



This is a repository copy of *The triggerless data acquisition system of the XENONnT experiment*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/226481/>

Version: Published Version

Article:

Aprile, E., Aalbers, J., Abe, K. et al. (162 more authors) (2023) The triggerless data acquisition system of the XENONnT experiment. *Journal of Instrumentation*, 18 (07). P07054. ISSN 1748-0221

<https://doi.org/10.1088/1748-0221/18/07/p07054>

Reuse

This article is distributed under the terms of the Creative Commons Attribution (CC BY) licence. This licence allows you to distribute, remix, tweak, and build upon the work, even commercially, as long as you credit the authors for the original work. More information and the full terms of the licence here:

<https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

PAPER • OPEN ACCESS

The triggerless data acquisition system of the XENONnT experiment

To cite this article: E. Aprile *et al* 2023 *JINST* **18** P07054

View the [article online](#) for updates and enhancements.

You may also like

- [Fast \$b\$ -tagging at the high-level trigger of the ATLAS experiment in LHC Run 3](#)
G. Aad, B. Abbott, K. Abeling *et al.*
- [Proton reconstruction with the TOTEM Roman pot detectors for high- \$\sqrt{s}\$ LHC data](#)
A. Hayrapetyan, A. Tumasyan, W. Adam *et al.*
- [Development of the CMS detector for the CERN LHC Run 3](#)
A. Hayrapetyan, A. Tumasyan, W. Adam *et al.*



ECS The Electrochemical Society
Advancing solid state & electrochemical science & technology

247th ECS Meeting
Montréal, Canada
May 18-22, 2025
Palais des Congrès de Montréal

ECS UNITED

Unite with the ECS Community

**Register to
save \$\$
before
May 17**

The triggerless data acquisition system of the XENONnT experiment

The XENON collaboration

E. Aprile,^a J. Aalbers,^{b,c,2} K. Abe,^d F. Agostini,^e S. Ahmed Maouloud,^f L. Althueser,^g B. Andrieu,^f E. Angelino,^h J.R. Angevaere,^{b,1} V.C. Antochi,^c D. Antón Martín,ⁱ F. Arneodo,^j L. Baudis,^k A.L. Baxter,^l L. Bellagamba,^e R. Biondi,^m A. Bismark,^k E.J. Brookes,^b A. Brown,ⁿ S. Bruenner,^b G. Bruno,^o R. Budnik,^p T.K. Bui,^d C. Cai,^q J.M.R. Cardoso,^r D. Cichon,^s A.P. Cimental Chavez,^k D. Coderre,ⁿ A.P. Colijn,^b J. Conrad,^c J.J. Cuenca-García,^{k,t} J.P. Cussonneau,^{o,3} V. D'Andrea,^{u,m} M.P. Decowski,^b P. Di Gangi,^e S. Di Pede,^b S. Diglio,^o K. Eitel,^t A. Elykov,^{n,t,1} S. Farrell,^v A.D. Ferella,^{u,m} C. Ferrari,^m H. Fischer,ⁿ M. Flierman,^b W. Fulgione,^{h,m} C. Fuselli,^b P. Gaemers,^b R. Gaior,^f A. Gallo Rosso,^c M. Galloway,^k F. Gao,^q R. Glade-Beucke,ⁿ L. Grandi,ⁱ J. Grigat,ⁿ M. Guida,^s R. Hamann,^s A. Higuera,^v C. Hils,^w L. Hoetzsch,^s N.F. Hood,^x J. Howlett,^a M. Iacovacci,^y Y. Itow,^z J. Jakob,^g F. Joerg,^s A. Joy,^c N. Kato,^d M. Kara,^t P. Kavrigin,^p S. Kazama,^z M. Kobayashi,^z G. Koltman,^p A. Kopec,^x F. Kuger,ⁿ H. Landsman,^p R.F. Lang,^l L. Levinson,^p I. Li,^v S. Li,^l S. Liang,^v S. Lindemann,ⁿ M. Lindner,^s K. Liu,^q J. Loizeau,^o F. Lombardi,^w J. Long,ⁱ J.A.M. Lopes,^{r,4} Y. Ma,^x C. Macolino,^{u,m} J. Mahlstedt,^c A. Mancuso,^e L. Manenti,^j F. Marignetti,^y T. Marrodán Undagoitia,^s K. Martens,^d J. Masbou,^o D. Masson,^{n,1} E. Masson,^f S. Mastroianni,^{y,1} M. Messina,^m K. Miuchi,^{aa} K. Mizukoshi,^{aa} A. Molinaro,^h S. Moriyama,^d K. Morå,^a Y. Mosbacher,^p M. Murra,^a J. Müller,ⁿ K. Ni,^x U. Oberlack,^w B. Paetsch,^p J. Palacio,^s R. Peres,^k C. Peters,^v J. Pienaar,ⁱ M. Pierre,^o V. Pizzella,^s G. Plante,^a J. Qi,^x J. Qin,^l D. Ramírez García,^k A. Rocchetti,ⁿ L. Sanchez,^v P. Sanchez-Lucas,^k J.M.F. dos Santos,^r I. Sarnoff,^j G. Sartorelli,^e J. Schreiner,^s D. Schulte,^g P. Schulte,^g H. Schulze Eißing,^g M. Schumann,ⁿ L. Scotto Lavina,^f M. Selvi,^e F. Semeria,^e P. Shagin,^w S. Shi,^a E. Shockley,^x M. Silva,^r H. Simgen,^s A. Takeda,^d P.-L. Tan,^c A. Terliuk,^{s,5} D. Thers,^o F. Toschi,^{n,t} G. Trincherro,^h C. Tunnell,^v F. Tönnies,ⁿ K. Valerius,^t G. Volta,^k C. Weinheimer,^g M. Weiss,^p D. Wenz,^w

¹Corresponding author.

²Now at: SLAC National Accelerator Laboratory, Menlo Park, CA 94025-7015, U.S.A. and Kavli Institute for Particle Astrophysics and Cosmology, Stanford University, Stanford, CA 94305-4085 U.S.A..

³Deceased.

⁴Also at: Coimbra Polytechnic — ISEC, 3030-199 Coimbra, Portugal.

⁵Also at: Physikalisches Institut, Universität Heidelberg, Heidelberg, Germany.



**C. Wittweg,^k T. Wolf,^s D. Xu,^a Z. Xu,^a M. Yamashita,^d L. Yang,^x J. Ye,^a L. Yuan,ⁱ
G. Zavattini,^{e,6} S. Zerbo,^a M. Zhong^x and T. Zhu^a**

^a*Physics Department, Columbia University, New York, NY 10027, U.S.A.*

^b*Nikhef and the University of Amsterdam, Science Park, 1098XG Amsterdam, Netherlands*

^c*Oskar Klein Centre, Department of Physics, Stockholm University, AlbaNova, Stockholm SE-10691, Sweden*

^d*Kamioka Observatory, Institute for Cosmic Ray Research, and Kavli Institute for the Physics and Mathematics of the Universe (WPI), University of Tokyo, Higashi-Mozumi, Kamioka, Hida, Gifu 506-1205, Japan*

^e*Department of Physics and Astronomy, University of Bologna and INFN-Bologna, 40126 Bologna, Italy*

^f*LPNHE, Sorbonne Université, CNRS/IN2P3, 75005 Paris, France*

^g*Institut für Kernphysik, Westfälische Wilhelms-Universität Münster, 48149 Münster, Germany*

^h*INAF-Astrophysical Observatory of Torino, Department of Physics, University of Torino and INFN-Torino, 10125 Torino, Italy*

ⁱ*Department of Physics & Kavli Institute for Cosmological Physics, University of Chicago, Chicago, IL 60637, U.S.A.*

^j*New York University Abu Dhabi - Center for Astro, Particle and Planetary Physics, Abu Dhabi, United Arab Emirates*

^k*Physik-Institut, University of Zürich, 8057 Zürich, Switzerland*

^l*Department of Physics and Astronomy, Purdue University, West Lafayette, IN 47907, U.S.A.*

^m*INFN-Laboratori Nazionali del Gran Sasso and Gran Sasso Science Institute, 67100 L'Aquila, Italy*

ⁿ*Physikalisches Institut, Universität Freiburg, 79104 Freiburg, Germany*

^o*SUBATECH, IMT Atlantique, CNRS/IN2P3, Université de Nantes, Nantes 44307, France*

^p*Department of Particle Physics and Astrophysics, Weizmann Institute of Science, Rehovot 7610001, Israel*

^q*Department of Physics & Center for High Energy Physics, Tsinghua University, Beijing 100084, China*

^r*LIBPhys, Department of Physics, University of Coimbra, 3004-516 Coimbra, Portugal*

^s*Max-Planck-Institut für Kernphysik, 69117 Heidelberg, Germany*

^t*Institute for Astroparticle Physics, Karlsruhe Institute of Technology, 76021 Karlsruhe, Germany*

^u*Department of Physics and Chemistry, University of L'Aquila, 67100 L'Aquila, Italy*

^v*Department of Physics and Astronomy, Rice University, Houston, TX 77005, U.S.A.*

^w*Institut für Physik & Exzellenzcluster PRISMA⁺, Johannes Gutenberg-Universität Mainz, 55099 Mainz, Germany*

^x*Department of Physics, University of California San Diego, La Jolla, CA 92093, U.S.A.*

^y*Department of Physics "Ettore Pancini", University of Napoli and INFN-Napoli, 80126 Napoli, Italy*

^z*Kobayashi-Maskawa Institute for the Origin of Particles and the Universe, and Institute for Space-Earth Environmental Research, Nagoya University, Furo-cho, Chikusa-ku, Nagoya, Aichi 464-8602, Japan*

^{aa}*Department of Physics, Kobe University, Kobe, Hyogo 657-8501, Japan*

E-mail: xenon@lngs.infn.it

ABSTRACT: The XENONnT detector uses the latest and largest liquid xenon-based time projection chamber (TPC) operated by the XENON Collaboration, aimed at detecting Weakly Interacting Massive Particles and conducting other rare event searches. The XENONnT data acquisition (DAQ) system constitutes an upgraded and expanded version of the XENON1T DAQ system. For its

⁶Also at: INFN, Sez. di Ferrara and Dip. di Fisica e Scienze della Terra, Università di Ferrara, via G. Saragat 1, Edificio C, I-44122 Ferrara (FE), Italy.

operation, it relies predominantly on commercially available hardware accompanied by open-source and custom-developed software. The three constituent subsystems of the XENONnT detector, the TPC (main detector), muon veto, and the newly introduced neutron veto, are integrated into a single DAQ, and can be operated both independently and as a unified system. In total, the DAQ digitizes the signals of 698 photomultiplier tubes (PMTs), of which 253 from the top PMT array of the TPC are digitized twice, at $\times 10$ and $\times 0.5$ gain. The DAQ for the most part is a triggerless system, reading out and storing every signal that exceeds the digitization thresholds. Custom-developed software is used to process the acquired data, making it available within ~ 30 s for live data quality monitoring and online analyses. The entire system with all the three subsystems was successfully commissioned and has been operating continuously, comfortably withstanding readout rates that exceed ~ 500 MB/s during calibration. Livetime during normal operation exceeds 99% and is $\sim 90\%$ during most high-rate calibrations. The combined DAQ system has collected more than 2 PB of both calibration and science data during the commissioning of XENONnT and the first science run.

