

Jornadas tecnológicas del IFIC 2020

Infraestructura Computacional en



Autor: José V. Carrión

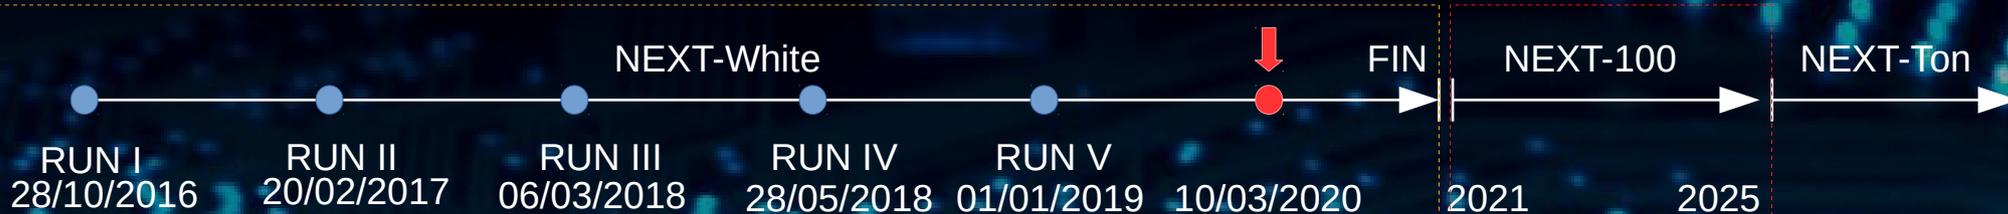


Trabajo realizado en el IFIC

Personal Técnico de Apoyo (PTA 2016) : duración 2018 a 2020

- **Reconstrucción de superficies 3D en mecanizados para la detección de anomalías. (J.J Hernández Rey)**
- **Sabor y Origen de la Materia (SOM) / 2011-2020:**
 - **Tareas:** Administración del cluster SOM, instalación y configuración de software científico, soporte IT a los miembros del grupo SOM, administración servidor web y disco (40TB).
 - Recursos SOM:
 - Cluster SOM (5 nodos, 144c/288t, 2.1 TB RAM, 1gb/E, 1 x GPU v100 + 1 x GPU T4)
 - Cluster Theory (48 nodos, 120c/240t, 307GB RAM)
 - **Total → CPU: 264c/528t + RAM: 2.4 TB + HD: 40 TB**
- **Neutrino Experiment with a Xenon TPC (NEXT) / 2015-2020:**
 - **Tareas:** Infraestructura computacional, almacenamiento y red del experimento NEXT. Administración clusters NEXT, instalación software científico. Soporte IT a usuarios NEXT, compras de equipos de supercomputación, administración servicios web y de control de versiones, etc.
 - Recursos NEXT:
 - LSC: CPU: 55c/110t - HD: 300TB - RAM: 450GB
 - Cluster DEMO ++ : CPU: 76c/152t - HD: 146TB - RAM: 172GB
 - Cluster neutrinos1: CPU: 64c/128t - HD: 83TB - RAM: 337GB
 - Equipos externos: CPU: 68c/200t - HD: 26TB - RAM: 1385GB
 - **Total → CPU: 263c/590t + RAM: 2.3 TB + HD: 555 TB**

Producción de datos en el detector NEXT-White (5 kg ^{136}Xe enriched)



