

Secure Web Portal SWP and Elasticsearch Integration for Enhanced Data Access.

Mr PYNDA, Yaroslav¹; Mr MACGREGOR, Robert¹; Mr OLYVA, Serhiy²; Mr SUDAKOV, Alexander¹

¹ - CTBTO Preparatory Commission; ² - USoft HTI

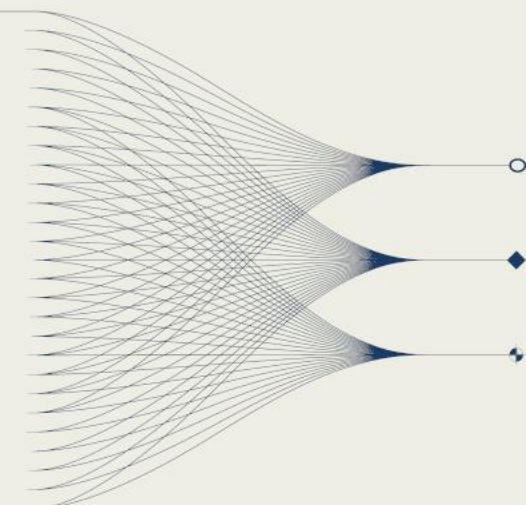


PUTTING AN
END TO NUCLEAR
EXPLOSIONS

INTRODUCTION AND MAIN RESULTS

The Secure Web Portal (SWP) is a key platform for Member States to access CTBT-related data. To overcome performance and scalability issues with the previous Oracle-based system, we integrated Elasticsearch as a distributed search and analytics engine. This upgrade improved data access speed, flexibility, and user experience.

With Elasticsearch, the SWP now offers faster search, real-time visualization, and improved system observability. It also sets the stage for future enhancements such as real-time dashboarding, machine learning, and broader use of the Elastic Common Schema (ECS), making the platform more modern and scalable.





SWP and Elastic Integration – Project Scope and Objectives

The main goal of this project is to integrate Elasticsearch into the Secure Web Portal (SWP) to address performance bottlenecks caused by traditional database queries on large historical datasets. By leveraging Elasticsearch's distributed architecture and full-text search capabilities, the system can now deliver significantly faster, more scalable, and more user-friendly access to IDC products and IMS data.

Key objectives included enhancing performance by optimizing the retrieval of long - term historical records, improving user experience through intuitive search and visualization interfaces, and ensuring the system can scale reliably as user demand and data volume grow. This shift supports the needs of Member States who rely on timely access to critical CTBT - related data.

Additionally, the integration focused on data availability and system resilience, minimizing downtime during peak access, and implementing robust monitoring and alerting with Elastic tools. This ensures early detection of issues and continuous visibility into system health, allowing for proactive maintenance and uninterrupted service delivery.



Fig 1. Elastic SHI index with 25 years of data

Data Pipelines for Populating Elasticsearch with IDC Products (SHI and RN)

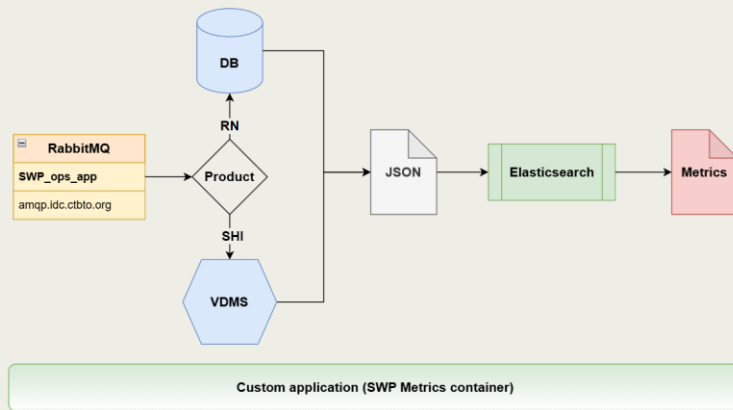


Fig 2. SWP Main operational pipeline

To efficiently index seismic and radionuclide (SHI and RN) products into Elasticsearch, we built three pipelines: operational, validation, and historical. Each addresses different data ingestion needs in the Secure Web Portal (SWP).

The Main Operational Pipeline ensures near real-time updates, pushing incoming data to Elasticsearch with low latency for up-to-date dashboards. The Validation Pipeline runs periodically, comparing Elastic data to the source Oracle database to detect and correct discrepancies. The Historical Pipeline handles bulk ingestion of archived records during initial loads or migrations. Together, they ensure complete, accurate, and scalable indexing of IDC data.

Data Visualization in SWP

The integration with **Elasticsearch** allows the SWP portal to deliver **fast, reliable, and intuitive access** to seismic and radionuclide data, including comprehensive search **across the entire 25-year historical dataset**. Users can interact with built-in visualizations that display product timelines, delivery status, and data availability, with the ability to filter, search, and download the data as needed. These features ensure quick access to critical CTBT information and support the daily workflows of Member States.