KEYWORDS: Dark Matter detectors (WIMPs, axions, etc.); Data acquisition concepts; Data processing methods; Front-end electronics for detector readout

ARXIV EPRINT: [2212.11032](https://arxiv.org/abs/2212.11032)

Contents

1	Introduction	1
2	From XENON1T to XENONnT	1
2.1	General DAQ upgrades	2
2.2	Triggerless data streams	2
2.3	Fast data processing & immediate data availability	4
2.4	Neutron veto DAQ	5
2.5	Three integrated DAQ subsystems	5
3	Data acquisition & readout	7
3.1	Analog electronics	7
3.2	Digital electronics	8
3.2.1	TPC	8
3.2.2	Muon veto	9
3.2.3	Neutron veto	9
3.3	Busy & high-energy veto	9
3.4	Acquisition monitors	11
3.5	Servers & software	11
4	Live processing	12
4.1	Data stream versus discrete events	12
4.2	Strax(en) data format	13
4.3	Online processing	14
4.4	Online monitoring	14
4.5	Data storage infrastructure	16
5	System control & oversight	16
5.1	Databases	16
5.2	User interface website	17
5.3	Readout coordination software	18
6	Performance	19
6.1	Noise levels	19
6.2	Livetime	19
6.3	Time synchronization	21
6.4	Live processing performance	21
7	Conclusion	24

1 Introduction

A variety of experiments use time projection chambers (TPCs) filled with liquid noble elements (usually xenon or argon) in the search for Weakly Interacting Massive Particle (WIMP) dark matter and rare radioactive decays [1–5]. While the details of each detector differ, common design features include arrays of photosensors at the ends of the drift region and accompanying readout systems.

Interactions in the TPC are observed via two processes: scintillation and ionization. When a particle interacts with either the electrons or nucleus of a target atom, prompt scintillation light and ionization electrons are produced, resulting in two signals. Two arrays of photosensors, photomultiplier tubes (PMTs) in XENONnT, are located above and below the cylindrical drift region of the TPC to capture these signals. The detected prompt scintillation light is referred to as the “S1” signal, while the electrons are drifted under an external electric field towards the liquid-gas interface. When the electrons reach this interface, a much stronger electric field extracts them from the liquid and causes electroluminescence in the gas, producing additional proportional scintillation that is detected and referred to as the “S2” signal. The time between the S1 and S2 signals, which is the drift time of electrons in the liquid phase, as well as the pattern of illumination on the top PMT array caused by the S2, are used to reconstruct the interaction vertex in the detector. The S2 is typically much larger than the S1, and the relative sizes of these two signals are used to discriminate between electronic recoil (ER) and nuclear recoil (NR) interactions.

The XENON collaboration has operated a series of increasingly larger dual-phase xenon TPCs at the INFN Laboratori Nazionali del Gran Sasso (LNGS) in central Italy for nearly two decades, probing WIMP-nucleon cross-sections down to $4.1 \times 10^{-47} \text{ cm}^2$ (for a $30 \text{ GeV}/c^2$ WIMP) [6]. The latest is the TPC of the XENONnT detector, containing 5.9 t in its active target mass and is expected to be sensitive to spin-independent WIMP-nucleon cross-sections down to $1.4 \times 10^{-48} \text{ cm}^2$ (for a $50 \text{ GeV}/c^2$ WIMP) [7].

2 From XENON1T to XENONnT

The upgrade from XENON1T to XENONnT saw the TPC increase in size from $\sim 1 \text{ m}$ in diameter and length to $\sim 1.3 \text{ m}$ and $\sim 1.5 \text{ m}$, respectively, to accommodate a larger target mass. This increase in TPC size was accompanied by a corresponding increase in the number of PMTs to 494, 253 in the top array in the gas phase and 241 in the bottom array in the liquid below the target. This constitutes a two-fold increase from XENON1T where the TPC was instrumented with 248 PMTs.

As detectors continue to grow in size, the maximum drift time (the drift length) of electrons in the TPC grows in accordance. A lower drift field of 23 V/cm in the first science run (SR0) of XENONnT [8] compared to 81 V/cm – 120 V/cm in XENON1T [6] means a further increase in the drift length. The need to store and read out one continuous drift length of data, which can exceed 2 ms , is alleviated by the firmware used by most of the readout hardware. This digital pulse processing with dynamic acquisition windows (DPP-DAW) firmware [9] was developed in collaboration with CAEN [10] for XENON1T [11] (first released in v4.17_137.7), and an updated version (v4.25_137.9) was used for XENONnT [12]. The DPP-DAW firmware affords many useful techniques such as baseline suppression or zero length encoding (ZLE), dynamically-sized acquisition windows that automatically extend as long as the input is above the digitization threshold, and the

independent and continuous readout of each channel. The digitization thresholds are set relative to dynamically-calculated baselines. However, increased drift time leads to an increased temporal width of S2 signals as the freed electrons diffuse over a larger amount of time. This increased temporal width directly increases data rates as signals remain above threshold for longer. Thus, new challenges arise as drift times increase and S2s become longer.

A new active neutron veto sub-detector was built to suppress the NR background from radiogenic neutrons generated through spontaneous fission and alpha-neutron reactions, as these mimic WIMP-induced signatures. It is made of an octagonal structure (3 m-high and 4 m-wide) placed inside the water tank around the cryostat that houses the TPC and is optically separated from the existing muon veto [7]. To improve the neutron detection efficiency, the water will be loaded with gadolinium (Gd). A total of 120 Hamamatsu 8” high quantum efficiency PMTs with low-radioactivity windows are placed along the lateral walls. Neutrons that leave the TPC volume are moderated by the water around the cryostat before being captured by Gd or H. A gamma-ray cascade with total energy of about 8 MeV is generated for capture on Gd and a single 2.2 MeV gamma is emitted in the case of capture on H. The gammas in water, mainly through Compton scattering, are converted into electrons and ultimately into Cherenkov photons. Monte Carlo studies indicate that the neutron veto is expected to reduce the total NR background by a factor of six [7]. Dedicated hardware was implemented to manage the neutron veto data readout. The water Cherenkov muon veto surrounding the cryostat, instrumented with 84 PMTs, is otherwise largely unchanged from XENON1T [11, 13].

2.1 General DAQ upgrades

The XENONnT data acquisition (DAQ) system is an evolution of that which was successfully used for XENON1T [11]. Many aspects of the system have received modifications and improvements based on the XENON1T system and the experience of operating it. An overview of the new system design is shown in figure 1.

One challenge of the general DAQ design is the range of sizes and shapes of signals the system must handle. The muon veto and neutron veto are Cherenkov detectors, registering photon signals over an interval of at most $\sim 1 \mu\text{s}$. The TPC, in contrast, must record both S1 and S2 signals. S1s can be very small, down to a single photon, and are very fast, lasting up to $\sim 100 \text{ ns}$. S2s are much larger, potentially millions of photons, and can have temporal widths exceeding $100 \mu\text{s}$. Representative signals in all three subsystems are shown in figure 2.

While XENON1T used the Phillips Scientific 776 amplifiers with dual $\times 10$ -gain outputs, XENONnT uses custom dual-gain ($\times 10$ and $\times 0.5$) amplifiers developed at the University of Zurich [14]. The low-gain signals from the top PMT array are digitized separately from the high-gain signals to try to improve energy and position reconstruction involving large signals that otherwise saturate the input stage of the digitizers. The low-gain signals from the bottom array are summed together and used by the high-energy veto, discussed in section 3.3. This results in the XENONnT TPC effectively having three times the number of PMT readout channels of the XENON1T TPC (747 compared to 248), as the number of PMTs is doubled and half are read out twice.

2.2 Triggerless data streams

The XENON1T TPC utilized a triggerless readout and a central buffer built from a MongoDB database [15]. Triggering software ran live over this database and determined “events” which were




Figure 1. The XENONnT DAQ layout at several stages. The TPC PMT signals are amplified before being digitized by CAEN V1724 modules. The neutron veto (NV) and muon veto (MV) PMT signals are digitized by the associated V1730S and V1724 digitizers, with a sampling rate of 500 MHz and 100 MHz, respectively. The top array is digitized twice at $\times 10$ and $\times 0.5$ gain, the signals from the latter constitute the high-energy system. The $\times 0.5$ gain signals of the bottom array are fed into the sum-signal fan cascade (Σ -fan cascade). The busy (V1495) and acquisition monitor (V1724) modules handle the busy logic and monitor the system performance. The reader servers read out the digitizers and write the data to a common (Ceph) storage disk. The eventbuilder servers do the processing of the reader data which get written to local storage where it can be distributed to other storage sites. Two MongoDB databases (section 5.1) are used for bookkeeping and inter-process communication, store their data long and short term, and are called the Runs and DAQ database, respectively. A portion of the processed data is also written to the Runs database where it can be accessed for online data monitoring. The control of the DAQ is achieved through the website that communicates commands (run control) to the readers/eventbuilders.

written to disk for later analysis, the remainder of the data was deleted. While this paradigm was successful at realizing a very low effective trigger threshold, the estimation of some backgrounds was complicated by the forced truncation of events after a certain maximum duration. Despite the very low effective trigger threshold for the WIMP search, several other searches were complicated or less sensitive by the enforced event trigger definition [16–18]. Additionally, the buffer database would not scale effectively to match the increased load foreseen by the demands of a larger system. The solution is to forego the software trigger, saving all the data and leaving the determination of events to much later in the data processing pipeline. This novel flexible triggerless framework was implemented in the XENONnT DAQ, facilitating a wide range of scientific searches, calibration options and detailed background studies. A versatile DAQ that omits any fixed event triggering logic

enables the science channels of liquid xenon TPC's to continue to diversify [19, 20], allowing each scientific search channel to rely on customizable and revertible event concepts.

The removal of all triggers except the per-channel digitization threshold does not lead to significantly increased storage requirements. The primary driver of data rate is not PMT dark counts or other small signals seen by only a few channels, rather it is large S2 signals that can remain above threshold for $\sim 10 - 100 \mu\text{s}$, and are seen by a large number of channels. For instance, in typical conditions during SR0, S2s from single electrons account for 30% of all reconstructed signals but only 2% of the data volume. In contrast, for very large S2s these values are reversed, accounting for 2% of the reconstructed signals but 30% of the data volume. Any trigger would be configured to save the large S2s, so additionally saving everything else (mostly S2s) does not represent a significant increase on the requirements of long-term storage. To support this, work was done studying data formatting and data compression, and a storage format was chosen that compresses more efficiently than the XENON1T storage format. Further, the removal of the software trigger eliminates the requirement for a database that can act as a base for fast triggering software, so the readout processes write data directly to high-speed disks in a continuous stream.

2.3 Fast data processing & immediate data availability

In addition to the hardware upgrades from XENON1T to XENONnT, the readout and processing software was also upgraded. To handle the continuous data stream of roughly three times the number of channels, the processing framework *PAX* [21] was replaced by the generic framework *strax* [22] and implemented for XENONnT in *straxen* [23]. *Strax* and *straxen* are referred to as *strax* for simplicity. *Strax* is written in python and was initially based on a re-write of *PAX* with a different memory model. It uses packages from the scipy stack [24, 25], just-in-time compilation (numba) [26], and a tabular data format to allow for fast processing by exploiting autovectorization. While *PAX* achieved processing speeds of $\sim 100 \text{ kB/s/core}$, *strax* can process data at rates of $\sim 10 - 100 \text{ MB/s/core}$. *Strax* achieves its highest per core processing speeds when running on only a few cores but also allows parallelization to tens of cores, albeit at lower per core performance. For the data rates observed during SR0, including the associated calibration periods where a higher rate was expected, the processing time was much lower than the data collection time. The *strax* framework will be further elaborated on in section 4.