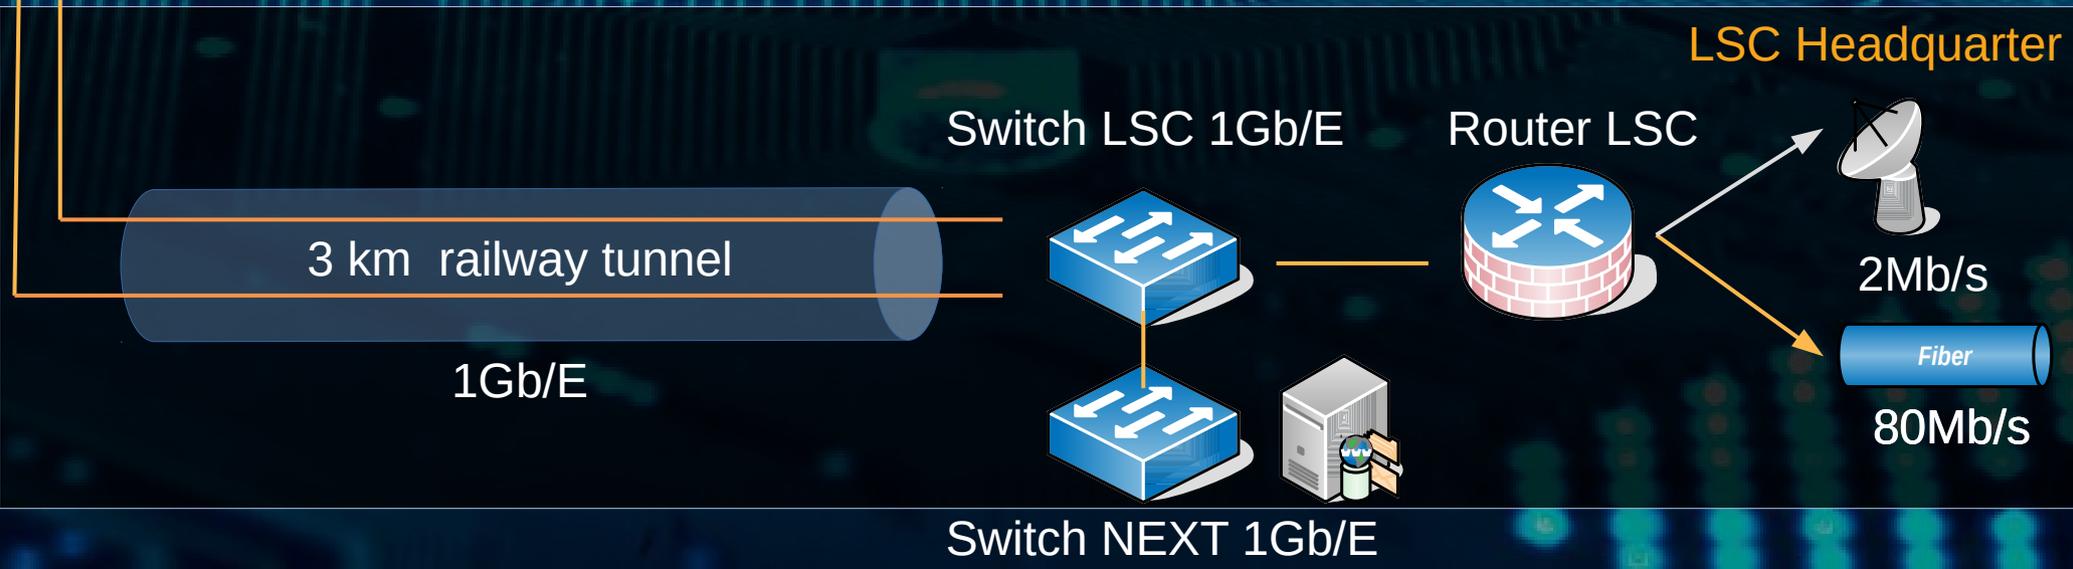
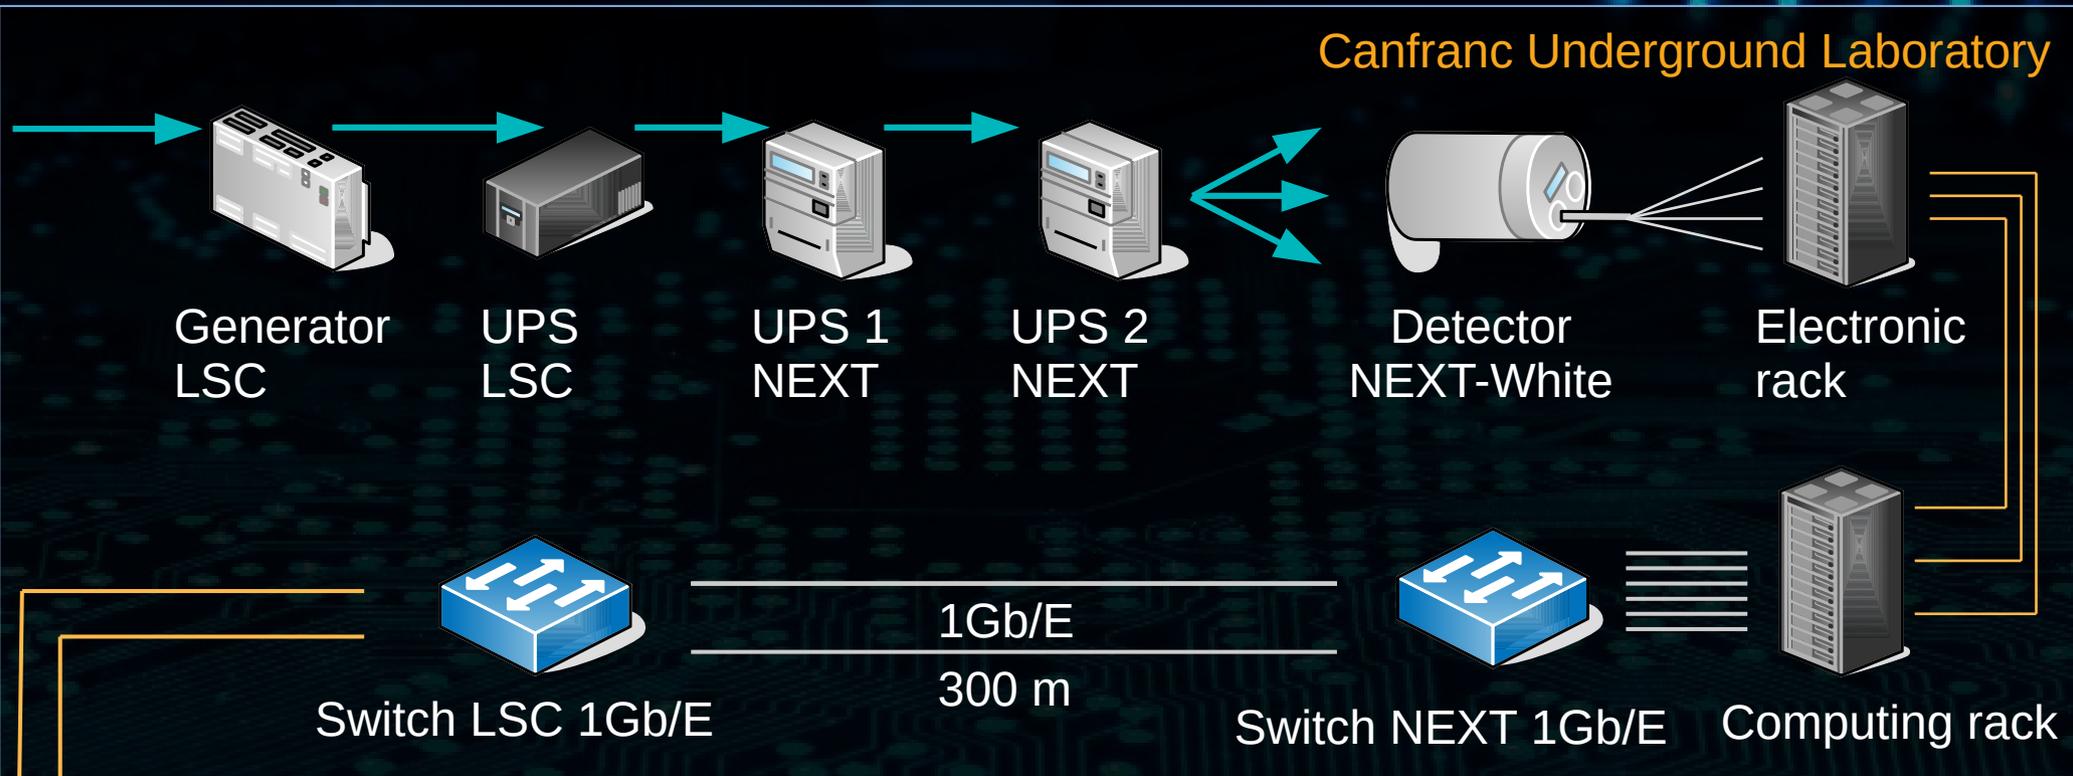
PRODUCED DATA			
Stage	Total Events	Produced Data	Duration
Run 1	14.796.565	51 TB	372 h
Run 2	67.939.259	421 TB	5.287 h
Run 3	13.850.468	78 TB	454 h
Run 4	485.500.724	1.403 TB	4.619 h
Run 5	775.467.577	1.886 TB	8.934 h
Total	1.357.554.593	3.839 TB	19.667 h

Datos procesados, analizados y almacenados en disco y cinta:

