OSIS 2.0: OPTIMIZING ANALYST-DRIVEN AUTOMATION OF OPEN SOURCE INFORMATION COLLECTION AND PROCESSING FOR SAFEGUARDS STATE EVALUATION

T. SKOELD
IAEA
Vienna, Austria
Email: t.skoeld@iaea.org

F. COURBON
IAEA
Vienna, Austria

K. SPENCE
IAEA
Vienna, Austria

Abstract

Collecting and processing open source (OS) information is an important aspect of the IAEA's mandate to implement safeguards based on all relevant information related to States' nuclear activities. Since the mid-1990s, the Division of Information Management (SGIM) has been collecting OS information into an internal database, the Open Source Information System (OSIS). In the early stages, the SGIM collection and review process was predominately manual with classic internet searching, PDF printing, and running scripts for uploading files to the OSIS database. Over time, the process has undergone numerous improvements to include elements of automation in order to increase efficiency against an ever-growing stream of open source information. While automation has unarguably been welcome for many of the processing steps, it has been essential to keep the analyst involved at key decision-making points such as judging information for relevance, categorization, and further distribution. With the technological advancement of computing and machine learning during the past five years, however, more options for additional automation of data processing have become available. In 2016, SGIM embarked on a project to integrate and further automate the continuous monitoring, collection, and processing of OS information. This paper describes the process that culminated in the launch in early 2018 of OSIS 2.0, an in-house developed tool that has provided numerous improvements, including: automation of manual steps of collecting and formatting files; creation of a centralized space for analysts to collaborate on information collection and processing; and, improvement of the categorization and distribution capability. Furthermore, automation has enabled analysts to focus efforts more on analysis than collection and processing. The paper will also discuss possible next steps in integrating additional information collection processes into OSIS 2.0 and how far automation can be taken before it starts to have a diminishing effect on reliable information collection and processing.

1. INTRODUCTION

One of the IAEA's objectives is to provide credible assurances to the international community that States are fulfilling their safeguards obligations in that all nuclear material remains in peaceful use. To fulfil this objective, the Safeguards Department analyses on a continuous basis all Safeguards-relevant information about States' nuclear activities available to the Agency. At the heart of the State evaluation process is the consistency analysis between States' declarations, results from in-field verification activities, and information from other sources, including open sources.

Over the past 20 years, the Safeguards Department has developed and maintained processes for collecting and processing open source information related to States' nuclear activities for the State evaluation process.[1] One

[1] For the purpose of the Safegaurds evaluation process, open source information can include publically available information such as news media, scientific journal articles, and company information, as well as information in IAEA internal databases such as travel reports and technical cooperation information.

of the tools at its disposal is the Open Source Information System (OSIS), an in-house developed data storage and retrieval system. OSIS has been using off-the-shelf commercial search engines, with full-text search and a categorization component. The user interface and indexing have been customized and tuned for the Safeguards domain. A project was launched in 2016 to automate and integrate multiple workflows into OSIS and create a collaborative workspace for open source analysts. Dubbed OSIS 2.0, the project is well underway with a fully functional database already implemented and in use. The new system also offers an opportunity for exploring more advanced techniques in an effort to expand the computer's role in, filtering, categorization, and knowledge discovery.

2. BACKGROUND

2.1. History of OSIS

In the second half of the 1990s, as part of a wide-ranging effort to improve the effectiveness and efficiency of the Safeguards system, the Safeguards Department started to systematically use open source information in its assessments of States' nuclear programmes and compliance with their Safeguards agreements.[2] SGIM[3] was given the task to routinely collect, process, and categorize Safeguards-relevant open source information and make it accessible for integrated analysis as part of the State evaluation process.[4] SGIM chose Verity's SEARCH'97 Information Server as the backbone for a new database, OSIS.[5]

The first iteration of OSIS was a web-based full-text search and retrieval system that allowed for the storing of large amounts of textual data in various formats. The input for the system came from the manual collection of Safeguards-relevant open source information identified in various external databases such as Factiva and Lexis-Nexis and by visiting relevant websites, as well as searching other IAEA Departments' databases (nuclear energy, safety, and waste, and technical cooperation-related information). Information deemed relevant was saved and uploaded to the OSIS system using File Transfer Protocol software. SEARCH'97 also included a spider that indexed remote websites and made pages available locally for searching.