Strax does complete online reconstruction of all the data within $\sim 30 \text{ s}$ after the PMTs detect light. This allows the detector performance and stability to be monitored with high-level data without the need for selections or triggers. To enable remote, online access to the data while it is being processed at the DAQ, several types of data are uploaded to the MongoDB database in a dedicated collection. This database is accessible from outside LNGS, such that it can be retrieved from anywhere. During normal operation, these data are available online within $\sim 30 \text{ s}$. With online data access the performance can be monitored using fully reconstructed data. This is especially useful for stability checks, as well as detailed feedback on operations with rapidly changing conditions such as calibrations or changing field configurations. The fully reconstructed events are blinded below 20 keV , and only when purposely unblinded and reprocessed, the science results [8] are obtained. All properties of the reconstructed events are set to unphysical values for the blinded data.

2.4 Neutron veto DAQ

The goal of the neutron veto in XENONnT is to detect the capture process of those neutrons responsible for NR background events, which can mimic the interaction of a WIMP. A neutron detection tagging efficiency greater than 85% is desired [7]. Since the expected Cherenkov signal in the case of neutron capture by H is of about 20 PE in total, it is important to have a high detection efficiency for each photon. To achieve such a high efficiency in a trigger-based DAQ architecture it would have been necessary to reduce the number of coincident PMTs that form the trigger. This, in turn, would have led to an increased number of triggers and acquired data, making it challenging for the DAQ readout.

Therefore, the neutron veto DAQ is designed around a triggerless data collection scheme like the TPC. Its ability to provide both the pulse shape and the timestamp of each PMT signal supports data collection with fully independent channels without the use of a global trigger, typically based on channel multiplicity. As will be described in section 4, the event building is done in software after data acquisition, where timestamps and coincidences between PMT signals are used to define events. This architecture based on a readout system of independent channels allows the acquisition of all the PMT signals above the digitization threshold and the lowering of the energy threshold.

PMT characteristics such as dark rate, afterpulsing and timing resolution are essential for the choice of front-end electronics. In particular, the dark rate puts a limit on the detection of very small signals, and can be used to estimate the accidental coincidence rate with a defined number of PMTs within a specific time window. Operating with a threshold of 0.5 photoelectrons (PE), the measured PMT dark rate during detector commissioning was about 0.96 kHz, generating an accidental coincidence rate that exceeded 4 kHz for a 2-fold coincidence between two random neutron veto PMTs. In addition, the materials in the sub-detector itself (PMTs and stainless steel structure) induce events from radioactivity that mimic NR signals in the neutron veto of ~ 100 Hz.

To efficiently tag neutron events, the electronics must be able to acquire signals ranging between 0.5 PE and ~ 100 PE, requiring a wide dynamic range. The small signals last about 100–200 ns, resulting mainly from the dark rate, and define the lower limit for the neutron veto data throughput. In contrast, gamma and beta particles from materials radioactive decays with a typical rate of ~ 100 Hz exhibit waveforms that last up to 10 μ s (considering signals and associated afterpulses) in many channels, requiring a much higher data collection rate. Therefore, the readout electronics must be able to handle extremely different time acquisition windows with the presence of sharp peaks of data rate.

In addition, the fast response (few ns) of PMTs used by the neutron veto to acquire single photoelectrons requires a fast waveform digitizer for signal sampling. This high time resolution is necessary to efficiently separate neutrons produced close to and far from the TPC cryostat, where the former are the primary target.

2.5 Three integrated DAQ subsystems

One requirement for the XENONnT DAQ system was that the three DAQ subsystems (TPC, muon veto, and neutron veto) should be able to operate both independently and as one combined system. While both the TPC and muon veto in XENON1T used the same 50 MHz clock signal, there was no synchronization of the start signals issued to the two readout systems, so there was some variation




Figure 2. An illustration of the variety of signals read out by the XENONnT DAQ in ADC counts (ADCc). The left and middle panels showcase raw signals from a selection of two TPC PMTs in a single event. The inset in the left panel zooms in on the S1 to emphasize its narrow width and short risetime, contrasting with the wide S2 in the middle panel. An inset in the middle panel zooms in on minute signals in the leading edge of the S2 waveform. In both panels, the black lines correspond to signals from a PMT in the top array, and the blue from the bottom array. The red and purple dashed lines represent the baseline and digitizer threshold, respectively. The rightmost panel shows signals that were recorded by individual channels in the muon veto (green) and neutron veto (dark red) DAQ subsystems. These signals are not correlated with the ones showcased for the TPC channels. Typical muon veto and neutron veto thresholds are depicted with dashed lines of matching color. The higher sample rate of the digitizers employed by the triggerless neutron veto subsystem is clearly visible in comparison to the triggered muon veto. In this case, a relative baseline is shown in red for illustration purposes only. As explained in subsubsection 3.2.2, this muon veto signal is read out, despite not being above threshold, because of the hardware coincidence trigger during this time interval.

in reconstructed timestamps between the two detectors, and data from the two subsystems were analyzed separately. The trigger signal of the muon veto was recorded in one of the TPC’s digitizers, but this did not provide the equivalent timestamp in the muon veto’s data stream, thus viewing the corresponding event as observed in the muon veto required analysts to perform additional steps.

To ameliorate this, the XENONnT DAQ was designed to allow for the start signal from one subsystem to be issued directly to one or both of the others, essentially combining them into one and ensuring that timestamps recorded in one can be directly compared to those from another. Subsystems can be combined or “linked” together as determined by the requirements of the data being taken, or can operate independently. Subsystems operating in linked modes are controlled as a single operational unit, and the data they record are combined at the readout level and processed together to facilitate handling and analysis. Combining multiple detectors into a single system allows XENONnT to explore novel ways of utilizing the veto systems in coincidence with the TPC, studying low-energy phenomena such as the Migdal effect.

3 Data acquisition & readout

The data acquisition is organized in two broad schemes following the system used in XENON1T. At the greatest scope are so-called “science runs” which represent a period of months or even years with a targeted science objective where the detector conditions are held constant. For the daily operation of the experiment, the organizational unit called “runs” is used, where each run represents a continuous period of a few minutes up to a few hours using a set of configuration options that remain constant for the duration of the run.

The three DAQ subsystems rely predominantly on commercially-available analog front-end electronic modules supported by custom hardware. The firmware and software include both custom components and some provided by CAEN. All commercially available CAEN products are marked with their model number throughout this work, the reader is referred to the company’s website for more details and manuals [10]. All DAQ hardware is installed within eight racks located in the DAQ room on the first floor of the XENON service building in Hall B of LNGS. An air conditioning system provides cooling for the electronics, maintains a constant air temperature in the room, and reduces the collection of dust.

3.1 Analog electronics

The 494 PMTs of the TPC employ separate signal and high-voltage (HV) lines and a dedicated grounding scheme implemented via custom connectors and cables [27]. In the DAQ room, an array of multi-channel CAEN A7030LN, A1536LN and A1535LN HV supplies provides the power for the TPC PMTs via Radial connectors, each hosting 24 distinct HV lines, 2 return lines and a distinct shield line [28–30]. HV boards are known to produce high-frequency switching noise. Hence, before being supplied to TPC PMTs the output from each HV board is passed through a custom filter box. Within each filter box, every HV channel line goes through a low-pass filter, removing electronic noise with frequencies greater than ~ 250 kHz.

Before installation, all TPC signal cables were grouped based on the location of their corresponding PMTs in the arrays and assigned to specific hardware modules. The resulting cable map ensures an equal distribution of the data load on readout electronics and was used as a guide throughout the hardware installation process. Signals from both PMT arrays are passed through custom amplifiers as mentioned in section 2.1. As illustrated in figure 1, both the high- and low-gain signals from the top PMT array are propagated to dedicated groups of digitizers. However, only high-gain signals from the bottom PMT array are passed to digitizers. The low-gain signals from the bottom array are summed up using a cascade of linear fan-in/fan-out modules and are used by the high-energy veto (HEV).

The 120 PMTs of the neutron veto are connected to the readout electronics and HV system located in the DAQ room by means of 30 m coaxial cables with separate grounding for signal and HV cable lines. A custom-made patch panel mounted on the back side of the neutron veto rack gathers HV lines in one section and signal lines in another. Signal lines are directly connected to the front-end electronics via a panel feedthrough. HV lines are low-pass filtered to reduce high frequency noise and connected to five CAEN A7435SP HV boards [31]. The 84 PMTs of the muon veto are powered by four CAEN A1535SP boards, and employ a similar signal and HV cabling scheme to the one used by the neutron veto. However, as the noise requirements for the MV system

are lower the HV to its PMTs is not low-pass filtered. The aforementioned HV boards are housed in separate CAEN SY4527 Universal Multichannel Power Supply System crates [32]. Lastly, each DAQ subsystem hosts a range of logic modules that are used for distributing initialization and trigger signals.

3.2 Digital electronics

For time synchronization across its subsystems the XENONnT DAQ relies on a CAEN DT4700 clock generator module [33]. Its 50 MHz low-voltage differential signaling (LVDS) outputs are propagated via seven shielded custom-manufactured cables to the first digitizer in each VME crate. Shielding the clock-carrying cables reduces the amount of external noise that can be injected into the cables, improving the stability of the clock signal. The propagated signals are then distributed within each VME crate by shorter clock cables from digitizer to digitizer, ensuring the temporal synchronization of the entire DAQ system. Time offsets in these clock chains were manually calibrated out, securing synchronization well below the digitizer temporal resolution.

Additionally, a GPS timing module [34] is used to distribute a 0.1 Hz trigger signal to dedicated digitizers in each DAQ subsystem. Each trigger is associated with a GPS timestamp (accurate to ~ 15 ns), providing another layer of time synchronization within the DAQ. The same signal can also be used for absolute time synchronization with other experiments.

3.2.1 TPC

The core of the TPC readout is formed by 95 CAEN V1724 digitizers [35] running the DPP-DAW firmware, an updated version of what was used in XENON1T [11]. The V1724 is an 8-channel board featuring a sample rate of 100 MHz, a dynamic input range of 2250 mV (input impedance 50 Ω) with 14-bits of resolution, and an input bandwidth of 40 MHz. Of these digitizers, 62 read the 494 high-gain signals from the top and bottom arrays, 32 read the 253 low-gain signals from the top PMT array, and 1 acts as the TPC's Acquisition Monitor detailed in section 3.4. The boards are distributed across five VME crates and connected to readout servers via daisy-chained optical links. Most optical links contain the maximum of 8 digitizers, while the acquisition monitor is read out via its own dedicated optical link to ensure it never goes busy.

One CAEN V2718 crate control module [36] is used as a synchronizing module to produce the sync/start/stop digital input (S-IN) signal that begins and ends the acquisition. This signal is distributed to all digitizers via logic fans, with the signals all reaching their respective digitizers with a spread of < 4 ns. Additionally, this module provides gate logic signals that control the propagation of the S-IN signal to the muon and neutron veto digitizers that are activated during linked-mode operation. A periodic external trigger signal can be generated by this module, which is distributed both to all digitizers and an external LED pulser used to calibrate the response of the PMTs. The third type of module is a general-purpose CAEN V1495 board [37] running custom firmware, which manages the busy subsystem detailed in section 3.3.

Finally, two NIM crates hold the logic fan modules used to distribute the S-IN and trigger/veto signals to all TPC digitizers, as well as a gate module, a NIM-TTL level converter, and a delay generator. These latter two are used to connect and synchronize the TPC DAQ with the LED calibration system.

