



Inside the Spring 2021 GOST® Release

GOST Automates Screening

The Spring 2021 GOST Release provides massive improvement in efficiency and expands scale to meet the demands of the largest enterprises.

GOST Spring 2021 Offers Revolutionized Entity Resolution and Scalability

The Spring 2021 GOST Release introduces new search processing infrastructure to support large-scale ongoing monitoring, allowing even the largest enterprise institutions the ability to effectively screen populations to reduce risk. The Spring 2021 release increases GOST processing speed by 4x, measured in searches per second. Additionally, updates to both link-level reliability and case-level reliability in the Spring 2021 GOST Release ensure an increase in the ease of use for the frontline workforce responsible for protecting enterprises from risk through the screening process.

A Deeper Look

Screening detects potentially problematic behaviors and information early. Screening is critical to meet legal, regulatory, and risk compliance requirements. To achieve maximum effectiveness, screening must be done continuously across entire populations. Unfortunately, the limitations of existing solutions drive institutions to make compromises as they try to meet their requirements.

There are two critical challenges to screening full populations continuously: the first is scalability, the ability to quickly process data on hundreds of thousands or millions of entities daily; and the second is entity resolution, the ability to reduce false positive results while still detecting adverse information. Public and private sector organizations cannot sustain workflows at the necessary scale without getting overwhelmed by the alert yield. At best, they can screen subsets of the population, based on risk segmentation or decision rules.

With critical enhancements to scalability and entity resolution, the GOST Spring 2021 Release enables the largest enterprises to screen full populations without making costly tradeoffs that increase risk exposure and decrease adverse media program effectiveness.

Scalability

When you have a small population to screen, professionals can manually run adverse media searches to accomplish the screening mandate. As the population increases, however, technology is required to achieve scale. Industry solutions have built tools to support list-based matching of names as large-scale for simple matching tasks like sanctions and PEP screening- data which is highly structured and mostly unchanging.

GOST has uniquely achieved intelligent scale for diverse, dynamic, and unstructured adverse media data, without restricting coverage to specific news media sources or relying on a predetermined list of names. The GOST Spring 2021 Release provides the infrastructure necessary for GOST to complete over 1,000,000 searches per day, an increase of 4x the previous rate. Institutions can screen and continuously vet entire populations quickly and with meaningful results.

The ability to process over a million searches per domain, per day is important for many reasons. First, institutions are no longer limited in the number of entities that can be screened. In other words, institutions no longer have to risk-accept major exceptions to their screening practices by subsetting populations and screening at different intervals. With the GOST Spring 2021 Release, screening at onboarding is virtually unlimited and periodic review intervals can be shortened for quicker time to insight. With these scale improvements, GOST makes due diligence backlogs a thing of the past.

Second, institutions now have the ability to scale their continuous vetting practices with GOST. One million searches per day means that GOST does not just scale to the volume requirements of the largest enterprise institutions, but also to the frequency requirements of monitoring workflows. A bank that refreshes adverse media on a customer population every two years can now do it quarterly. An investment management company that monitors sources of reputational and financial stress for their registered investment advisers and broker-dealer can refresh their data weekly instead of annually. A government agency conducting continuous evaluation of cleared professionals can act on changes in near-real time as opposed to when a significant change is reported or a review is conducted. GOST identifies potential sources of risk as fast as the information changes, empowering analysts to follow up on the highest value leads on an ongoing basis.



Entity Resolution

GOST's unique approach to entity resolution through the use of artificial intelligence (AI) and machine learning (ML) is unmatched in the current screening and vetting market.

GOST performs entity resolution at two levels -- first in the extraction and ordering of informational content that is returned, and second in the aggregation of that information to a top-level score that allows you to differentiate between entities with more or less reliable information that is verifiably about them. This means that GOST can filter on results that are less likely to be false positive matches.

At the content-level, GOST creates concepts and measures the conceptual distance between the entity as-searched and the entity in the data GOST returns. This allows the GOST user to see a rank-ordered list of content in order of how closely related the content and the searched entity are.

GOST puts the user in control of whether they want to see content about the right person, or if they want to prioritize derogatory content regardless of precise entity resolution. In most tools on the market, matches are limited to exact matches on structured data, like watchlists. Without GOST, the user will miss a lot of risk-indicative information readily discoverable in unstructured data. Because GOST uses content-level entity resolution capabilities, the customer is given access to better coverage and greater variety in the data. The GOST Spring 2021 Release includes significant improvements to the rank-ordered presentation of this information, and the user's ability to make explicit tradeoffs in the prioritization of information.