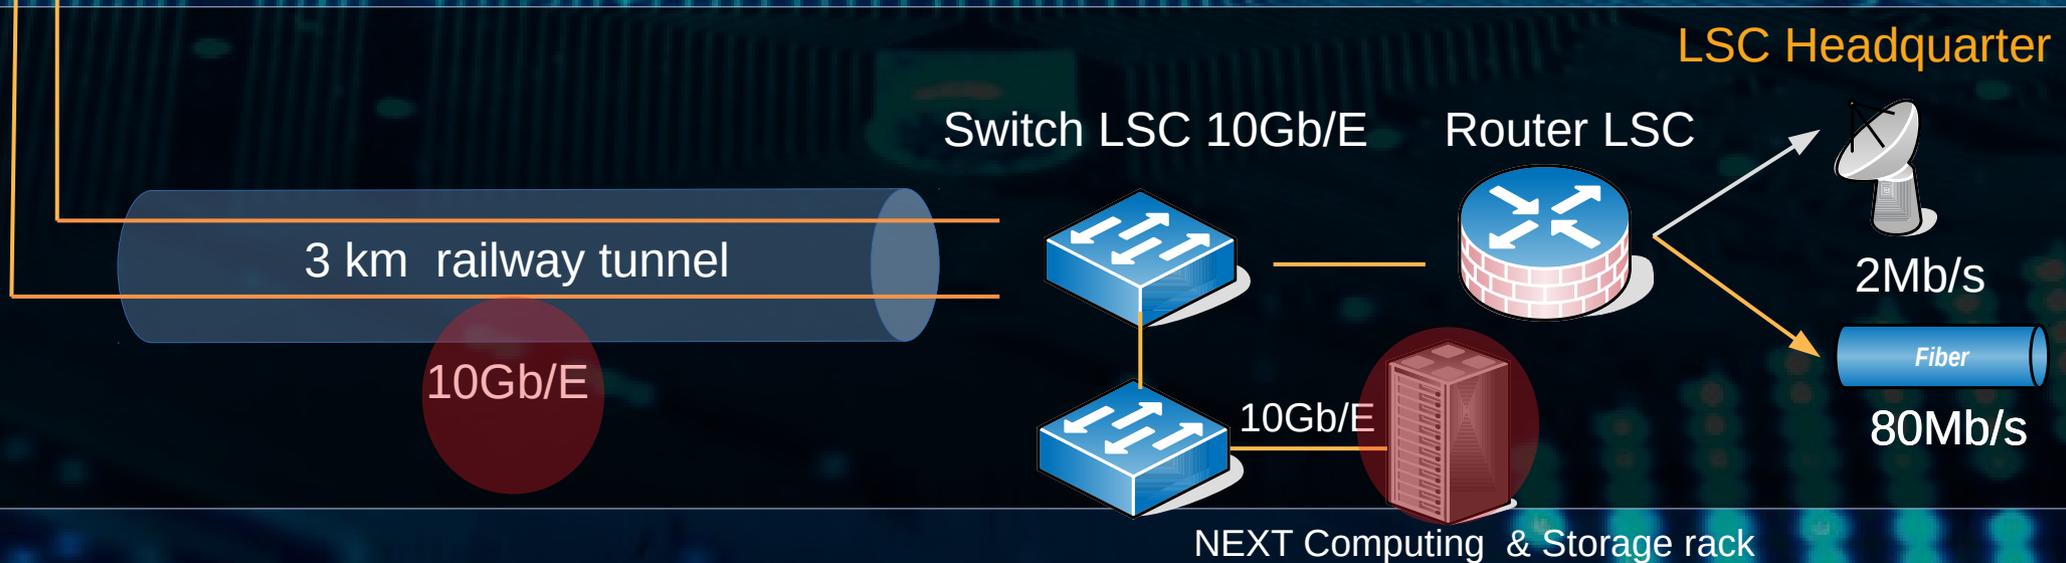
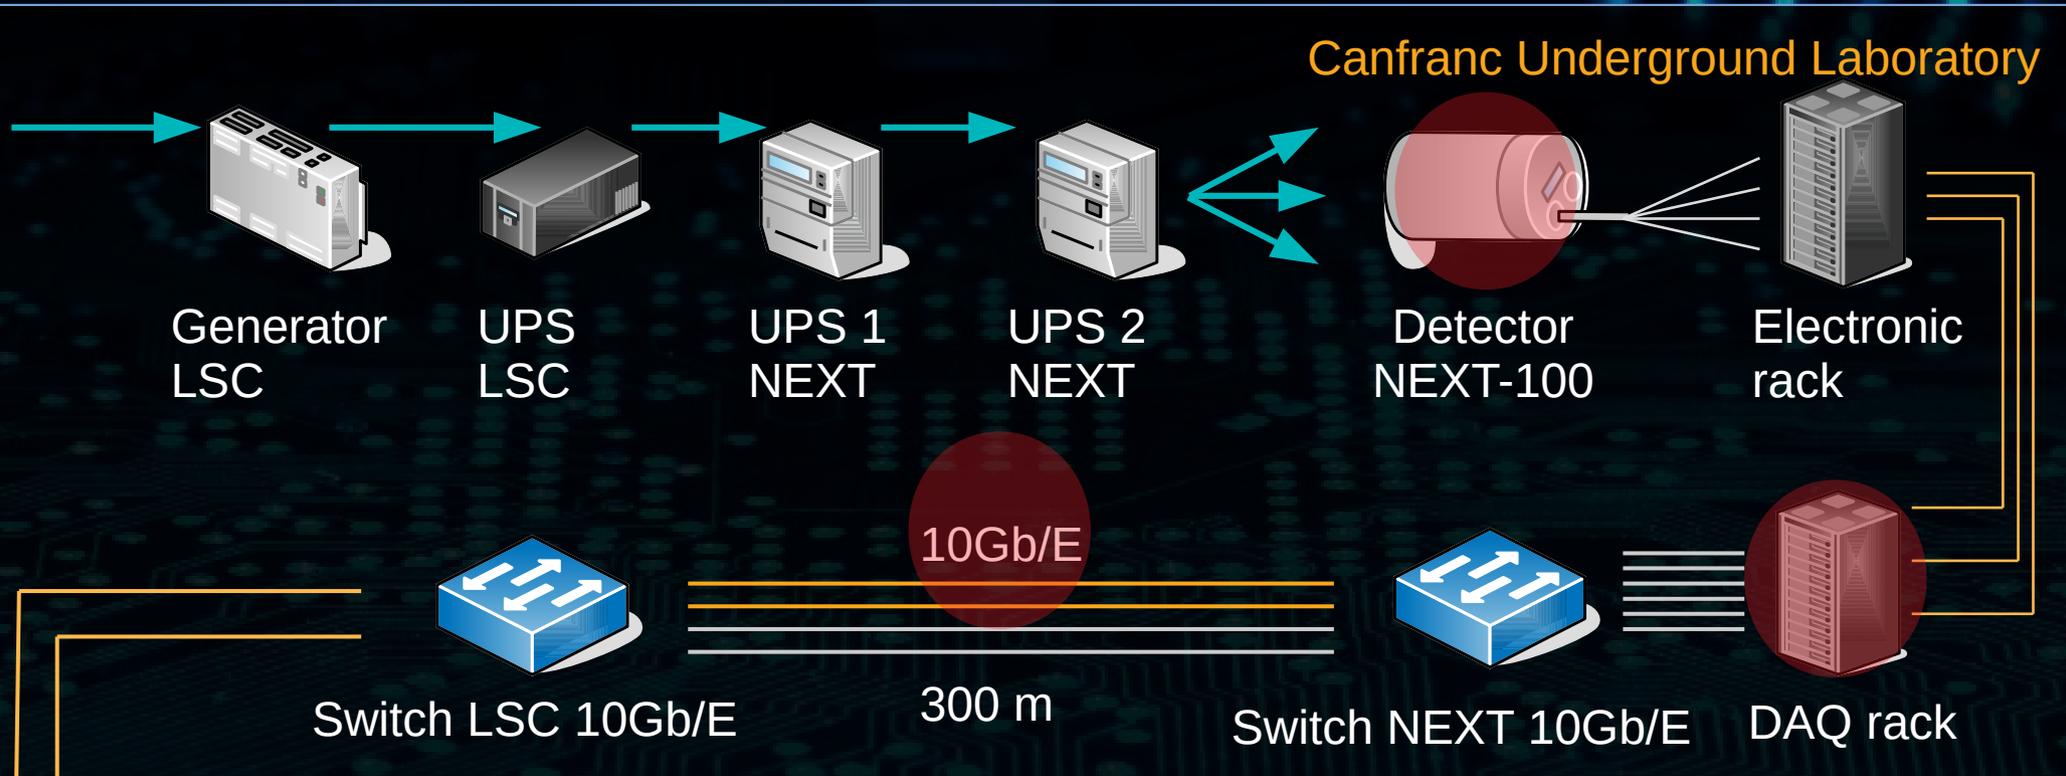
225 TB : Almacenados y accesibles.