3.2.2 Muon veto

The muon veto readout is unchanged from XENON1T as described in [11], though some additional connections were made between this subsystem and those of the TPC and neutron veto. Eleven V1724 digitizers with the default ZLE firmware form the readout system, although the zero length encoding features are not used. Three optical fibers are used to read out these digitizers. A V2718 module provides the S-IN signal for these digitizers during unlinked operation. A CAEN V976 unit [38] serves as a logic fan to distribute both this S-IN signal and that of the TPC during linked operation to all muon veto digitizers. A V1495 board acts as a programmable trigger unit, allowing the user to specify both the number of participating channels and the coincidence window necessary to generate a hardware trigger.

3.2.3 Neutron veto

To take advantage of the fast response of the 120 PMTs and to efficiently reconstruct the fast component of Cherenkov photons in the neutron veto sub-detector, eight CAEN V1730S new generation digitizers [39] are used to acquire PMT signals. Each V1730S board is a VME 6U module housing a 16-channel 14-bit 500 MHz flash ADC. The input dynamic range can be set to either 2 V or 0.5 V on single ended MCX coaxial connectors. During commissioning and SR0, the 2 V dynamic range was used. The input section is $50\,\Omega$ -coupled and feeds a programmable gain amplifier to select the suitable analog range. In case all the buffer memory is filled, a busy condition occurs and a logic module inhibits the data acquisition for all the boards (as described in section 3.3). The V1730S digitizers are operated with the DPP-DAW firmware like the TPC digitizers. An exemplary neutron veto waveform is shown in figure 2.

The V1730S module is also able to work with a common global trigger, either coming from the digital input external trigger (TRG-IN) input or a coincidence trigger. In particular, the external trigger mode is used by the neutron veto system during calibration. A V2718 board hosted in the VME crate generates several control signals (mainly the start-of-acquisition and calibration signals) that are subsequently distributed to the digitizers via logic fan-in/fan-out modules. Two additional boards are hosted in the neutron veto crate: a V1495 to manage the V1730S busy signals and provide the veto signal, and a V1724 digitizer that serves as an acquisition monitor.

The neutron veto digitizers are connected to a readout server via two optical links; one daisy-chains the V1730S digitizers while the second is for the V1724 acquisition monitor. In order to synchronize all the digitizers in the neutron veto DAQ and limit the clock uncertainties to below ~ 1 ns, an external common clock reference feeds all the modules. The V1724 digitizer receives the common 50 MHz clock signal (see section 3.2), which is then upconverted to 62.5 MHz via a phase locked loop device and propagated through the V1730S boards. Lastly, several auxiliary electronic modules are hosted in a NIM crate, managing the distribution of calibration triggers and run start signals.

3.3 Busy & high-energy veto

The V1724 and V1730S digitizers have a limited on-board memory buffer for storing data between digitizing and readout, amounting to 1 MB/channel and 10.24 MB/channel, respectively. If incoming data accumulates in the digitizer's memory buffer faster than it is read out, the buffer will become

full and the digitizer will no longer be able to acquire new signals, rendering it *busy*. To ensure the integrity of individual events the triggerless TPC and neutron veto DAQ subsystems employ a hardware-based *busy veto*. Each subsystem hosts a general-purpose V1495 board, equipped with field-programmable gate array (FPGA) firmware developed in-house. When a digitizer enters the busy state it emits an LVDS signal via a pair of connectors on its front panel. These signals are propagated via ribbon cables from each digitizer in each subsystem to its respective V1495 module.

Whenever the V1495 recognizes that any of the digitizers emits an LVDS busy signal, it outputs a veto NIM signal for a fixed duration. This veto signal is distributed to all the digitizers within the subsystem, inhibiting data acquisition for 1 ms or until none of the digitizers are busy. Within the FPGA firmware, busy intervals are assigned with start and stop NIM signals. These are also output from the V1495 board and propagated to the relevant acquisition monitor digitizer, as explained in section 3.4. The TPC V1495 board has a more advanced version of this firmware. Besides being responsible for the busy veto, it is also capable of generating an artificial periodic veto, with user controlled duration and frequency. During detector commissioning, the water tank was empty and the detector was not shielded from radiation in the experiment hall. This periodic hardware-induced veto allowed the TPC DAQ to handle high background rates in addition to taking $^{83\text{m}}\text{Kr}$ calibration data. Additionally, the V1495 module performs several other important functions. In the TPC DAQ it collects and propagates the HEV signal, and assigns it with start and stop NIM signals that are read by the TPC acquisition monitor. In both the neutron veto and TPC subsystems the V1495 board is also responsible for the propagation of the LED trigger to the digitizers during LED calibration.

The TPC DAQ also employs a hardware veto to reduce the load on the system during acquisition of high-rate calibration data. This HEV was developed based on a commercially-available multipurpose digital pulse processor DDC-10 from SkuTek [40]. It hosts a variety of chips and daughter cards on a BlackVME S6 motherboard, including a Spartan-6 FPGA and a 100 MHz, 10-channel, 14-bit ADC. Mounted on the FPGA is custom-developed firmware whose main goal is to identify and veto high-energy S2 signals. The HEV digitizes the analog sum signal from the low-gain TPC channels of the bottom array, determining the risetime, width and integral of signals that surpass a threshold of 200 ADCc (ADC counts) above the baseline. If any identified S2 exceeds predefined parameter threshold settings, the HEV issues a 3 ms veto NIM signal. The thresholds set on the parameters vary depending on the utilized calibration source and the data rate induced by it in the TPC. For example, during ^{220}Rn calibrations thresholds of 350 ns, 200 ns and 40 000 ADCc were set on the width, risetime and integral parameters, respectively. To provide the HEV module with enough time to make the veto decision, data readout is delayed within the TPC digitizers by 10 μs . The veto signal generated by the HEV is propagated to the TPC V1495 module, from where it is distributed to each TPC digitizer.

Drift field conditions in the TPC during SR0 produced broad S2 signals with widths greater than $\sim 10 \mu\text{s}$, which were found to have the largest contribution to the DAQ rate. It is difficult to identify and characterize the shape parameters of such signals within the 10 μs time limit. Hence, the HEV firmware has an additional operation mode, whose purpose is to veto low-amplitude high-width S2 signals that might last for $\sim 10 \mu\text{s}$. In this mode, if the HEV is not able to determine the width and the risetime of the signal within $\sim 5 \mu\text{s}$ and the signal's amplitude is still above the HEV threshold it will consider the signal to be a high-width S2, and will issue a veto. The aforementioned HEV operation modes can be utilized separately or run in parallel.

It should be noted that similar hardware veto implementations were employed by the XENON1T DAQ [11], yet, for XENONnT both systems were expanded and their firmware was refurbished. A schematic view of the hardware-based veto systems described above is shown in figure 1. The operation of the HEV results in raw data reduction at the readout stage of up to 40%, depending on the utilized HEV settings. Throughout SR0 the HEV was utilized during AmBe and ^{220}Rn calibration data taking.

3.4 Acquisition monitors

The TPC and the neutron veto DAQ subsystems each host a dedicated V1724 digitizer, whose aim is to collect information about the status of the DAQ itself and the operation of its hardware veto modules. Both the TPC and neutron veto acquisition monitors receive the start and stop NIM signals from their respective V1495 veto modules, which indicate the boundaries of busy veto intervals. Additionally, these digitizers also acquire the 0.1 Hz NIM synchronization signal from the GPS module [34]. Uniquely for the TPC, its acquisition monitor also digitizes the same analog sum waveform signal that is seen by the HEV module. To prevent the TPC acquisition monitor from ever going busy a relatively high threshold of 100 ADCc or 14 mV is set on this channel. Furthermore, acquisition monitors are also excluded from the busy veto distribution scheme.

Acquisition monitor data are read out identically to the rest of the digitizers and incorporated into the overall data processing chain. These data are then used to diagnose the performance of the busy and HEV systems, and to determine the deadtime they induce. The measured deadtime under several operational modes is discussed in section 6.2. Moreover, the same data are used as a basis for a data quality cut. The cut removes any events that could be misreconstructed due to missing information as a result of their proximity to a busy or a HEV veto interval. The cut decreases the livetime and is accounted for in the exposure rather than the cut acceptance [8].

3.5 Servers & software

Five server computers are responsible for the readout of all the digitizers, three for the TPC and one each for the muon veto and neutron veto. Two additional servers provide backup capacity. The TPC readout servers each have four 960 GB write-intensive solid state drives which are configured together as a Ceph cluster [41] to form a single high-speed buffer disk with approximately 10 TB of capacity that is accessible from all servers within the DAQ network. While replication is possible using Ceph, it is not necessary for a short-term buffer disk, so the configuration is equivalent to RAID0 (data striping) to provide the highest access speeds. This buffer disk can sustain simultaneous read and write operations from multiple sources at rates exceeding 1 GB/s. Data are stored on the Ceph buffer from the start of acquisition until the live processing for that run has successfully concluded, which is typically one or two hours, so very little data are lost in the event of disk failure. The combined data rate from the three subsystems during science data taking is approximately 40 MB/s, so the disk can potentially buffer data for a considerable amount of time in case of issues in the live processing.

Each readout server is equipped with at least one CAEN A3818 PCIe interface card [42]. Each A3818 supports up to 4 optical fibers, with each fiber capable of daisy-chaining up to 8 digitizers and supporting a maximum data throughput of 80 MB/s–90 MB/s. Digitizers and optical links are distributed to approximately balance the load on each of the readout servers.

The readout servers all run the *redax* software package [43], which copies data from the digitizers and transforms it from the digitizer-native format into one compatible with the strax data processing package [22, 23]. Data are read from digitizers in block transfers via the CAENVMELib C++ library, where each optical fiber is read out exclusively by a dedicated readout thread. A round-robin technique is used, where each board on an optical fiber is successively polled. When a digitizer has data available for readout, block transfers are performed until all data stored on that digitizer have been read into the server’s memory. The readout threads then transfer data asynchronously to processing threads, where the binary format transformation is performed. Each processing thread periodically compresses its buffered output data and writes it to the Ceph buffer in fixed-time intervals called chunks following the chunking paradigm in strax as described in section 4. Chunks are labeled with the name of the readout process, the chunk number, and also the ID of the thread that wrote that chunk, which acts as a unique identifier. Additionally, *redax* is responsible for programming the digitizers in preparation for each run via configurations it obtains from a central database. *Redax* also writes status snapshots to this database once per second, including quantities such as the current state of that instance of *redax*, the amount of data currently buffered in memory, and the data rate for each channel of each digitizer being read out.

Six additional servers, called the *eventbuilders*, are responsible for the live processing. In case of high data rates or unlinked operation, when the three DAQ subsystems run independently, multiple hosts can process data simultaneously. For low data rates only a single host is required for the processing. Three eventbuilders are PRIMERGY RX2540 M4 Fujitsu servers with two Intel® Xeon® Gold 6128 CPUs at 3.40 GHz and 202 GB of RAM each. Additionally, there are three backup PRIMERGY RX2540 M1 Fujitsu servers with two Intel® Xeon® E5-2660 v3 CPUs at 2.60 GHz and 135 GB of RAM each. These backup servers were also used in XENON1T. Two of these three mainly serve as extra redundancy, while the third acts as a general purpose machine with access to the latest data. This machine, for example, automatically produces online monitor plots and handles requests for making and retrieving said plots (as explained in section 4.4).

4 Live processing

The data stream of raw data from the digitizers is fully processed onsite at LNGS. The triggerless data stream is handled by the data stream processor, strax [22, 23]. Using live processing and online data storage, data can be accessed while their collection is still ongoing.