Apart from rich full text search and capability to search meta-data (such as date, title, and source), SEARCH'97 offered weighted queries based on so-called topic trees, which were customized with the Agency's internal technological taxonomy, the Physical Model. The topic trees used a function called "Concept Retrieval", a technology that enabled searches for subjects or concepts, as opposed to specific words or phrases.[6]

---

[2] For more information about the history of implementing information management-related measures of the IAEA's "Programme 93+2", see: NILSSON, A., CHITUMBO, K., COOLEY, J., SCHRIEFER, D., *"*Information analysis - a key element in integrated safeguards: progress and advances*"*, 40th INMM Annual Meeting, Phoenix, Arizona, USA, July 25-29 (1999); and CHITUMBO, K., BRAGIN, V., *"*Strengthening safeguards through enhanced information analysis*"*, 24th ESARDA Annual Meeting : Workshop on R&D Responses to the New Safeguards Environment, Luxembourg, May 28-30 (2002).

[3] Up until January 2007, SGIM was called SGIT (Safeguards Information Technology). The name was changed to reflect a reorganization within the Safeguards Department. The current name SGIM will be used throughout the paper for convenience.

[4] LEPINGWELL, J., Nicholas, M., BRAGUINE, V., *"*Strengthening safeguards through open source information collection and analysis*"*, INMM 44th Annual Meeting, Phoenix, AZ, July 13-17 (2003).

[5] CHITUMBO, K., and BRAGIN, V. (2002).

[6] COSTANTINI, L., HILL, J., *"*Information collection strategies to support strengthened safeguards*"*, Symposium on international safeguards: Verification and nuclear material security, Vienna, 29 Oct – 2 Nov (2001).

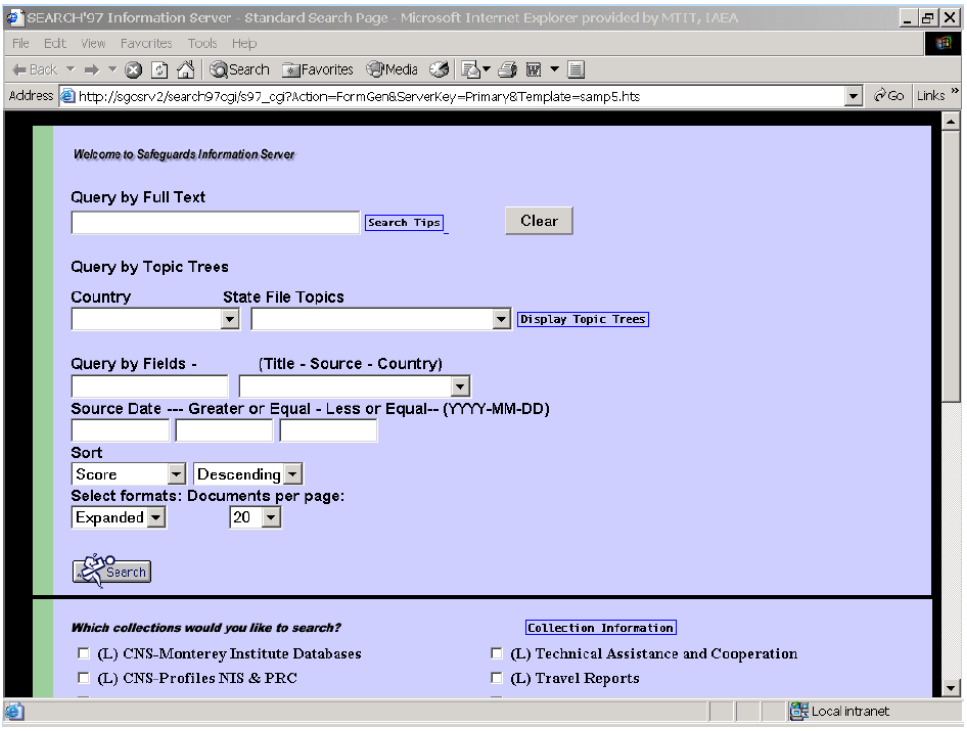

FIG. 1. OSIS user interface in 2004. Source: "Open Source Information System SEARCH'97 Information Server" Training Course, I. Tchepkasova, SGIT, 2004.