457 TB : Archivados en cinta LTO7 (LTFS).

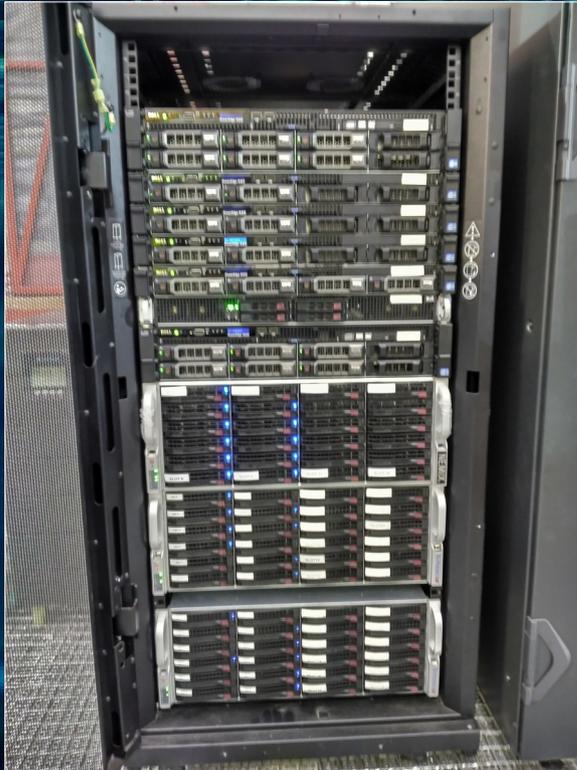
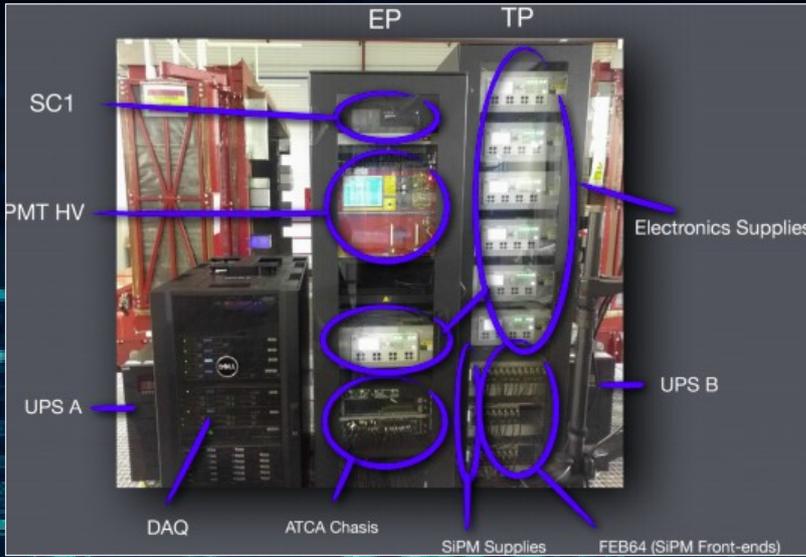
Infraestructura NEXT-White (5kg ^{136}Xe enriched) en el LSC



Infraestructura NEXT-100 (100kg ^{136}Xe enriched) en el LSC



Experimento NEXT en el LSC



DAQ I: NEXT-White a NEXT-100

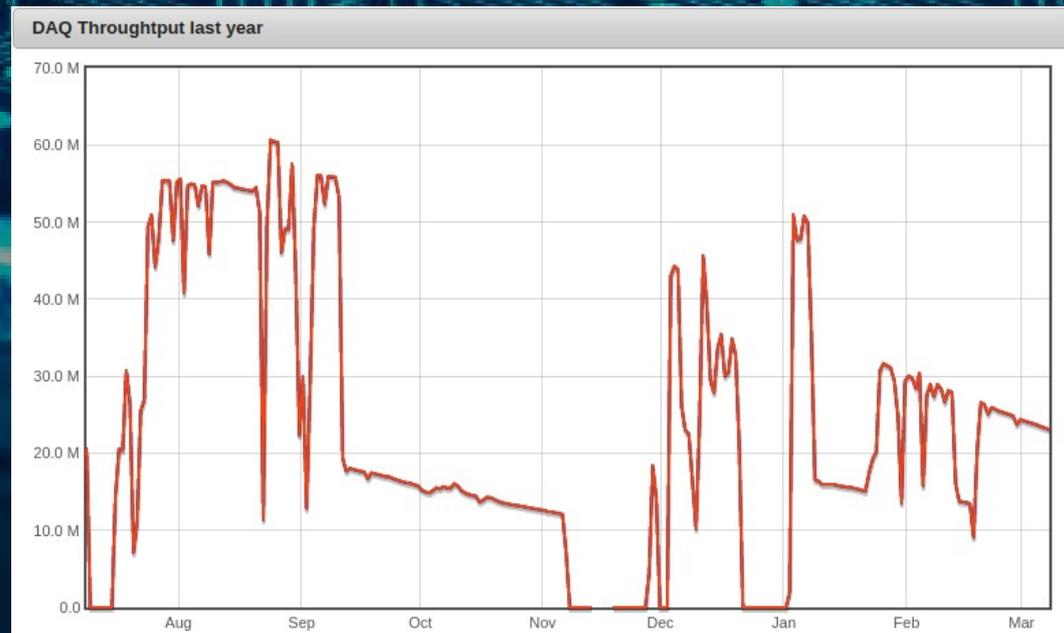
NEXT-White (5Kg enriched ^{136}Xe)

- Calibración: **140 MB/s**
 - SiPM → 2.000
 - PMTs → 12
 - Data Links → **12**
- Software **DATE** (CERN)
- RUN V → Zero suppression + compression **hardware**
 - Reducción **80%** bytes/evento
 - Mayor trigger hertz
 - Mayor número de eventos/s
- Automatización DAQ:
 - Parada/reinicio automático de DATE
 - Registro automático de eventos en ELOG (CERN) y NEXT FileCatalog.