4.1 Data stream versus discrete events

The triggerless design of the XENONnT DAQ results in a continuous data stream. For processing as well as storage purposes, handling discrete time intervals of data is advantageous as it allows for parallelization. To this end, the digitizer data which are read out by *redax* [43] are partitioned in 5 s–20 s time intervals called chunks. Each chunk is accompanied with an overlap region of ~ 0.5 s to the previous and following chunk. As such, the overlap region is being saved twice, once with the previous chunk and once with the following. These overlap regions, called the pre- and post-chunk, are processed together with a chunk to ensure that each process has access to sufficient data to do the reconstruction. This is important for reconstructing S1/S2 signals (peaks) whose data might otherwise be split into consecutive chunks. Strax searches for time regions where there are no data

for $1\ \mu\text{s}$ within the pre- and post-chunk and discards the data before (pre-chunk) or after (post-chunk) this time region. This discarded time region will instead be processed together with the previous or next chunk. If no time interval of $1\ \mu\text{s}$ is found within the overlap region, artificial deadtime would be inserted, which was never required for the entire SR0 dataset including calibration data. Due to this temporal separation between chunks, single chunks of low-level data are handled independently, allowing for processing in parallel. For high level data, such as events, the processing is based on stateful algorithms which are therefore single threaded and may rearrange chunk boundaries.

4.2 Strax(en) data format

The processing at LNGS is handled by the eventbuilders and uses the publicly-available strax framework [22, 23]. Strax is a purely python-based streaming processor. Autovectorization, just-in-time compilation, and a tabular data format make the processing fast. The tabular data format is achieved by fixing the shape of the data fields in software. At the level of PMT traces (figure 2), this is achieved by splitting one variable-length PMT trace into a sufficient number of fixed-length intervals. The data are organized in a hierarchical structure of “datatypes”. At higher level datatypes, like S1/S2-peaks, the summed waveform of all PMTs is down-sampled to a fixed number of samples.

There are several steps in the processing, from PMT-traces as in figure 2 at the lowest level, to a fully reconstructed S1 and S2 pair originating from one physical interaction within the TPC at high level. The PMT-traces are stored as *raw-records* as they are the lowest level (raw) datatype that is stored long term. The S1s and S2s are saved as *peaks* level data which can be grouped in time to form *events*. Peaks primarily store the sum waveform from all PMTs, as well as per-PMT information like the area per channel, and shape parameters from the sum waveform [44]. Events store the information on the largest two S1s and largest two S2s. Additional properties of peaks and events can be added through the modular plugin system described below. The time scales of the typical objects in raw-records, peaks and events differ by orders of magnitude. For example, figure 2 shows that raw-records can be of $\sim 1\ \mu\text{s}$ and an S2 peak of $\sim 30\ \mu\text{s}$. The duration of an event is set to be at least as long as the drift length (2.2 ms). Correspondingly, for higher level data, the number of items and the data size decreases by orders of magnitude. For instance, an hour of data may amount to 250 GB of raw-records, 5 GB of peaks, and only 30 MB of events.

The different levels of data processing are organized in software modules called *plugins*, each producing one or more datatypes which can serve as the input data for subsequent (higher level) plugins. This structure allows for a modular design and a flexible processing framework. When a chunk of data is processed, it is transferred between processing threads to any higher level plugins requiring it as input. Using this structure, the versioning of the data is handled per datatype by tracking the dependency chain. This has the benefit for reprocessing that, for example, a new or modified plugin at event level can only require event level input data to (re)compute and does not affect any lower level datatypes. During processing, auxiliary information on several quantities required for data processing, like PMT gains, are queried from a dedicated collection within the MongoDB database. This collection is frequently updated with the latest values to ensure the data are processed with as up to date corrections and detector variables as possible.

4.3 Online processing

The processing by strax includes several stages, a full description of all its aspects is beyond the scope of this paper. Instead, some of the aspects are briefly discussed to illustrate that the full reconstruction of all the data is done live. The lowest level data, the raw-records, are all written to disk without further processing, allowing to always go back to the unprocessed data. After the raw-records level, PMT-traces are baseline-subtracted, inverted and integrated. For the TPC, time intervals are obtained wherein photon hits are extracted from the PMT traces to build peak sub-clusters, *peaklets*. The peaklets are classified and re-clustered according to their type to obtain peaks; S1-peaks are assumed to consist of only one peaklet, while S2-peaks can consist of many. Using this two-step clustering, strax is able to deal with the very short S1 signals while also being able to reconstruct the longer S2 signals as single peaks. Three different neural networks are applied on the peak level data for xy-position reconstruction based on the PMT hit pattern, which allows for cross-validation of their results. Due to the lack of a global triggering scheme, a reconfigurable “software trigger” defines the concept of events. Events are built on the basis of a large S2 peak (the “triggering” peak). The triggering peak should be >100 PE and there should be fewer than 8 other peaks with at most 50% of the area of the triggering peak in a ± 10 ms window around the triggering peak, to prevent building or extending events based on too small S2 peaks. An event is the time region from 2.45 ms before and 0.25 ms after the triggering peak. This time region is set to be longer than the drift length of 2.2 ms (in SR0) and all peaks within the time region are considered part of the event. Multiple triggering peaks within the event window extend the event duration to capture multi-scatters. This is effectively the event trigger, which is set as a high-level configuration in the processing chain, in stark contrast to XENON1T [11], where an event window was fixed once at the DAQ and all other data was discarded. As a result, the event trigger is easily re-optimized in a high-level analysis.

Processing of muon veto and neutron veto data is also performed within strax using dedicated veto plugins which are applied similarly to both types of veto data. These plugins reconstruct veto-events based on the number of PMT hits. Additionally, a software coincidence trigger for the neutron veto reduces the data stream at a low data level. This software trigger is not used for the muon veto because it has a hardware coincidence trigger.

The program *bootstrax* is responsible for processing on the eventbuilders and is optimized per host machine to provide the maximum performance under a wide variety of data rates. As soon as a new run is issued by the dedicated program (the dispatcher, see section 5.3), bootstrax looks for newly written chunks on the Ceph buffer disk. Bootstrax marks a set of data ready for uploading into long-term storage after completing the processing.

4.4 Online monitoring

The live processing on the eventbuilders is able to keep up with the data rates observed during SR0, including all the calibration periods. This opens up possibilities to use fully reconstructed data to monitor the state of the detector while data collection is ongoing. To this end, several datatypes are uploaded while data are being collected. This includes acquisition monitor data, all the fully reconstructed events and selections of data from the muon veto, neutron veto, and a selection of the peaks data from the TPC.

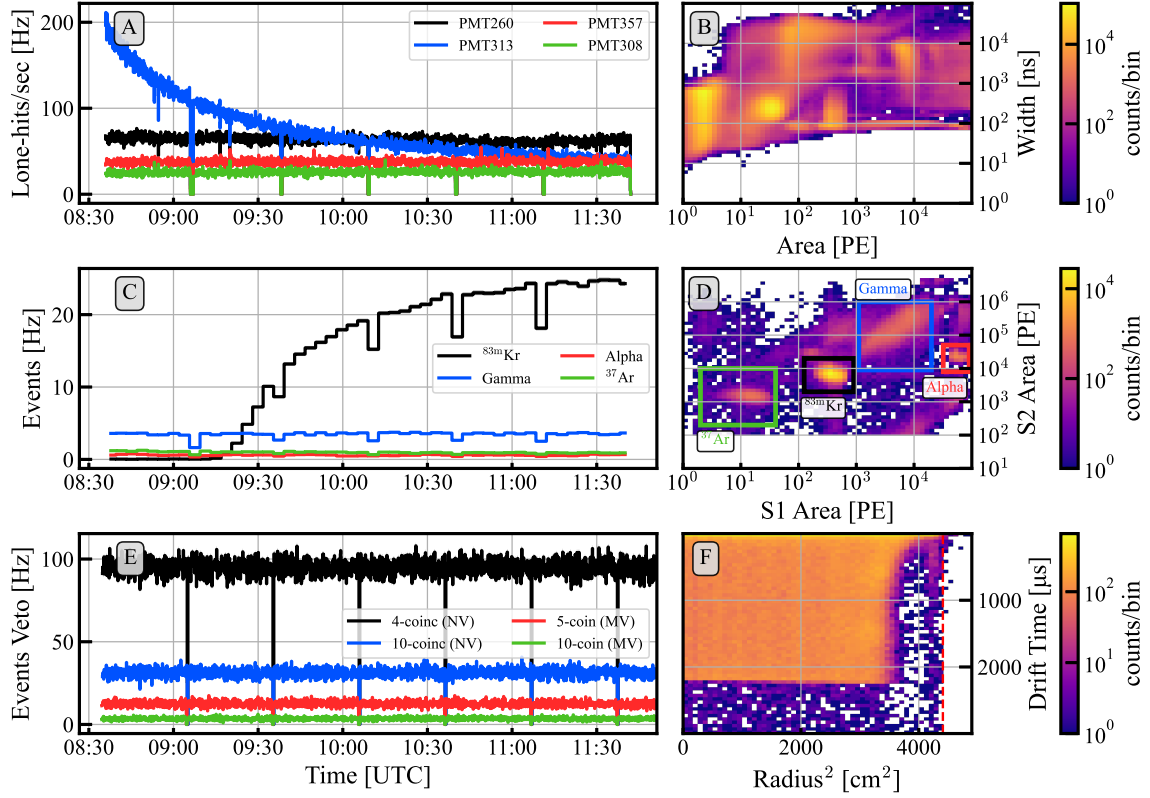


Figure 3. Online monitor plot for monitoring the detector status. Panel A shows the per PMT *lone-hit* rate, which are pulses that are seen in one PMT without any pulses in other PMTs within a short time interval. PMT313 (blue) just “flashed” [45], and is slowly returning to a rate comparable to the other PMTs. Panel B shows the area of peaks versus the range of 50 percent decile (known as the width) of the sum-waveform of a peak. This parameter space is useful for identifying peak populations, e.g., the peaks from $^{83\text{m}}\text{Kr}$ S1 are visible in the range 80-800 PE at roughly 100 ns width. Panel C shows the evolution of reconstructed events, which are roughly selected on their S1 and S2 area as shown in figure panel D. Since this was the start of a calibration period with $^{83\text{m}}\text{Kr}$ following an ^{37}Ar calibration, the event rate of $^{83\text{m}}\text{Kr}$ is increasing in panel C while the ^{37}Ar remnants are being removed via online distillation to a negligible level [46]. Panel E shows the evolution of the number of veto events in the veto-systems over time. Panel F shows the reconstructed event positions throughout the TPC, where at larger drift times (deeper into the detector), the events are reconstructed inward due to an inhomogeneous drift field inside the detector [47]. The drops in the rate to 0 Hz (panels A and E) mark the periods where the DAQ is switching from one run to another. As the data in panel C are re-binned, the run transitions manifest in $\sim 10\%$ drops in the rate.

Redax buffers at least two chunks in memory, which first have to be written to the Ceph buffer disk before that data can be processed. Several chunks are usually combined in memory during processing before writing it to disk to reduce the number of small files in long-term storage. However, when chunks of processed data are uploaded to the MongoDB database, there is no such limitation and a chunk of processed data is therefore immediately uploaded after processing. This usually results in the data being available in the database ~ 30 s after light has been detected by PMTs.