By 2000, OSIS was fully developed and implemented. Over the following decade, the search engine was first replaced by Verity's subsequent product K2 Enterprise Server and then later by Autonomy's IDOL server, both upgrades offering increased performance and additional features.[7] The information analysts in SGIM meanwhile further developed and fine-tuned processes for collecting, processing, and analysing open source information using OSIS as one of the main repositories for storing and retrieving textual data.

2.2. Addressing information flooding

With the continued digitization of information sources and the advent of new sources such as social media, the need for automation became increasingly apparent throughout the second half of the 2000s and early 2010s. As technology advanced and offered new tools to automate collection and processing, SGIM tested and implemented a number of software to assist in the collection phase; e.g. Copernic Agent and Tracker, Teleport Pro, Grokker, Website-Watcher, and various RSS aggregators such as Google Reader, Feed Demon and Feedly, to name a few. Assistance from the EC's Joint Research Centre (JRC) in Ispra, Italy, also contributed to automate information collection in a project to exploit the JRC's Europe Media Monitor (EMM) and NewsDesk applications.[8] These tools made the front-end process more efficient, enabling analysts to review and process an ever-increasing amount of incoming information on a daily basis. A number of more advanced tools promising more automatic and intelligent categorization of information were considered but not implemented within the context of routine open source information collection in SGIM.[9] The conclusion reached in all assessments was

---

[7] Autonomy acquired Verity and all its products in 2005. "Autonomy timeline", Autonomy Accounts website, https://autonomyaccounts.org/autonomy-timeline [accessed 2018-09-05].

[8] Cojazzi, G, G.M., et.al., "Collection and Analysis of Open Source News for Information Awareness and Early Warning in Nuclear Safeguards" ESARDA Bulletin, No. 50, (December 2013).

[9] The advanced tools under consideration were more relevant for analysis-related tasks. For more details on some of the nVision sub-projects, see Nicholas, M., Hilliard, J., Murray, J., "Advanced information analysis technologies for safeguards", Addressing Verification Challenges (Proceedings of an International Safeguards Symposium on Addressing Verification Challenges, IAEA, Vienna, Austria, 2006).

that the analyst was still best suited to categorize incoming textual data and the risk of "missing" important information was judged as too high to allow the machine to take over initial filtering.

3. THE NEW OSIS SYSTEM

3.1. Development of OSIS 2.0

In 2016, SGIM started a project to further automate as many steps as possible in its open source information collection processes to minimize manual non-analytical tasks for analysts. Apart from the need to address the information-overload issue described above, SGIM was facing staffing issues and needed to find ways to make work processes more efficient.

One of the key factors to the successful development of OSIS and user adoption was the frequent interactions between users and the analytics specialist, who was embedded in the team. In close collaboration, users and the analytics specialist mapped existing processes in detail and identified steps that could be candidates for automation. As development proceeded and prototypes were built and tested, requirements were modified and taken into account in a pragmatic give and take relationship between users and the analytics specialist. The result was a tool that has provided:

- Automation of several manual steps of collecting and formatting files;
- Creation of a centralized space for analysts to collaborate on information collection and processing;
- Improvement of the categorization and distribution capability.

The team used a commercially available project management tool to track issues and communicate requirements and prioritize sub-projects. This collaborative workspace helped with planning, documentation, and keeping an overview of the project, including knowing which issues still needed attention.