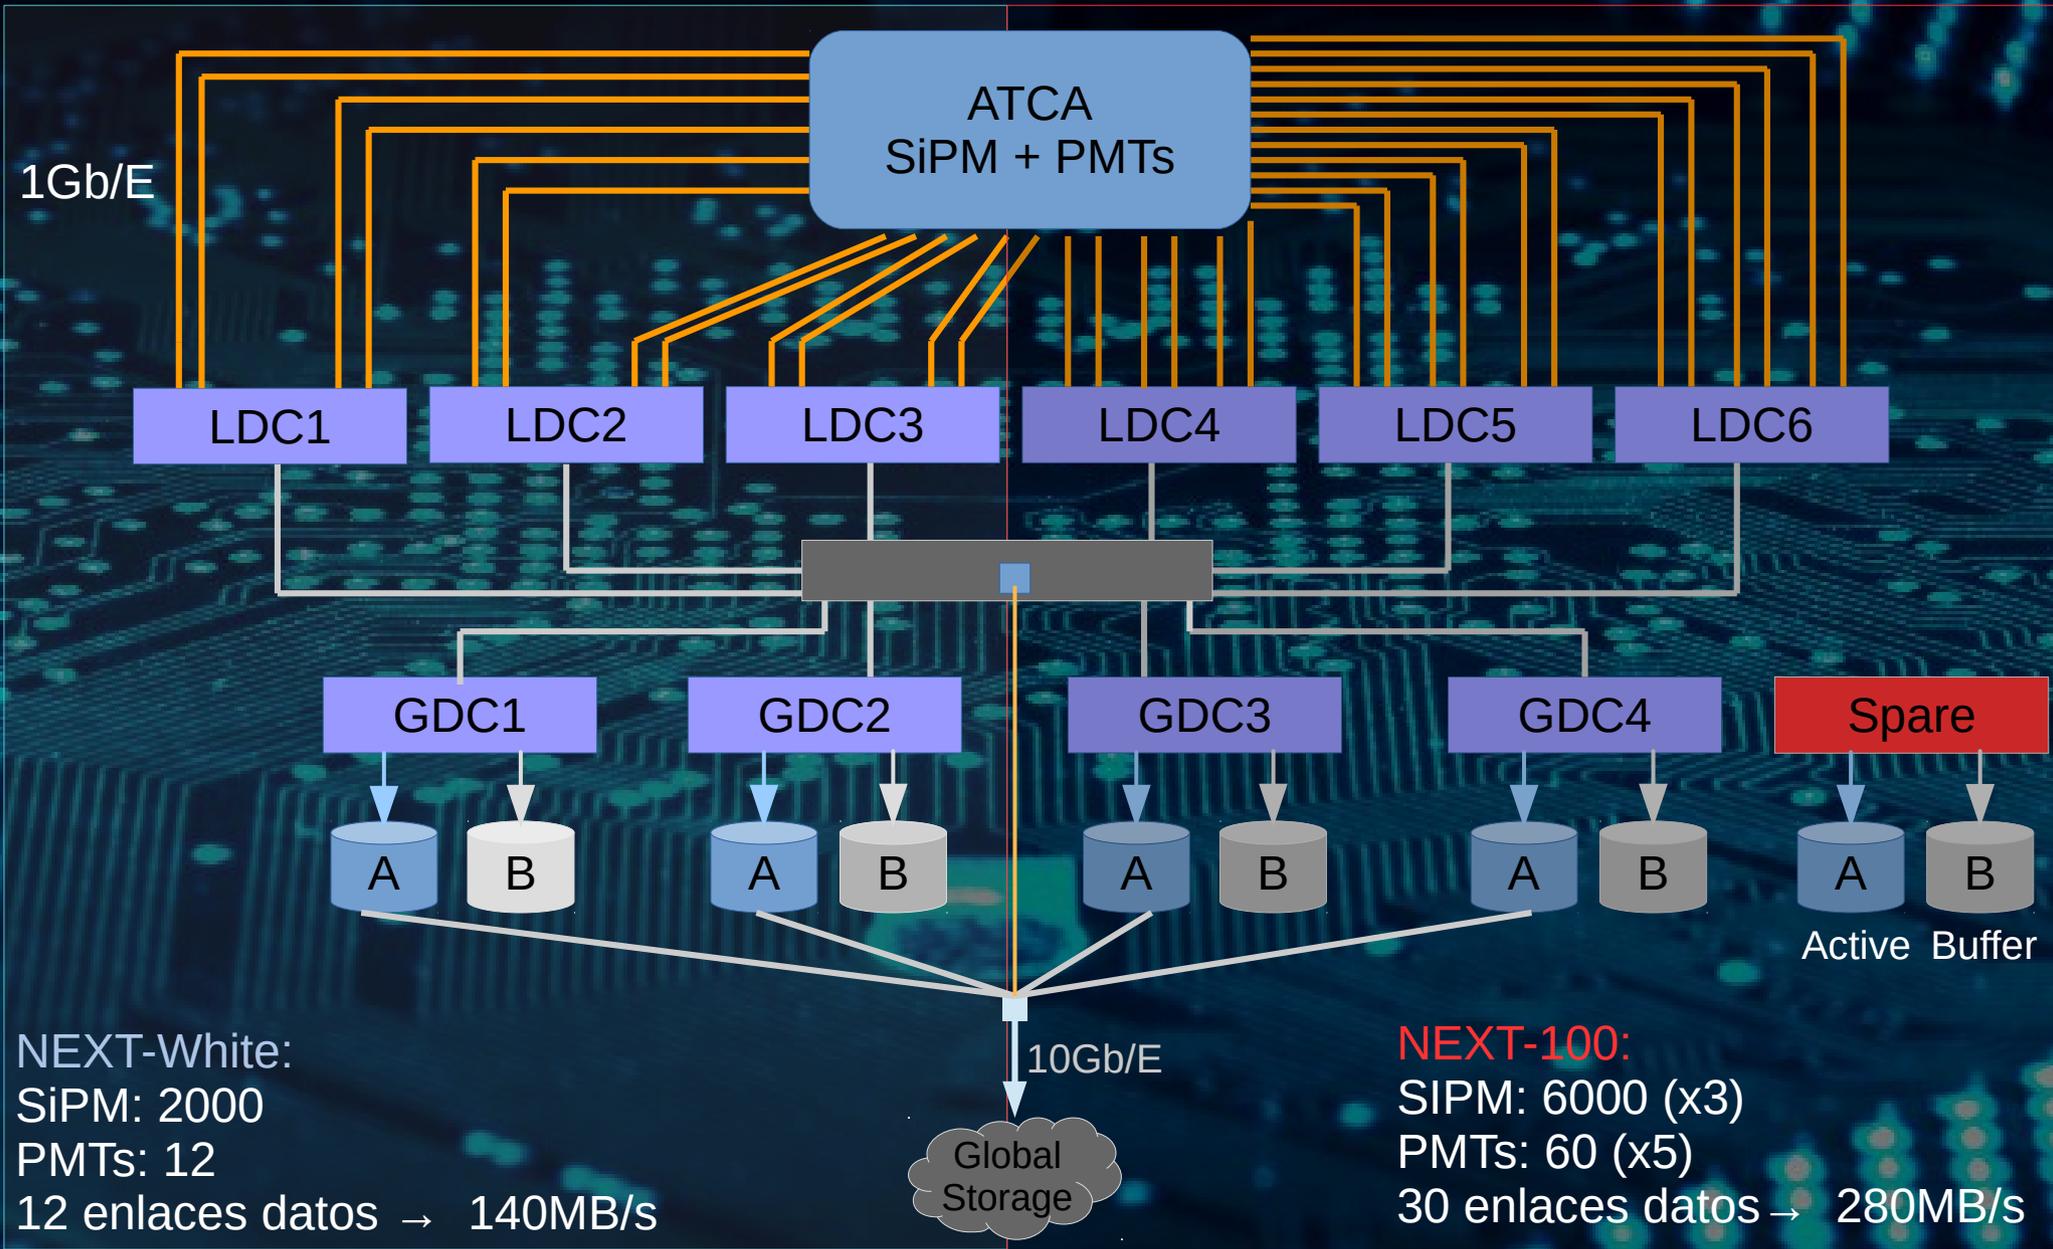


NEXT-100 (100Kg ^{136}Xe enriched)

- Calibración: **280 MB/s**
 - SiPM → 6.000
 - PMTs → 60
 - Data Links → **30**
- Software **DATE** (CERN) vs **RCDAQ** (PHENIX)