Status overview plots to monitor the detector conditions and data quality are made with the open-source infrastructure [22, 23] and additional XENONnT software. As an example, figure 3

shows two unrelated changes in detector conditions close in time: a period of intermittent light emission (“flash”) of PMT313 [45] (apparent from panel A) and the start of a calibration period with $^{83\text{m}}\text{Kr}$ (most clearly visible in panel C). This figure can be produced continuously to see changing detector conditions live. Additionally, each hour a plot is automatically produced and sent to the XENONnT-Slack [48] workspace which is used as the common chat room for the entire experiment. On Slack, one can also easily request periods of time to make this plot for, which is handled automatically by one of the backup eventbuilders. Alternatively, the data can be directly retrieved from the MongoDB database or via strax to do custom analysis, for example to create SuperNova Early Warning System (SNEWS) [49] warnings.

4.5 Data storage infrastructure

During the commissioning of XENONnT and the first science run (SR0), the DAQ collected >2 PB of uncompressed data. To reduce the required amount of long term storage, aggressive compression algorithms like *bz2* (in case of low data rates $\lesssim 65$ MB/s) and *zstd* (in case of higher data rates) are used to compress the low-level data. The *bz2* and *zstd* algorithms compress the raw data by factors of 4.7 and 3.5, respectively. While this increases CPU usage on the eventbuilders compared to the faster compression algorithms used for high level data, such as *blosc*, CPU usage is usually not the constraining factor for the eventbuilders. The XENON1T compression format, that used a combination of *zlib* and *pickle* compressed the data by a factor of 3.4.

The eventbuilders write to their own hard disks configured in a RAID5 configuration for performance and redundancy resulting in 22 TB of storage per eventbuilder. These hard disks, shown as local storage in figure 1, can be accessed by other hosts within the LNGS network. The data are uploaded from these disks into long-term storage as soon as bootstrax marks them ready for upload in the MongoDB database.

5 System control & oversight

Control and oversight of the DAQ and its associated subsystems are handled via databases, a user-interface website, and a software controller that coordinates the readout processes. Two additional servers are used in these roles, one in the LNGS surface server room and one underground with the other DAQ servers in the DAQ room. The surface server hosts the necessary databases for the DAQ, and also acts as a secure gateway through which experts can remotely access the DAQ subnet underground. The underground server hosts the DAQ website and the software controller.

5.1 Databases

Two databases are used for system control, monitoring, and interprocess communication, both implemented in the NoSQL-based MongoDB [15]. These are referred to as the “DAQ” and “Runs” databases. Each database is subdivided into “collections”, analogous to tables in an SQL-based database, each containing “documents” which are analogous to rows. Unlike SQL-based databases, documents in one collection are not required to have the same schema, which allows for considerable flexibility.

The Runs database is a three-node replica set of servers located in LNGS, University of Chicago, and Rice University, the latter two being the primary XENONnT analysis facilities. This ensures that

analysts have access to the database in the event of transient network disruption and protects against data loss due to hardware failure. One collection in this database contains metadata for each discrete run, which includes quantities such as the run start and end times, which of the three detectors were assigned to this run, the full readout configuration of each detector, and a listing of all datatypes for this run and their storage locations. Another collection contains data produced for online monitoring as described in section 4.4.

The DAQ database, in contrast, is expressly for the operation of the readout and not required for analysis. All data in this database are either only stored temporarily, or change very infrequently and can be restored from periodic backups in the event of data loss. This database, therefore, is neither replicated nor directly accessible outside of LNGS. Several collections in this database contain the regular status snapshots used to monitor the various components of the DAQ. These collections include the status snapshots of the redax readout processes, the health and performance of all the servers and NIM/VME crates in the DAQ system, and the status of the live processing. These collections are configured with time to live (TTL) indexes to only store data for 3 days, primarily to ensure that queries against these collections remain fast, and also because the information contained can either be reconstructed from the processed data or is additionally written to disk for long-term availability. Other collections store all available readout configurations, the system operational goal state set by the website, commands being issued to the readout processes, and important logging messages from the various DAQ processes.

5.2 User interface website

To facilitate easy use of the DAQ for the day-to-day operation of the experiment, a front-end website was developed using NodeJS [50]. A variety of pages allows users to view the current readout performance, set the system’s operation goal state, and monitor the data rates from each readout channel to identify potential problems in the detector, such as localized regions of sustained electron emission, also known as “hotspots”. The *status* page displays instantaneous data rates for each readout process in the entire DAQ system, information about the current activity of each eventbuilder, and the current status of the Ceph high-speed buffer disk. Additionally, a plot displays the recent data rate for each readout process, which allows a user to identify transient behavior in the system that may not be clear from the instantaneous rates alone. This page is shown in figure 4. The *control* page allows users to specify the operational goal state of all three subsystems, such as selecting an operational mode and a desired run duration. The *monitor* page displays the instantaneous data rate for each channel in the TPC in a convenient layout mirroring that of the physical locations of PMTs in the TPC (shown as the inset in figure 4). In this view, a hotspot will appear as a localized region, typically three adjacent PMTs, with a data rate significantly higher than other nearby PMTs. Additionally, a plotting function is provided to allow for a direct comparison of recent rates between different channels. Another page provides a user interface to the Runs database, where metadata about each run can be viewed. Other pages allow experts to modify and create preset operational modes and configurations, monitor the status of all the servers in the DAQ network, and interface with the dispatcher (described below).

Finally, an application programming interface (API) is provided to enable programmatic control of the DAQ and access to part of the DAQ database. This is used, for instance, by Slow Control to perform the periodic automatic calibration of the PMTs via a pulsed LED. Slow Control continuously

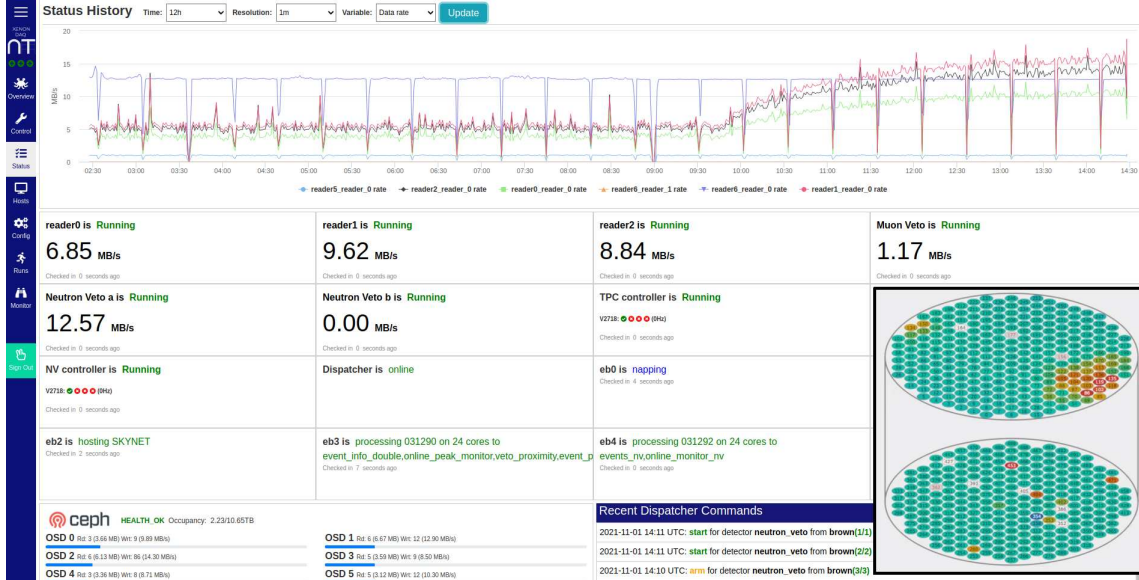


Figure 4. The Status page on the DAQ interface website. The plot shows the data rates for the past 12 hours, and the status cards show the instantaneous statuses of all readout and live processing elements in the DAQ system. An increase in the rate due to the start of regular detector calibration with $^{83\text{m}}\text{Kr}$ is clearly visible. A navigation bar on the left provides convenient links to other pages in the website. The inset in the bottom right shows the per-channel rate as displayed on the monitor page.

queries the API, and notifies experts if any aspects of the performance of the system deviate from what is expected, if disks are full, or a hotspot is suspected based on the per channel data rate.

5.3 Readout coordination software

To oversee and coordinate the readout processes, a program called the *dispatcher* was developed. The primary responsibility of the dispatcher is to convert the desired operational goal state as specified on the website into direct commands issued to the various readout processes. To do this, the dispatcher retrieves the most recent status snapshot of each readout process. These are aggregated together to determine an overall status for each subsystem, such as if the subsystem is idle, running, in a transitional state, or if processes are not responding. This aggregated status is then compared to the desired operational goal state of each subsystem, and commands are issued to the readout processes to make the former match the latter. For example, if a user wants the readout to begin with a certain operational mode, the dispatcher will ensure that the necessary processes are capable of starting, issue commands to begin the digitizer-programming sequence, wait until all necessary processes report the successful completion of this sequence, and then issue the start command. When an active run reaches the desired length as specified by the website, the readout is stopped, and the cycle is repeated. In rare cases where a readout process or digitizer stops responding properly, the dispatcher will automatically kill and restart delinquent readout processes and power-cycle VME crates as necessary to restore normal behavior. Additionally, if such action is necessary during linked-mode operation and restarting a process or VME crate fails to rectify the situation, the dispatcher will unlink the detectors so the readout of the detectors that are responding normally can

continue. Experts receive notifications whenever automatic actions such as these are taken. At the start of every run, the dispatcher creates an entry in the Runs database containing a copy of the readout configuration and other metadata necessary for the live processing. At the end of a run, the corresponding entry is updated with additional quantities such as the end time of the run and the average data rates for all contributing detectors.

6 Performance

In the first two years of operation, the TPC subsystem collected more than 1280 TB of data, while the muon veto and neutron veto collected 28 TB and 680 TB, respectively. The performance of the overall DAQ system can be measured in several ways. Most obvious are the maximum data rate the system can maintain and the livetime with which the system operates, but other criteria such as inter-detector synchronization and noise levels are also important. An additional key performance metric is the speed of the live processing, as it is crucial that the data are made available for transfer off-site at least as fast as it is recorded.

6.1 Noise levels

Externally triggered, short runs were taken weekly throughout SR0 to assess the noise levels in all TPC channels, using fixed windows of approximately 1 ms duration. The mean RMS noise level was found to be stable for each PMT array with values of 0.23 mV (~ 0.01 PE) and 0.34 mV (~ 0.02 PE) for the high-gain channels in the top and bottom arrays, respectively, and 0.16 mV for the low-gain channels in the top array. The installed filter boxes are effective in suppressing electronic noise > 250 kHz, which is related to HV power supplies. However, as expected, the filter boxes have a negligible effect on the low-frequency noise peak at 24.41 kHz, which is correlated with intrinsic noise produced by CAEN digitizers. Channels from the bottom PMT array on average exhibited an RMS noise level which was ~ 1 ADCc higher when compared to top PMT array channels. This effect could be attributed to either a different resistor type and assembly procedure that was employed for filter boxes used for the bottom PMT array, or different and noisier HV power supplies. These low levels of noise support low digitization thresholds for the PMTs. Over 98% of TPC channels have thresholds set at 2 mV (~ 0.13 PE) and only 1 PMT has a threshold higher than 3.4 mV (~ 0.3 PE), giving an average acceptance to single PE $> 90\%$. For the neutron veto, 109 PMTs (91%) have thresholds set at 1.8 mV (~ 0.3 PE) and the other 11 at 2.4 mV (~ 0.4 PE). The average noise for the neutron veto PMTs is 0.3 mV (~ 0.05 PE). Lastly, for the muon veto the average RMS noise is 0.18 mV, the thresholds are set for all channels at 3 mV (~ 1 PE). All above voltages were estimated for an input impedance of $50\ \Omega$.