At the entity-level, in addition to the standard Ranking score GOST provides as a measure of risk-relevant information, GOST also provides an aggregate entity Reliability score that represents the likelihood that the entity being searched in fact matches across the body of information returned.

To produce this Reliability score, GOST's machine learning algorithms aggregate the reliability of the content-level information to summarize whether the GOST user should expect highly-resolved information to be available in the respective entity's search results.

GOST's unique approach to entity resolution through the use of artificial intelligence (AI) and machine learning (ML) is unmatched in the current screening and vetting market.

Content-Level Reliability (CLR)

The probability score calculated by GOST determining the likelihood that the information contained within a link is associated with the searched entity.

Entity-Level Reliability (ELR)

The aggregate reliability score provided by GOST representing the likelihood that the entity being searched matches across the body of information returned.

The Reliability score allows users to rank-order entities on the dimension of entity resolution, for every entity in the customer population. This means that GOST guides you to where you can expect a higher share of false positive matches, and where you can expect a higher share of true positive matches. This can be effective both ways. For organizations with limited resources that need to spend time on relevant and reliable information only, they may want to focus on the extremely high end of the Reliability spectrum. For organizations looking for synthetic identities or potential sources of fraud, they may want to focus on the lower end of the Reliability spectrum, where the entity search inputs do not effectively resolve to a discernible identity. This also means that unlike industry-standard tools that only provide information conditional on a validated or manually resolved match, GOST will return information on every entity searched. This can be useful search content across the entire enterprise, and not just for the screening function where it is deployed. No other tool is this transparent and adaptable for a number of different functions and workflows across the enterprise.

Conversely, other tools currently yield a many-to-one alert ratio - i.e. because alerts are on content, not on entities, you can receive redundant alerts. This volume of alerts is unwieldy and unnecessary. The user needs an easily-accessible summary of the entity in question, and a method to prioritize entities for review.

GOST is unique in the way it summarizes adverse media information under fixed entity profiles that the user can access, eliminating the overwhelming nature of multiple alerts per entity.

The combined improvements of content-level and entity-level reliability put the user firmly in control of how many false positives they choose to see for different use cases. This is illustrated in the four scenarios below, all of which are achievable using the slider scale in the GOST user interface.

As a User:

1. I only want to see high Reliability entities so that every high Ranking result I review is likely to have relevant information about my customer, and include adverse media.
2. I only want to see high Reliability entities so that I can guarantee the information I use for my risk-rating model is positively associated with my customer, and that negative but potentially false positive information is not contributing to my risk rating.
3. I only want to see information that can be positively associated with my customer, and I only want those artifacts to contribute to the Ranking score.
4. I want to allow information with low or medium content-level reliability to contribute to Ranking so that I can review and apply expert judgement to whether or not there is a risk or threat from doing business with that entity.

the right person

the right topic



Update Summary

Scalability:

Infrastructure is now in place that allows GOST to complete 4x the searches in the same time. This amounts to 1,000,000 or more searches per domain, per day. The new infrastructure allows any institution, including the largest enterprise institutions, to be certain that their ongoing monitoring needs, as well as their onboarding needs, can be met effectively and at scale for volume and frequency.

Entity Resolution:

Updates to content-level reliability include an increase in the performance of entity resolution at the information level. Changes to the features and methodology of the content-level reliability model allow the user to better discern whether links are more likely to be about the entity in question.

Updates to the scoring mechanism result in an increase in sensitivity of the link slider bar in the Organization Administrator settings. Updates to the link-level reliability model allow the slider bar to accurately sort links in the Domain-Specific Results tab. Users can now change the link slider bar to weigh either derogatoriness or entity resolution to see more apparent changes in the sorting of the links in results specific to your topic of interest. Users can make a tradeoff between content that is more likely to be about the entity in question or content that is more likely to have risk-relevant information regardless of entity resolution. No other solution puts this capability in the user's hands.

Updates to entity-level reliability include Reliability scores that are more reflective of content in results related to the topic of interest -- i.e. what is risk-relevant information for your domain. The entity resolution at the document level now contributes more directly to the entity-level Reliability score.

Updates to the querying of entities in GOSTCrawl and outside data holdings bring back higher-quality content.

To learn more about GOST or to speak to one of our experts, visit us at:
www.giantoak.com