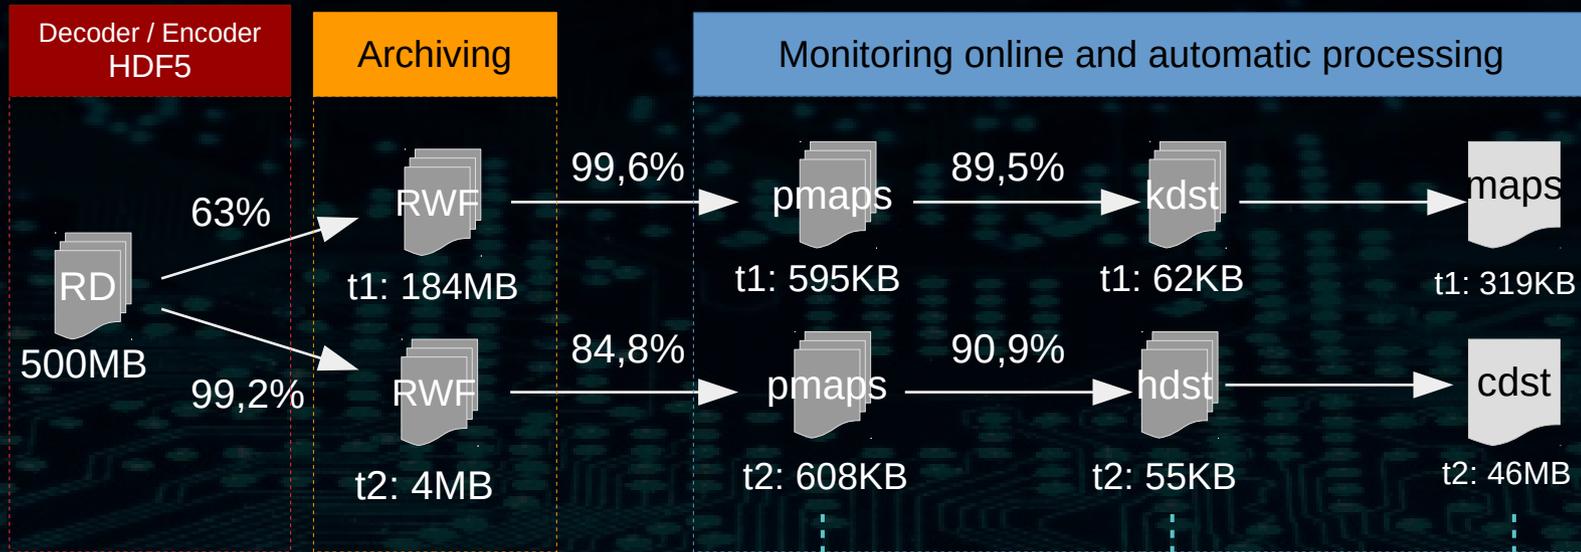


DAQ II: NEXT-NEW a NEXT-100



Análisis de datos en NEXT-White

Canfranc Underground Laboratory



Temporal Raw Data

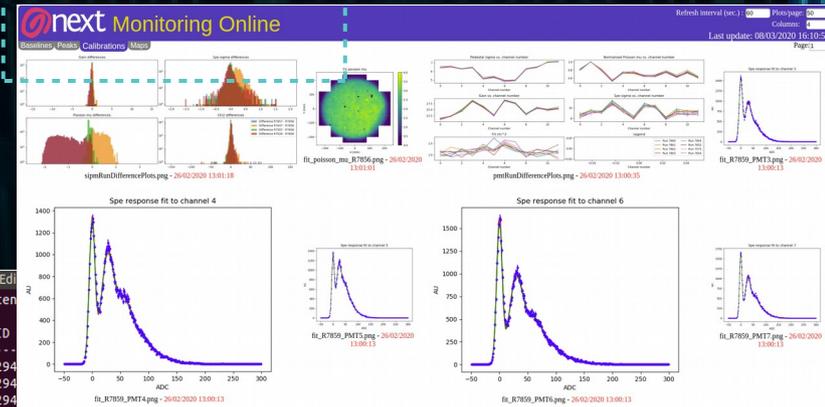


Reprocessing

Analysis



- Desarrollo propio NEXT: Invisibles Cities (Python)
- Monitorización online de la producción de datos.
- Sistema de procesado automático (batch).
- Control de versiones: Github + Travis CI , Git-LFS, Citlab

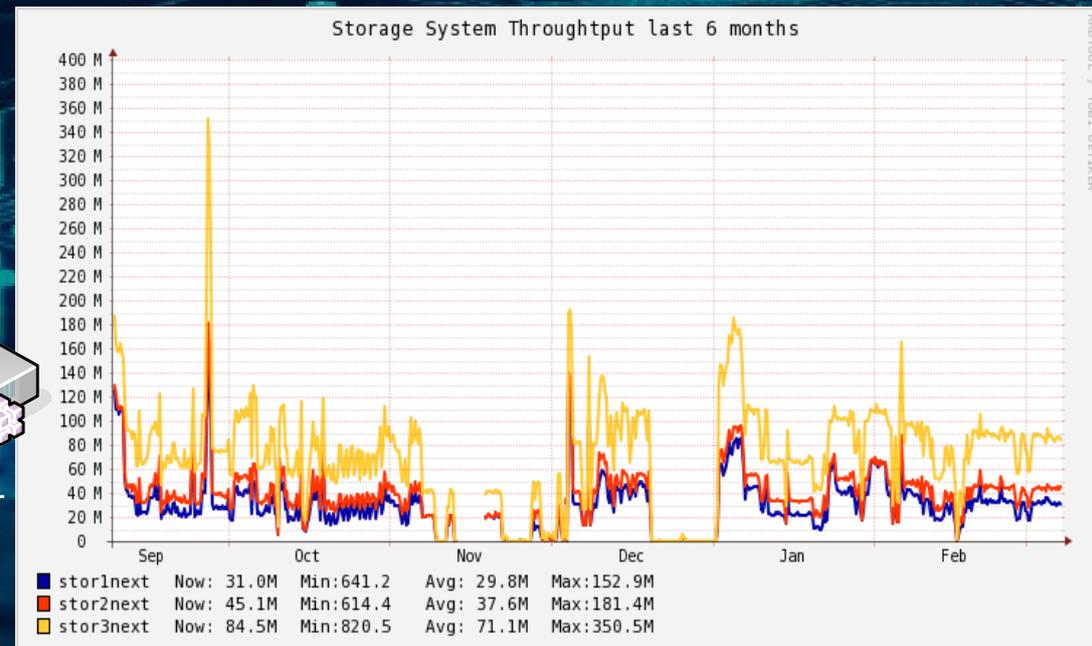
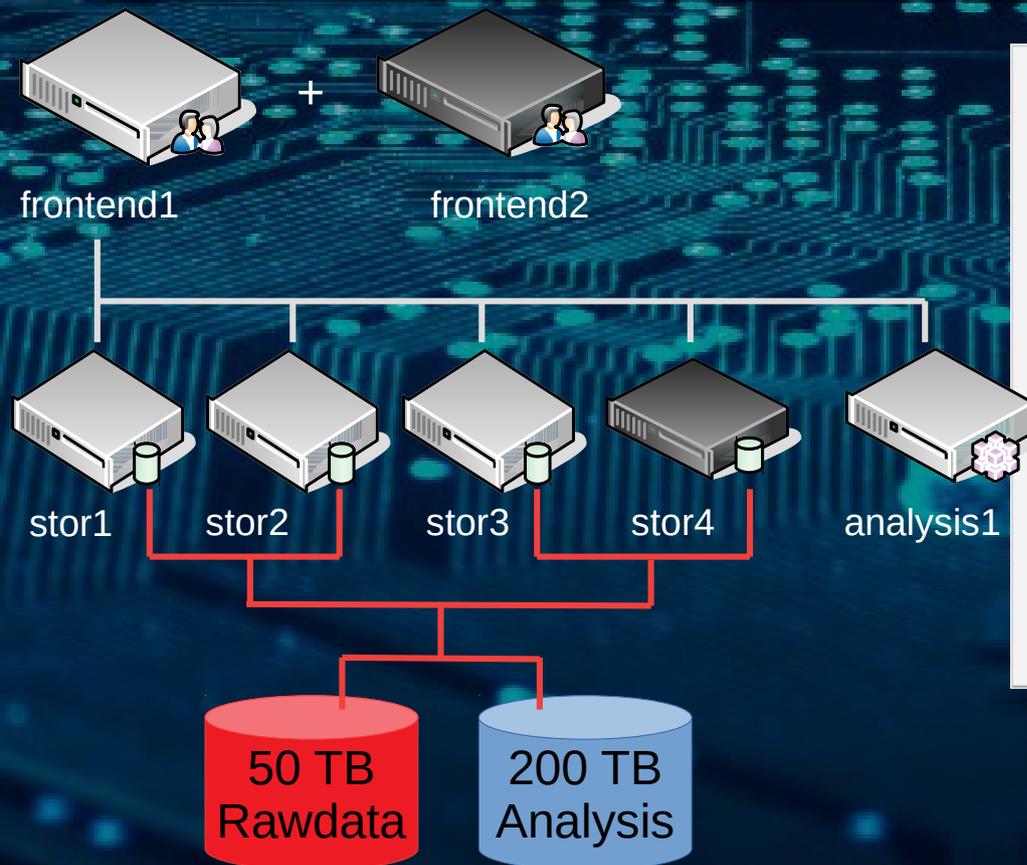


```

File Edit View Help
frontend
Job ID
-----
3360294
3360294
3360294
33602956.frontendInext shifter short 7885_irene_3212 2481 1 1 1gb 24:00:00 R 00:00:50
33602958.frontendInext shifter short 7885_irene_3213 2590 1 1 1gb 24:00:00 R 00:00:46
33602960.frontendInext shifter short 7885_irene_3214 2638 1 1 1gb 24:00:00 R 00:00:43
33602961.frontendInext shifter short 7885_irene_3215 47299 1 1 1gb 24:00:00 R 00:00:41
33602962.frontendInext shifter short 7885_irene_3216 47325 1 1 1gb 24:00:00 R 00:00:39
33602963.frontendInext shifter short 7885_irene_3217 2688 1 1 1gb 24:00:00 R 00:00:36
33602964.frontendInext shifter short 7885_irene_3218 2714 1 1 1gb 24:00:00 R 00:00:35
33602965.frontendInext shifter short 7885_irene_3219 47359 1 1 1gb 24:00:00 R 00:00:32
33602966.frontendInext shifter short 7885_irene_3220 47385 1 1 1gb 24:00:00 R 00:00:30
33602967.frontendInext shifter short 7885_irene_3222 47412 1 1 1gb 24:00:00 R 00:00:27
33602968.frontendInext shifter short DECODER_7885_323 2782 1 1 1gb 24:00:00 R 00:00:16
33602969.frontendInext shifter short DECODER_7885_323 47440 1 1 1gb 24:00:00 R 00:00:16
33602970.frontendInext shifter short DECODER_7885_323 2823 1 1 1gb 24:00:00 R 00:00:14
33602971.frontendInext shifter short OLIVIA_WF_7885_3 47466 1 1 1gb 24:00:00 R 00:00:05
33602972.frontendInext shifter short OLIVIA_WF_7885_3 47482 1 1 1gb 24:00:00 R 00:00:05
33602973.frontendInext shifter short OLIVIA_WF_7885_3 47516 1 1 1gb 24:00:00 R 00:00:03
[root@frontendInext ~]#
    
```