6.2 Livetime

Throughout SR0 and the commissioning of XENONnT all three DAQ subsystems operated stably, collecting in total more than 200 days of commissioning and science data, and close to 100 days of various calibration data. The deadtime fraction induced by the operation of the busy veto is 2×10^{-5} for the majority of SR0 science data (which is typically $\lesssim 25$ MB/s), as illustrated in figure 5. The average deadtime fraction for all SR0 science data is 3×10^{-4} . Muons traversing the TPC cause short periods in time with high a higher data rate ($\gtrsim 25$ MB/s during science data taking) and a




Figure 5. Mean total deadtime per chunk (top panel), and relative frequency of chunks (bottom panel) as a function of data rate per chunk for several run modes. Each chunk is a time interval of 5 s to 20 s. AmBe calibrations are performed by keeping the source at several positions with respect to the TPC, leading to distinct populations in the bottom panel. The High-Energy Veto (HEV) run modes (AmBe HEV and ^{220}Rn HEV) have higher deadtime fractions as a result of the inserted deadtime by the HEV, see section 3.3. For typical (98%) science data ($\lesssim 25$ MB/s) the deadtime fraction is 2×10^{-5} . In science data, higher rate data points are caused by short periods in time following a muon traversing the TPC, leading to high data rates and deadtime fractions of $\sim 1\%$. ^{220}Rn has a lower deadtime fraction for $\gtrsim 25$ MB/s than science data, since these higher rates are caused by a higher S2 rate, rather than muons. Above ~ 250 MB/s the data quality deteriorates due to the onset of pileup.

higher deadtime. During high-rate ^{220}Rn and AmBe calibration periods the deadtime resulting from the combined operation of the busy and HEV on average amounts to $\sim 10\%$. The above deadtime values describe only the intrinsic deadtime produced by the operation of the busy and HEV modules, and not the data reduction caused by the analysis cut described in section 3.4. Lastly, it should be noted that the busy veto-induced deadtime of the NV was found to be negligible.

The highest sustained data rate in SR0 was ~ 500 MB/s during an AmBe calibration (population near 500 MB/s in the bottom panel of figure 5). The DAQ was designed to withstand higher data rates, but performing high rate calibrations in SR0 was inhibited by the long event duration which

leads to pileup of events where they start overlapping. A few chunks were digitized with data rates of up to 600 MB/s.

6.3 Time synchronization

Throughout SR0 data-taking the three DAQ subsystems operated in linked mode with full temporal synchronization. However, as was described above, to facilitate the operation of the HEV the TPC data are delayed within the V1724 digitizers by 10 μ s, as compared to the TPC acquisition monitor and the other sub-detectors. The time synchronization across all sub-detectors was verified using the 0.1 Hz signal generated by the GPS. It was supplied to a dedicated TPC digitizer channel, as well as to its acquisition monitor. Additionally, this signal was acquired by the acquisition monitor of the neutron veto and a digitizer in the muon veto. A comparison between the timestamps of these signals recorded on a time period of about three months was used to measure with high precision (<0.1 ns) the average temporal difference between sub-detectors. The average time difference between the TPC and neutron veto was measured to be 10 157 ns, while the time difference between the TPC and muon veto was found to be 5283 ns. The variation in the measured delay time is related to the trigger formation time used by the muon veto. After SR0 these constant offsets are subtracted out during readout, resulting in time synchronization with a precision of ~ 10 ns, which is comparable to the sampling time of the digitizers.

To illustrate the inter-detector synchronization a muon event passing through all XENONnT sub-detectors is presented in figure 6. The signals obtained in each detector are aligned by accounting in software for the time differences described above. In addition to the sub-detector signals, also the analog sum waveform that is acquired by the acquisition monitor of the TPC is shown in the top panel. As seen in the inset of the top panel of figure 6, the prompt S1 from the relativistic muon's interactions is followed by a sustained S2, which lasts for the full drift time of the detector. This extended S2 indicates a vertically traversing muon interacting along most of the drift column to produce ionization electrons. The subsequent long tail is formed from photoionization electrons that follow the S2 for multiple milliseconds. The drop in the digital sum waveform at ~ 1400 μ s is caused by a baseline fluctuation. Such baseline fluctuations were observed only after violent muon events and the exact cause is still under investigation, but they may be potentially caused by effects such as PMT base saturation. Monitoring digitizer thresholds and conducting weekly LED calibrations ensures the stability of the DAQ and PMTs, as well as the quality of acquired data. Additionally, any data in close proximity to muons is excluded from scientific analysis by dedicated cuts. The muon track in the muon veto and neutron veto sub-detectors corresponds to the Cherenkov photons detected by the PMTs of these detectors, followed by PMT afterpulses lasting up to ~ 10 μ s. The general shape of the waveform in all three panels is the same, indicating clearly that the same event is shown. Lastly, the GPS synchronization signal was also used to evaluate the clock drift of the DT4700 clock module, yielding approximately 2 μ s over a period of 10 s, or 0.2 ppm.

6.4 Live processing performance

The DAQ and its software were designed to run under high data rates during detector calibrations using a combination of internal and external radioactive sources to quantify the detector performance. However, due to a limited voltage on the TPC cathode during SR0 compared to the design value, high rate calibrations could not be performed as events quickly piled up because of the long drift

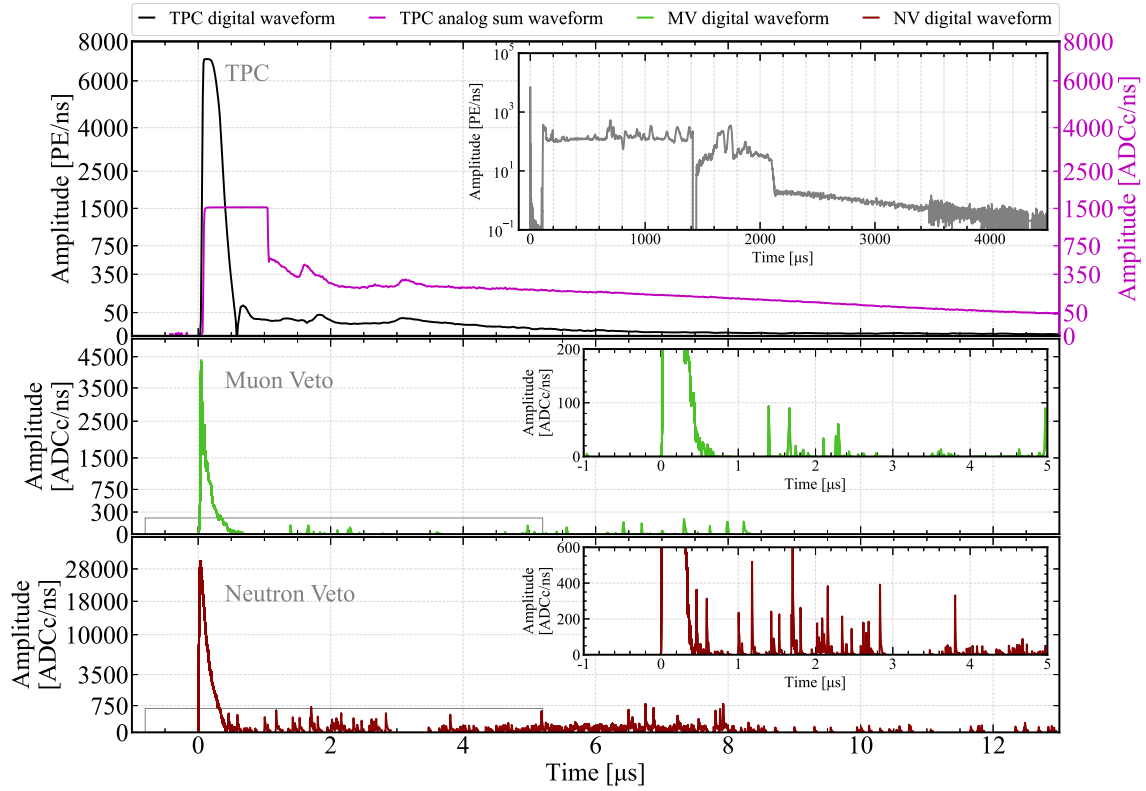


Figure 6. An illustration of a muon event recorded by all three DAQ subsystems. The top panel shows a zoomed-in view of the beginning of the muon waveform that was recorded by the TPC digitizers (black). An analog sum waveform of the bottom PMT array that was recorded by the TPC’s acquisition monitor digitizer is also shown in the same panel (magenta). The inset in the top panel showcases the digital sum waveform from data acquired by all TPC PMTs and recorded by the digitizers. The waveform is composed of downsampled data with different time resolution, resulting in the apparent amplitude fluctuation in its tail. The drop at $\sim 1400 \mu\text{s}$ is caused by a baseline fluctuation. The middle panel shows the same muon event recorded by the muon veto, while the bottom panel shows the data recorded by the neutron veto. Insets in both middle and bottom panels show a zoomed-in view of the respective waveforms. All three waveforms were aligned based on the TPC signal (black).

time of 2.2 ms. As a consequence, the system did not have to work under persistent high data rates and live processing was always able to process the data faster than it could be collected. During SR0, the data rate never exceeded $\sim 50 \text{ MB/s}$ for extended periods of time. To quantify the performance of the eventbuilders in high data rate conditions, pre-SR0 commissioning data that were taken during a high rate $^{83\text{m}}\text{Kr}$ calibration are used.

Figure 7 shows bootstrax cumulative processing time for different datatypes. Here, raw-records is the lowest level datatype, followed by peaks and finally events. The results for higher level datatypes include the time to compute the lower level datatypes as well. There are some additional datatypes of intermediate data in between, which one can find in the straxen documentation [44], several of which were briefly discussed in section 4. The total processing time comprises the time of starting bootstrax for a given run, decompressing the redax data, processing the data until the specified datatypes, compressing and writing all of the processed data to disk.




Figure 7. Live processing time as a function of the raw data rate for several target datatypes. The raw-records datatype is the lowest level datatype, followed by peaks and finally events (for simplicity this is called events, even though the benchmarks were obtained for the event-basics datatype [44]). When a higher level datatype is computed, all the lower level datatypes are also produced, so processing event also includes raw-records and peaks. All SR0 science data in this figure are 50 MB/s (gray band). Pre-SR0 commissioning data were used for the high rate data points where the DAQ was operated with a fractional livetime mode (discussed in section 3.3) during a high rate $^{83\text{m}}\text{Kr}$ -calibration. For data rates $\lesssim 250$ MB/s the live processing keeps up with one eventbuilder (EB) server as the processing time is lower than the acquisition time for any datatype. For any data rate in this plot the points are below the break-even line of the three eventbuilders, meaning that live processing at the DAQ could keep up with the readout.

Figure 7 shows that for data rates below 250 MB/s, the processing time is shorter than the collection time and a single eventbuilder can manage the entire data stream regardless of the datatype considered. For higher rates, the processing up to the events or peaks datatype is not fast enough to keep the processing live as each chunk would be processed slightly later than it is acquired. At these data rates, the finite RAM of the servers and increased disk read/write operations prevent processing at the same rate as at lower data rates, since processing each new chunk on a separate core starts requiring more memory than available on the host. The break-even line for one eventbuilder in figure 7 lies around ~ 250 MB/s for events, ~ 400 MB/s for peaks, and ~ 550 MB/s for raw-records. Additionally, the work is divided among three eventbuilders (with two additional as backup) and the combined eventbuilders can keep up with much higher data rates.