3.2. Description of the system

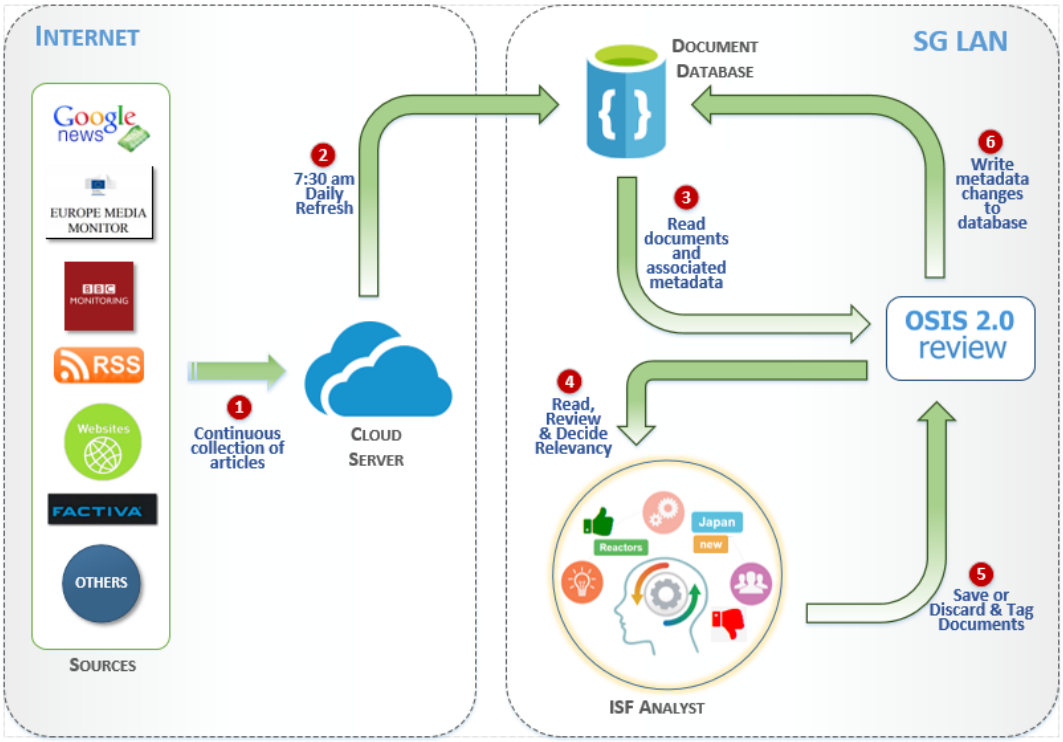

FIG. 2. OSIS 2.0 simplified architecture.

The new OSIS application consists of a web crawler that collects and sanitizes targeted web pages, and a user interface to browse, review and annotate the documents collected by the crawler.

The web crawler is a back-end process running in the cloud that continuously monitors hundreds of RSS URLs and static web pages for new links. The content/target page of each new link is collected and sanitized (removal of superfluous content like navigation bars and ads). The document is converted into a PDF, a rich format supported on virtually any device. The PDF and associated metadata (source, published date, URL etc.) is stored in a NoSQL database (MongoDB), feeding the user interface.

The user interface is a web application for browsing, reviewing and annotating the collected documents, similar to a RSS reader, with additional features specific to Safeguards workflows. At its core, a user can review new documents in analyst-specific queues, retain and further annotate relevant documents by adding tags (country name, category such as "background information" or a specific step in the nuclear fuel cycle). The remaining documents are marked as discarded. The application features classical full-text searching capabilities (Boolean, wildcard, proximity, highlighting, etc.) as well as faceted search and aggregated results. It is built on top of Elasticsearch, currently the most popular enterprise search engine.

The OSIS application currently consolidates a number of time consuming, low value operations previously carried out by analysts via a variety of tools, into a single interface allowing them to focus on the core activity of assessing/enriching the value of open source information.