Computación y Almacenamiento: NEXT-White a NEXT-100

- Sistema de cómputo: 4 nodos (55 cores / 110 threads) + front-end
- Sistema de almacenamiento: GlusterFS (3 nodos, 300TB, 84 HD)
 - Ancho de banda agregado: 683 MB/s.
 - Dos volúmenes de datos optimizados para ficheros medianos y pequeños.
 - No hay tolerancia a fallos en el DFS (300TB) → backup de RWF en cinta.
- NEXT-100: 1 nodo de almacenamiento más → redundancia a datos 1+0 → tolerancia a fallos → mayor disponibilidad
- NEXT-100: GlusterFS o Ceph



Herramientas de monitorización remota, desarrollo de software y sistema de información NEXT

- Monitorización y gestión remota: IPMI-iKVM, VNC, Nagios, Ganglia.
- Sistema de documentación: docDB
- Desarrollo: Python, TensorFlow, PyTorch, Keras
- Control de versiones: GitHub y GitLab, Travis CI, Git-LFS.

The image displays a comprehensive monitoring dashboard for the 'next' system. The top section features a 'Current RUN Status' panel with details like Run Number (5375), Start Time (12/03/2018 18:28:06), and Duration (68 seconds). It also includes 'RUN Conditions' such as Source (51 No physics) and Hardware Conditions (CPU 41.00, Mem 20.14, etc.).

Below this is a 'Current Trigger/Event Rate' section showing an average trigger rate of 6,790 and a total of 482 events. The 'Current Storage Status' section indicates 5.78 GB of space exceeded and 64.21 MB/s of recording speed.

The central part of the dashboard is dominated by a 'Service Status Details For All Hosts' table, which lists various services (e.g., Current Load, System Status, RAID Disk) across multiple hosts (lbc2next, gbl2next, etc.) with their respective status (OK, Warning, Critical) and last check times.

The bottom half of the image shows a grid of network and system performance graphs, including 'Lbc2next Network last hour', 'gbl2next Network last hour', and 'Aggregate bytes_in last hour'. These graphs plot metrics like network traffic and storage usage over time.

On the right side, there is an 'IPMI View' window for a Supermicro server, displaying real-time hardware data such as CPU temperature, fan speeds, and power consumption.

The footer of the dashboard includes the 'next' logo and version information: 'Ganglia Web Frontend version 3.6.2', 'Ganglia Web BackEnd (j version) Check for Updates', and 'Powered by Datto 1.1.1'.

Herramientas de registro y catalogación de runes en NEXT

- Registro de runes: ELOG
- Catalogación de runes y búsqueda de archivos: NEXT FILECATALOG.

next-new-check-list | next-new-production | next-new-calibration | next-new-test | next-new-commissioning | next-new-troubles

NEXT-NEW ELOG CALIBRATION Logged in as "next"

Run No: 7911 | Author: ATCA | Type: Run Conditions | Run Type: 23 - Kr-83 + Double beta No Source

DETECTOR CONDITIONS - RUN 7911

Pressure (bar)	Temperature (°C)	Cathode Voltage (volts)	Gate Voltage (volts)
1.00550REV1	20.11	3001.50	7699.95

JAVA CONDITIONS

GENERAL CONFIGURATION					
Mode of operation	Buffer Size	Pretigger	Number of Triggers	PMT Buffer Size	Run Type Code
Normal Mode 2	1600	800	0	1500	23 - Kr-83 + Double beta No Source

TRIGGER CONFIGURATION

GLOBAL PARAMETERS				INTERNAL TRIGGER 1				INTERNAL TRIGGER 2							
Trigger Mask	Trigger freq	Auto External Trigger	External Trigger	Double Trigger	Max. Time Trg A/B	Events Trg A	Events Trg B	CW size Trg A	CW size Trg B	Double Trigger	Max. Time Trg A/B	Events Trg A	Events Trg B	CW size Trg A	CW size Trg B
OFF	15	OFF	OFF	OFF	0	2	1	48	1	OFF	0	2	1	48	1

FEC DISCONNECTED CARD

FEC 0 disconnected	FEC 2 disconnected	FEC 4 disconnected	FEC 6 disconnected	FEC 8 disconnected	FEC 10 disconnected	FEC 12 disconnected
NO	NO	NO	NO	NO	YES	YES

next - NEXT FILE CATALOG logged in as: admin

Search OPTIONS

Show entries Run number:

Run Number	Source	Start Date	End Date	Pressure (bar)	Cath. Voltage (V)	Gate. Voltage (V)	Total Events	Total Size (GB)
7911	Kr-83 + Double beta No Source	10-03-2020 13:54:15	RUNNING	10,07	30,001,50	7,699,95	RUNNING	RUNNING

Filter by TYPE: Filter by NAME: Filter by VALUE:

BACKUP Archiving Location
 TRG Auto External Trigger ON/OFF OFF
 Ch1 PMT FEC 2 PMTs TRG A Baseline dev. 1 10
 Ch0 PMT FEC 2 PMTs TRG A Baseline dev. 1 10
 Ch0 PMT FEC 2 PMTs TRG A Baseline dev. 2 10
 Ch1 PMT FEC 2 PMTs TRG A Baseline dev. 2 10
 FE general Baseline Ref 50
 GEN Buffer Size 1600
 PMT FEC 2 Ch0 selected TRG YES
 PMT FEC 2 Ch1 selected TRG YES
 PMT FEC 2 Ch10 selected TRG YES
 PMT FEC 2 Ch11 selected TRG YES
 PMT FEC 10 Ch12 selected TRG YES

7910		00-00-0000 00:00:00	10-03-2020 13:51:10	0,00	0,00	0,00	12	0,00
7909	Kr-83 + Double beta No Source	10-03-2020 13:36:39	00-00-0000 00:00:00	10,07	30,001,60	7,700,01	0	0,00
7908	Kr-83 + Double beta No Source	10-03-2020 13:12:52	00-00-0000 00:00:00	10,06	3,28	1,00	0	0,00
7907	Kr-83 + Double beta No Source	10-03-2020 13:03:56	10-03-2020 13:05:36	10,06	3,27	1,04	262	0,75
7905	Kr-83 + Double beta No Source	10-03-2020 13:00:24	10-03-2020 13:00:41	10,06	3,13	0,91	17	0,01
7904	Kr-83 + Double beta No Source	10-03-2020 12:57:12	00-00-0000 00:00:00	10,06	3,10	1,15	0	0,00
7903	Kr-83 + Double beta No Source	10-03-2020 12:54:54	00-00-0000 00:00:00	10,06	3,10	1,15	0	0,00
7902	Kr-83 + Double beta No Source	10-03-2020 12:41:56	00-00-0000 00:00:00	10,06	3,18	1,10	0	0,00
7901	Kr-83 + Double beta No Source	10-03-2020 12:38:48	00-00-0000 00:00:00	10,06	3,18	1,10	0	0,00
7900	Kr-83 + Double beta No Source	10-03-2020 12:35:27	00-00-0000 00:00:00	10,06	3,13	1,00	0	0,00