The total rate of the neutron veto and muon veto subsystems was found to be relatively constant and amounts to 10 MB/s–20 MB/s and ~ 1 MB/s, respectively. It takes about 80 s to process a 1800 s run for the neutron veto data, and $\lesssim 40$ s for a 1800 s run of muon veto data.

7 Conclusion

The XENON collaboration has designed and commissioned the triggerless XENONnT DAQ. By forgoing a trigger and relying instead on fast software to handle the continuous data stream, all data exceeding the digitization threshold is written to disk. The TPC, muon veto and neutron veto subsystems that constitute the DAQ can be operated independently, or as one linked system sharing the same 50 MHz clock signal. The increased number of PMTs and the double digitization of the top PMT array leads to roughly three times the number of channels with respect to XENON1T for the TPC. While the triggered muon veto-subsystem remains virtually unchanged, the new neutron veto-subsystem was successfully built to enable tagging of neutron events, one of the main backgrounds for the XENONnT WIMP search. A 500 MHz sampling rate enables the required characterization of neutron signals.

The DAQ is able to operate at the highest data rates observed during the first science run of XENONnT (SR0) of ~ 500 MB/s with the potential to go higher. The deadtime fraction is as low as 3×10^{-4} for science data and $\lesssim 10\%$ for calibration data at high rates of < 350 MB/s.

Using online processing, high level data are directly made available to monitor the detector. This enables analysts to have immediate feedback on changing detector conditions with fully processed data. The online processing is able to handle all data rates observed during SR0, where each of the three dedicated servers is able to process the data with a rate of up to ~ 250 MB/s. The maximum observed data rates during SR0 were limited by the low drift-field conditions. However, the DAQ was designed and is capable of dealing with data throughput rates greater than 750 MB/s.

During commissioning and SR0, the XENONnT DAQ has collected more than 2 PB of both science and calibration data, and it will continue to operate in subsequent science runs. The successful operation of the XENONnT DAQ and the implementation of the triggerless readout paradigm provides a solid basis for the development of DAQ systems for the next generation of liquid xenon dark matter experiments [19, 20].

Acknowledgments

We gratefully acknowledge support from the National Science Foundation, Swiss National Science Foundation, German Ministry for Education and Research, Max Planck Gesellschaft, Deutsche Forschungsgemeinschaft, Helmholtz Association, Dutch Research Council (NWO), Weizmann Institute of Science, Israeli Science Foundation, Binational Science Foundation, Fundacao para a Ciencia e a Tecnologia, Région des Pays de la Loire, Knut and Alice Wallenberg Foundation, Kavli Foundation, JSPS Kakenhi and JST FOREST Program in Japan, Tsinghua University Initiative Scientific Research Program and Istituto Nazionale di Fisica Nucleare. This project has received funding/support from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 860881-HIDDeN. Data processing is performed using infrastructures from the Open Science Grid, the European Grid Initiative and the Dutch national e-infrastructure with the support of SURF Cooperative. We are grateful to Laboratori Nazionali del Gran Sasso for hosting and supporting the XENON project.

References

- [1] XENON collaboration, *The XENONIT Dark Matter Experiment*, *Eur. Phys. J. C* **77** (2017) 881 [[arXiv:1708.07051](#)].
- [2] LZ collaboration, *The LUX-ZEPLIN (LZ) Experiment*, *Nucl. Instrum. Meth. A* **953** (2020) 163047 [[arXiv:1910.09124](#)].
- [3] PANDAX-4T collaboration, *Dark Matter Search Results from the PandaX-4T Commissioning Run*, *Phys. Rev. Lett.* **127** (2021) 261802 [[arXiv:2107.13438](#)].
- [4] nEXO collaboration, *nEXO: neutrinoless double beta decay search beyond 10^{28} year half-life sensitivity*, *J. Phys. G* **49** (2022) 015104 [[arXiv:2106.16243](#)].
- [5] DARKSIDE collaboration, *DarkSide-50 532-day Dark Matter Search with Low-Radioactivity Argon*, *Phys. Rev. D* **98** (2018) 102006 [[arXiv:1802.07198](#)].
- [6] XENON collaboration, *Dark Matter Search Results from a One Ton-Year Exposure of XENONIT*, *Phys. Rev. Lett.* **121** (2018) 111302 [[arXiv:1805.12562](#)].
- [7] XENON collaboration, *Projected WIMP sensitivity of the XENONnT dark matter experiment*, *JCAP* **11** (2020) 031 [[arXiv:2007.08796](#)].
- [8] XENON collaboration, *Search for New Physics in Electronic Recoil Data from XENONnT*, *Phys. Rev. Lett.* **129** (2022) 161805 [[arXiv:2207.11330](#)].
- [9] CAEN, *User Manual, 724 DPP-DAW Registers Description*, <https://www.caen.it/download/?filter=DPP-DAW>, (July 2019).
- [10] CAEN, <https://www.caen.it/>.
- [11] XENON collaboration, *The XENONIT Data Acquisition System*, *2019 JINST* **14** P07016 [[arXiv:1906.00819](#)].
- [12] CAEN, *DPP-DAW Firmware Release Notes*, <https://www.caen.it/?downloadfile=6349>, (November 2021).
- [13] XENONIT collaboration, *Conceptual design and simulation of a water Cherenkov muon veto for the XENONIT experiment*, *2014 JINST* **9** P11006 [[arXiv:1406.2374](#)].
- [14] J. Wulf, *Direct Dark Matter Search with XENONIT and Developments for Multi-Ton Liquid Xenon Detectors*, Ph.D. thesis, Universität Zürich, Swiss (2018).
- [15] MongoDB documentation, <https://www.mongodb.com/docs/atlas/>.
- [16] XENON collaboration, *Emission of single and few electrons in XENONIT and limits on light dark matter*, *Phys. Rev. D* **106** (2022) 022001 [[arXiv:2112.12116](#)].
- [17] XENON collaboration, *Search for Light Dark Matter Interactions Enhanced by the Migdal Effect or Bremsstrahlung in XENONIT*, *Phys. Rev. Lett.* **123** (2019) 241803 [[arXiv:1907.12771](#)].
- [18] XENON collaboration, *Light Dark Matter Search with Ionization Signals in XENONIT*, *Phys. Rev. Lett.* **123** (2019) 251801 [[arXiv:1907.11485](#)].
- [19] DARWIN collaboration, *DARWIN: towards the ultimate dark matter detector*, *JCAP* **11** (2016) 017 [[arXiv:1606.07001](#)].
- [20] J. Aalbers et al., *A next-generation liquid xenon observatory for dark matter and neutrino physics*, *J. Phys. G* **50** (2023) 013001 [[arXiv:2203.02309](#)].

- [21] J. Aalbers et al., *PAX data processor*, <https://zenodo.org/record/4290785> [DOI:10.5281/ZENODO.4290785].
- [22] J. Aalbers et al., *AxFoundation/strax: v1.1.5*, <https://zenodo.org/record/5833114> [DOI:10.5281/ZENODO.5833114].
- [23] J. Angevaere et al., *XENONnT/straxen: v1.2.4*, <https://zenodo.org/record/5834311> [DOI:10.5281/ZENODO.5834311].
- [24] C.R. Harris et al., *Array programming with NumPy*, *Nature* **585** (2020) 357 [arXiv:2006.10256].
- [25] P. Virtanen et al., *SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python*, *Nature Meth.* **17** (2020) 261 [arXiv:1907.10121].
- [26] S.K. Lam, A. Pitrou and S. Seibert, *Numba: A LLVM-based python JIT compiler*, in the proceedings of the *Second Workshop on the LLVM Compiler Infrastructure in HPC*, 2015, p. 1–6.
- [27] A. Brown, *Search for Elastic and Inelastic Dark Matter Interactions in XENON1T and Light Detection for XENONnT*, Ph.D. thesis, Universität Zürich, Swiss (2020).
- [28] CAEN, *User Manual, A1535/AG535 — 24 CH 3.5KV/ 3 mA High Voltage Boards*, <https://www.caen.it/download/?filter=A1535>, (October 2015).
- [29] CAEN, *User Manual, A1536 AG536 Family Power Supply Boards*, <https://www.caen.it/download/?filter=A1536>, (February 2015).
- [30] CAEN, *User Manual, A7030-AG7030 3kV/ 1mA (1.5W) HV Boards*, <https://www.caen.it/download/?filter=A7030>, (September 2019).
- [31] CAEN, *User Manual, A7435 AG7435 24 Channel 3.5 kV/3.5 mA (9W)*, <https://www.caen.it/download/?filter=A7435>, (September 2022).
- [32] CAEN, *User Manual, SY4527/SY4527LC Universal Multichannel Power Supply Systems 16/10 slots Mainframes*, <https://www.caen.it/download/?filter=SY4527>, (August 2022).
- [33] CAEN, *User Manual, DT4700 — Clock Generator*, <https://www.caen.it/download/?filter=DT4700>, (January 2020).
- [34] M. De Deo et al., *Accurate GPS-based timestamp facility for Gran Sasso National Laboratory*, *2019 JINST* **14** P04001 [arXiv:1903.00233].
- [35] CAEN, *User Manual, V1724 8 Channel 14bit — 100MS/s Digitizer*, <https://www.caen.it/download/?filter=V1724>, (December 2019).
- [36] CAEN, *User Manual, V2718/VX2718 — VME/VME64X — PCI Optical Link Bridges*, <https://www.caen.it/download/?filter=V2718>, (July 2018).
- [37] CAEN, *User Manual, V1495 General Purpose VME Board*, <https://www.caen.it/download/?filter=V1495>, (August 2019).
- [38] CAEN, *User Manual, V976 — Quad 4 Fold AND/OR/MAJ, NIM-TTL TTL-NIM Translator, Fan-In Fan-Out*, <https://www.caen.it/download/?filter=V976>, (November 2015).
- [39] CAEN, *User Manual, V1730/VX1730 & V1725/VX1725*, <https://www.caen.it/download/?filter=V1730>, (March 2023).
- [40] SkuTek Instrumentation, *Digital Pulse Processor for Nuclear Physics*, <https://www.skutek.com/ddc10.htm>.
- [41] Ceph documentation, <https://docs.ceph.com/en/quincy/start/intro/>.

- [42] CAEN, A3818 PCI Express optical link, <https://www.caen.it/download/?filter=A3818>, (January 2020).
- [43] D. Masson et al., *AxFoundation/redax: Version 2.3.0*, <https://zenodo.org/record/5882717> [DOI: 10.5281/ZENODO.5882717].
- [44] Straxen documentation, <https://straxen.readthedocs.io/>.
- [45] V.C. Antochi et al., *Improved quality tests of R11410-21 photomultiplier tubes for the XENONnT experiment*, 2021 *JINST* **16** P08033 [arXiv:2104.15051].
- [46] XENON collaboration, *Application and modeling of an online distillation method to reduce krypton and argon in XENONIT*, *PTEP* **2022** (2022) 053H01 [arXiv:2112.12231].
- [47] XENON collaboration, *XENONIT Dark Matter Data Analysis: Signal Reconstruction, Calibration and Event Selection*, *Phys. Rev. D* **100** (2019) 052014 [arXiv:1906.04717].
- [48] Slack documentation, <https://api.slack.com/docs>.
- [49] SNEWS collaboration, *SNEWS 2.0: a next-generation supernova early warning system for multi-messenger astronomy*, *New J. Phys.* **23** (2021) 031201 [arXiv:2011.00035].
- [50] NodeJS documentation, <https://nodejs.dev/en/api/v19/documentation/>.