3.3. Use cases

OSIS 2.0 has provided a platform in which analysts can collaboratively collect, process, and analyse open source information, for both broad collections and State-specific monitoring, within a single system. As the main tool of the Global Monitoring Team (GMT), a small group of analysts are able to monitor pre-defined feeds organized in specific "baskets" that cover nuclear-related topics throughout the world spanning thousands of sources. The information is divided among the GMT and can be reviewed simultaneously by several GMT members and other open source analysts in SGIM. This information is tagged and modified as needed and further disseminated to the appropriate state analysts through a daily automated e-mail. Additionally, the GMT's collaborative work has streamlined the SGIM Open Source Highlights newsletter production process. With OSIS 2.0, the Editor of the Highlights newsletter is able to view all relevant articles in one system allowing for a more comprehensive and efficient understanding of the day's collection.

For State evaluation, OSIS 2.0 allows analysts to set up broad feeds specific to their assigned states in addition to more targeted searches. While the GMT baskets are aimed at a broader range of sources and topics, the State specific baskets use feeds for sources and topics unique to the State's nuclear landscape. The analyst or analysts are able to easily review and process the collection with higher confidence as to its relevance. Furthermore, the analyst can easily share and save documents as well as incorporate the information directly into the State evaluation process.

After analyst approval, all articles saved by the GMT or other analysts are saved in the OSIS 2.0 Library, which is accessible Section-wide, and can be easily transferred to the Integrated Safeguards Environment (ISE) for use by the State Evaluation Group (SEG) or others in the Safeguards Department.

3.4. Integrating additional workflows

Additional uses of OSIS 2.0 for analysis include the incorporation of multimedia and trade information and geospatial information. SGIM recently launched the Safeguards Multimedia Information Analysis and Integration (SG-MM) project with the intention of fully exploiting non-text-based information.[10] Although SGIM has been collecting and analysing multimedia information (images, videos, schematics, etc.) for years, a centralized repository that allows for collection, processing, and analysing has yet to be established. Through modifications of the OSIS 2.0 interface, analysts will soon be able to automatically and collaboratively collect

[10] See paper by Fowler, M., et al., "Optimizing the use of multimedia information in IAEA Safeguards". IAEA Safeguards Symposium, November 2018.

and process multimedia information as well as extract the full resolution image or video for further processing and analysis.

Another area in which OSIS 2.0 will be utilized is the collection of trade-related information such as dual-use items and import/export data. Using detailed tagging corresponding the Additional Protocol Annex II, both trade and open source analysts will be able to more easily identify and analyse trade concerns.

4. MOVING FORWARD: OPTIMIZATION OF AUTOMATED INFORMATION COLLECTION

A key objective of the development of the new OSIS system was to build a database of documents, enabling a variety of processing previously not possible with the file-based structure. In particular, in addition to the primary goal of identifying and annotating Safeguards relevant documents (positive samples), it became possible to keep those documents discarded by the analysts (negative samples) at no extra cost other than disk space, something that was not possible in the previous workflow where only relevant (positive) documents were kept. Thus, after two years of usage, SGIM has built a sizable dataset suitable for exploring advanced filtering options.

A simple classifier based on document similarity has already been deployed and is being used by analysts. The classifier predicts with good accuracy the country name and the step in the Nuclear Fuel Cycle of an input document. A more sophisticated "Safeguards relevant" classifier prototype is in development. If found suitable, it could improve the SGIM open source information collection workflow by filtering out significantly the number of irrelevant documents to review, as well as making the annotation process more efficient, freeing up even more time for the analyst to concentrate on tasks of higher value.

An important element of successfully implementing more advanced automated classification and filtering will be buy-in from users. Intrinsically, information analysts tend to hold a healthy dose of scepticism towards “too much” automation of analytical tasks, such as assessment of relevance and importance of information. To gain the trust of analysts and implement more automation, the project must continue to rely on constant and responsive interaction between all stakeholders during development whereby users are kept in the loop and trying new features as they are being built, reporting back to the analytics specialist whether it works or not and how it can be improved.

Careful and gradual implementation of automation marked by proven successes along the way will go a long way to gain analysts’ trust and leverage advanced automation to infuse more data-driven decisions in the open source information collection and processing workflows in SGIM.