Run Number Source Start Date End Date Pressure (bar) Cath. Voltage (V) Gate. Voltage (V) Total Events: 2.002.417 Total Size: 5.999 GB

Showing 1 to 25 of 2,202 entries (filtered from 25 total entries) Previous ... Next

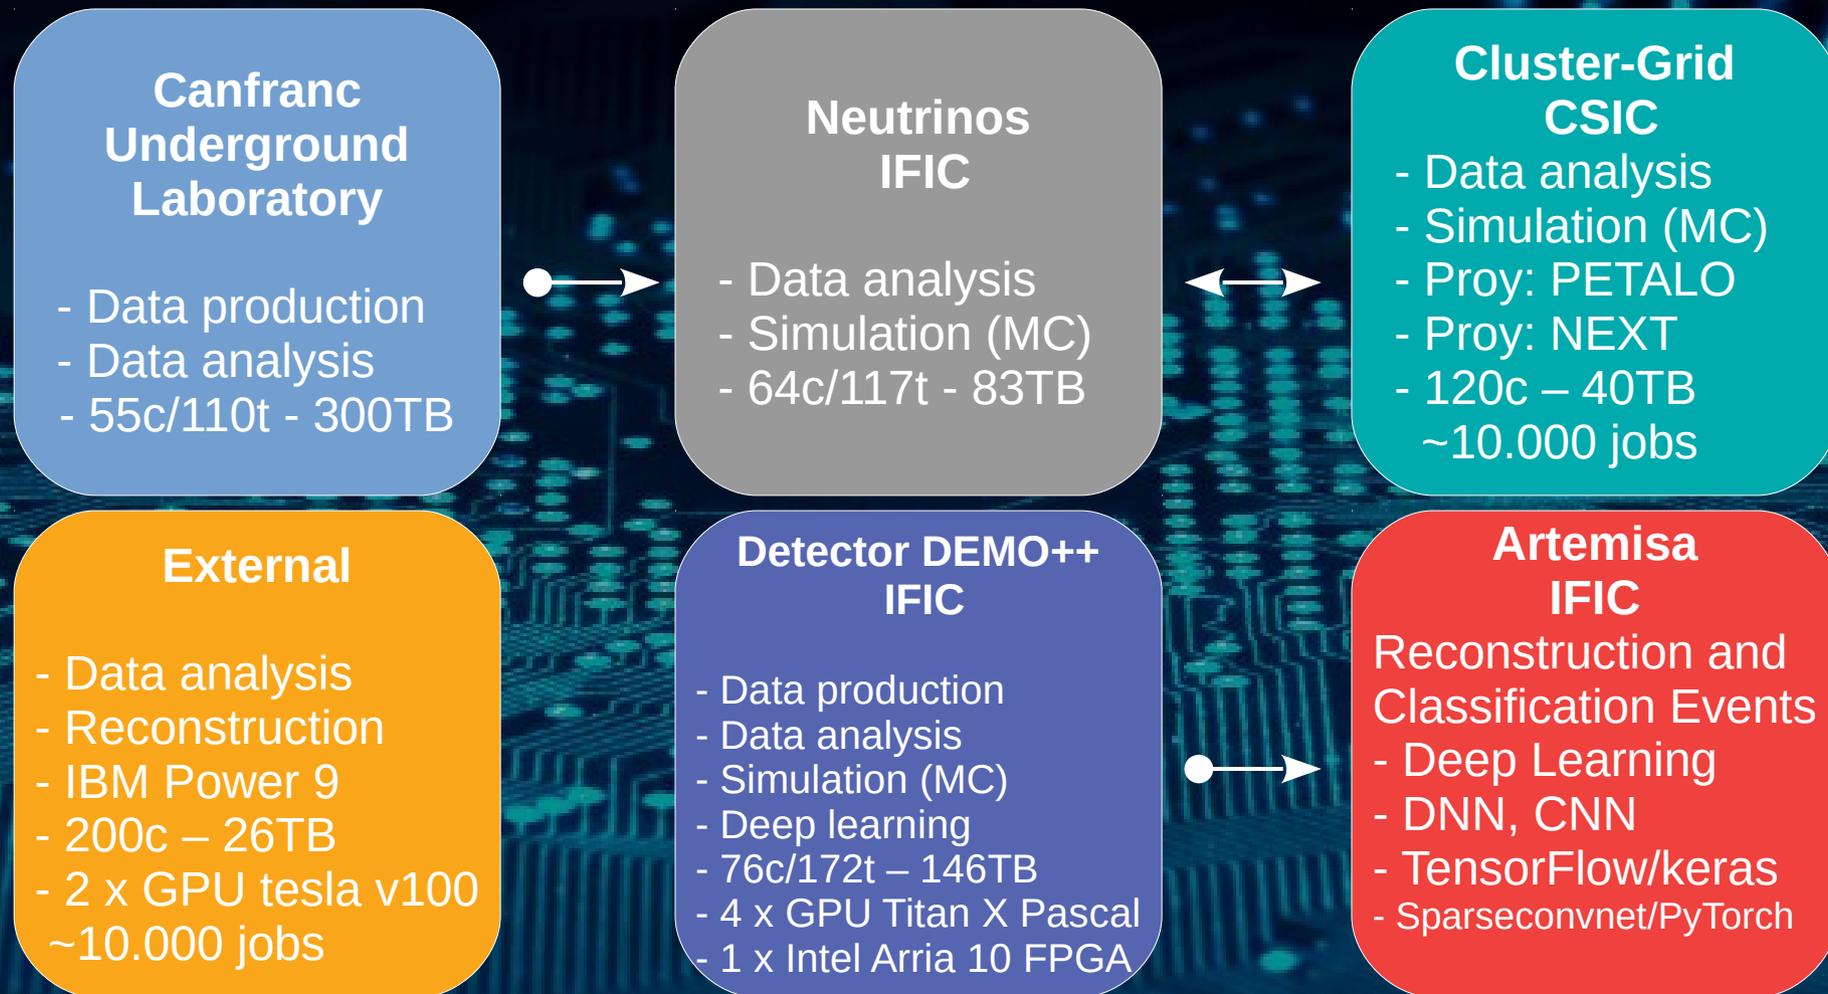
ANALYSIS RAW DATA / ATCA ARCHIVE / BACKUP

PMT 20	PMT 21	PMT 22	PMT 23
NO	NO	NO	NO

Min Time Thr (ns)	Pulse Valid Ext. (ns)
2000	50
2000	50

Min Time Thr (ns)	Pulse Valid Ext. (ns)
2000	50

Recursos de computación distribuidos en NEXT



- **Análisis de datos:** Comportamiento de la electrónica, búsqueda de picos S1 y S2 (centelleo primario, centelleo secundario), energía, posición, tiempo del evento, etc.
- **Simulación por modelización de Montecarlo:** Cálculo deposición de energía, simulación de waveforms.
- **Deep learning:** Clasificación y caracterización de eventos (señal vs background).

<https://arxiv.org/abs/1609.06202>

The background is a complex, glowing blue circuit board. It features a dense network of lines and nodes, with a central glowing element that resembles a small screen or a data hub. The overall aesthetic is high-tech and digital.

PREGUNTAS