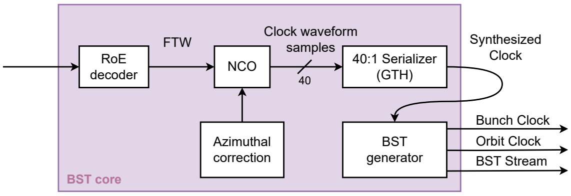

Figure 4: Architecture of the BST core (SansDAC).

alignment between the two timescales. Another important aspect is that converting from RF-synchronous to beam-synchronous operation requires accounting for:

- Network cable delay: deterministic when using RoE, similar to WR-streamers [15].
- Azimuthal position of the observing node relative to the master RF clock source: handled in the SansDAC.

This approach delivers better performance compared to the current BST receiver. For bunch clock reconstruction, the jitter of the WREN (referenced to the WR2RF) is measured at 50 ps RMS, half the jitter observed with the current BST receiver.

WREN as BST Transmitter

Apart from reconstructing RF from FTWs, a WREN can also generate the legacy BST Stream, effectively acting as a BST transmitter. To achieve this, the WREN combines the recovered Turn / Bunch clock with GMT data following the TTC [16] protocol. The TTC protocol defines the use of a unidirectional optical fiber link in which two information channels, A and B, are Time Division Multiplexed (TDM) and Bi-Phase Mark (BPM) encoded using the Bunch clock as the carrier frequency. Channel A exclusively carries the Turn clock, while Channel B transmits packaged address and data information for various reset commands, calibration, control, and test parameters, along with a dedicated field containing redundant bits for error detection and correction. In the gateware implementation, this coding scheme is realized as a Hamming code capable of double-error detection and single-bit error correction. Since the Bunch clock must be transmitted alongside the data, the Channel B information is jointly encoded with the Bunch clock using the BPM encoding scheme.

BEYOND CERN

Within the White Rabbit Collaboration [17], a new project [18] aims to extend WREN towards a general-purpose timing node for any facility requiring precise synchronisation.

A key limitation of the current WREN implementation is the lack of support for precise RF-synchronous timing, which is essential for light sources. The system must therefore provide equivalent functionality in both UTC and RF-synchronous modes. For example, it should allow the generation of events with RF timestamps and the loading of a pulser at an RF-synchronous instant. Additionally, optional mounting of a DAC or PLL should be supported to achieve performance comparable to WR2RF when required.

Regarding form factors, WREN currently exists in PCIe, PXIe, and VME formats, with the WREN functionality to be implemented on an AFCZ carrier with a digital I/O FMC for µTCA . Nevertheless, a dedicated µTCA-compliant WREN could offer a simpler, cheaper, and more maintainable solution, and its development is proposed within the project scope.

Finally, to improve accessibility, a Python API will complement existing Linux drivers and C/C++ libraries. This will allow users to configure and embed WREN functionality in their applications without requiring in-depth knowledge of White Rabbit, lowering the entry barrier.

CONCLUSION

The WREN is a versatile WR node developed primarily for upgrading CERN's timing systems with the following capabilities:

- Synchronisation to a TAI-traceable Grandmaster.
- Reproduction of RF clocks.
- Generation of BST streams.
- Event processing which, combined with local configuration, enables actions such as pulse generation and host interrupts, executed in absolute TAI or RF time.
- Event transmission within the WRT network.
- Input timestamping and output generation with fine-delay control.

This functionality also makes it suitable for a wide range of other applications currently being explored through the WR Collaboration.

REFERENCES

[1] G. Moscardi *et al.*, "White Rabbit Timing: The New CERN Accelerator Timing System", presented at ICALEPCS'25, Chicago, IL, USA, Sep. 2025, paper WEBR003, this conference.

[2] M. Suminski, I. Kozsar, M. M. Lipinski, and J. Palluel, "Evolution of White Rabbit Network for large-scale deployment at CERN", presented at ICALEPCS'25, Chicago, IL, USA, Sep. 2025, paper TUPD070, this conference.

[3] White Rabbit Event Node (WREN), `https://ohwr.org/projects/wren`

[4] Simple PCIe FMC carrier SPEC, `https://ohwr.org/projects/spec`

[5] Simple VME FMC Carrier (SVEC), `https://ohwr.org/projects/svec`

[6] FMC DIO 5 ch TTL A, `https://ohwr.org/projects/fmc-dio-5chttla`

[7] Ideas for BST over White Rabbit, `https://indico.cern.ch/event/1249710/contributions/5251693/attachments/2603199/4495528/BI-TB_BSToWR.pdf`

[8] AMC FMC Carrier AFCZ, `https://ohwr.org/projects/afcz/`

[9] KiCad — Cross-Platform Open-Source EDA Suite, `https://www.kicad.org/`

[10] CERN Open Hardware Licence, `https://cern-ohl.web.cern.ch`

[11] Open Hardware Repository (OHWR), `https://ohwr.org/`

[12] White Rabbit Trigger Distribution (WRTD), `https://ohwr.org/projects/wrtd/`

[13] T. Włostowski *et al.*, "White Rabbit and MTCA.4 Use in the LLRF Upgrade for CERN's SPS", in *Proc. ICALEPCS'21*, Shanghai, China, Oct. 2021, pp. 847–852. `doi:10.18429/JACoW-ICALEPCS2021-THBR02`

[14] Radio Frequency over Ethernet (RoE), `https://roe-mapping.web.cern.ch`

[15] White Rabbit Streamers, `https://gitlab.com/ohwr/project/wr-cores/-/wikis/WR-Streamers`

[16] P. Gällnö, "Timing, trigger and control distribution and dead-time control in ATLAS", in *6th Workshop on Electronics for LHC Experiments*, Krakow, Poland, Sep. 2000, pp. 415–418. `doi:10.5170/CERN-2000-010.415`

[17] The White Rabbit Collaboration, `https://www.white-rabbit.tech/`

[18] White Rabbit Phase and Synchronization Overview, `https://cernbox.cern.ch/pdf-viewer/public/RuBTbQbGjFYNMYt/WRT_PSO.pdf`