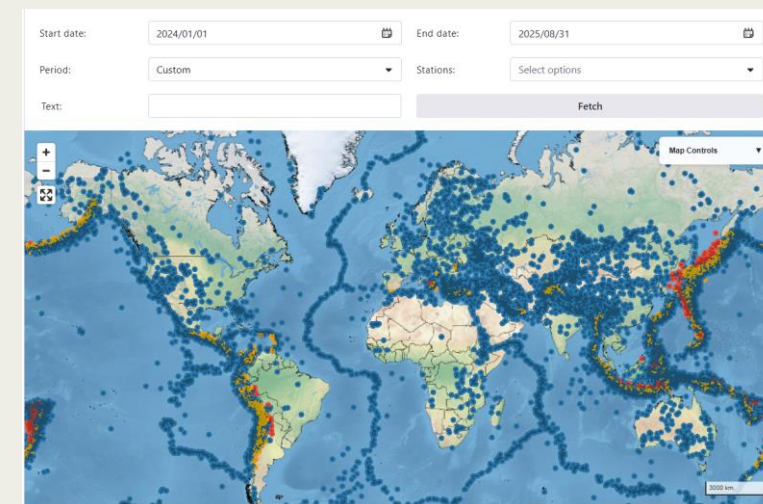


Fig 3. SWP Dashboard with REB products for 2024-2025 60923 events

The views expressed on this e-poster are those of the author and do not necessarily reflect the view of the CTBTO



SWP Data Extraction comparison: VDMS vs. Elastic

Our comparison shows that while **VDMS and Elastic deliver identical results in terms of data accuracy**, Elastic provides a **significant advantage in speed and scalability**. Queries that take **30-40 seconds in VDMS** are completed in just a **few seconds in Elastic**, even as the data volume increases from a few thousand to over one hundred thousand events. This consistent performance gain across all SWP products (SEL1, SEL2, SEL3, REB, SEB, SSEB) makes Elastic far more suitable for the **interactive and large-scale requirements of the SWP website**. By adopting Elastic as the primary method, users will experience **faster access, smoother navigation, and improved responsiveness**, ensuring that the SWP website remains reliable and efficient for both short-term monitoring and long-term historical analysis.

SWP using VDMS

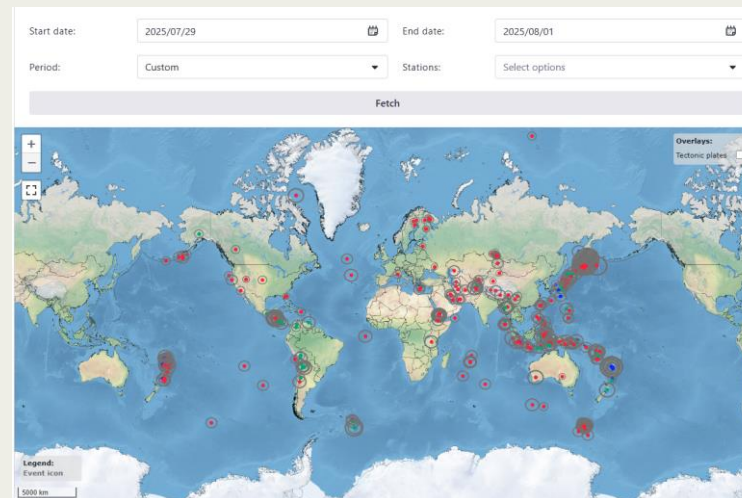


Fig 4.REB product from 2025-07-29 to 2025-08-02 time **35 sec**

SWP using Elastic

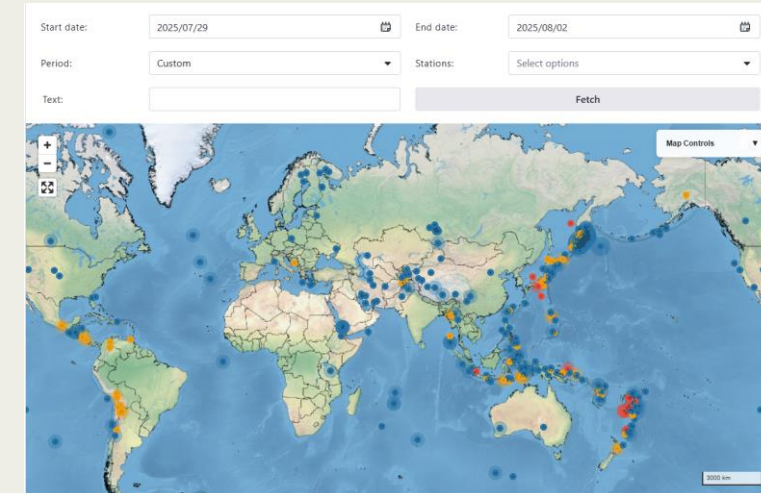


Fig 5.REB product from 2025-07-29 to 2025-08-02 time **3 sec**

Method	Start Date	End Date	SEL1	SEL2	SEL3	REB	SEB	SSEB
			Events Time	Events Time	Events Time	Events Time	Events Time	Events Time
SWP/VDMS	29/07/2025	01/08/2025	2428 38 sec	2356 35 sec	2368 28 sec	1991 35 sec	1991 38 sec	723 29 sec
SWP/Elastic	29/07/2025	01/08/2025	2428 3 sec	2356 3 sec	2368 3 sec	1991 3 sec	1991 3 sec	723 3 sec
SWP/Elastic	29/07/2025	31/08/2025	9742 3 sec	9316 3 sec	9463 3 sec	3272 2 sec	3272 2 sec	3087 2 sec
SWP/Elastic	01/06/2025	31/08/2025	20852 4 sec	19132 4 sec	19435 4 sec	9350 3 sec	9350 3 sec	7350 3 sec
SWP/Elastic	01/01/2025	31/08/2025	51615 8 sec	45859 8 sec	46687 8 sec	26074 5 sec	26074 5 sec	19468 4 sec
SWP/Elastic	01/01/2024	31/08/2025	118494 20 sec	103003 18 sec	104946 19 sec	60923 12 sec	60923 12 sec	43885 7 sec

Important Note:

The current SWP/VDMS cannot handle requests **longer than 3 days** without significant delays. Queries spanning more than **1 month cannot be executed at all** - which is why the comparison table only contains one longest VDMS example.