QINETIQ

ThreatCanaryAI

Artificial Intelligence/Machine Learning Open-Source Intelligence Threat Detection Solution

QinetiQ helps our clients achieve their key objectives with individualized consulting, technical, and programmatic support services. We specialize in building experienced teams who deliver long-term results for our clients in the Civilian, Homeland Security, National Security, and Defense spaces, including through Artificial Intelligence (AI) and Machine Learning (ML).

Machine Learning is a branch of artificial intelligence and computer science that uses data and algorithms to enable AI to imitate the way that humans learn.

ThreatCanaryAI is an AI/ML solution developed by QinetiQ's Data and Technology Group to detect explicit threats and inappropriate communications against U.S. government officials and law enforcement.

QinetiQ deployed ThreatCanaryAI for use by our partners at the U.S. Marshals Service (USMS) under the Department of Justice in October 2022. ThreatCanaryAI supports the work of the Judicial Security Division's Open-Source Intelligence (JSD-OSINT) unit in their large protective mission, which provides security for the country's 94 Judicial districts, including 2,200 sitting judges and 26,000 Federal prosecutors and court officials.

ThreatCanaryAI augments the capabilities of OSINT analysts by automating web-scraping, text processing, and threat prediction tasks, presenting collated results to the user in a convenient web interface for review. With coverage of over a dozen social media platforms, ThreatCanaryAI yielded a dramatic expansion to the JSD-OSINT unit's capability, upgrading the unit's manual screening workflow to process hundreds of thousands of internet comments each week. ThreatCanaryAI accounts for 59% of all open-source threats against protected personnel referred to USMS district investigators and has created a unit-productivity improvement of over 200%.

ThreatCanaryAI's threat prediction algorithm leverages a language model that has been trained on over 1 million labeled records, as well as threat keyword identification of around 30K words and phrases. Using modern Natural Language Processing (NLP) techniques, ThreatCanaryAI simulates human reasoning to accurately identify violent threats amidst an ocean of innocuous internet content. In direct comparison with third-party OSINT tools, ThreatCanaryAI's predictions were 42% more accurate than its closest competitor. Other notable features include doxing detection, Google dork automation, and video transcription.

QINETIQ

Increased Productivity and Enhanced Threat Detection

After launching in November 2022, ThreatCanaryAI immediately made an impact on unit productivity, precipitating a fundamental change in how the unit performs its work.

Referrals Over Time By Source %

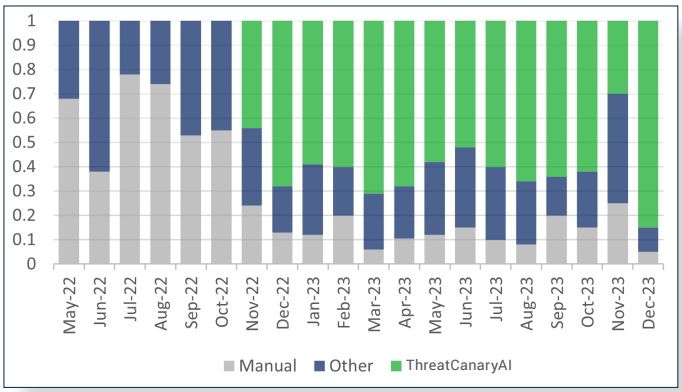

ThreatCanaryAI immediately made an impact on unit productivity, precipitating a fundamental change in how threat detection work is performed. Manual work decreased and ThreatCanaryAI usage increased, surpassing other tools.

Benefit	Description
Timesaving	Analyst work that previously required hours can now be completed in minutes
Productivity Gains	ThreatCanaryAI enabled JSD-OSINT to review more content and refer more threats leading to a 200% improvement in unit productivity
Cost Savings	With ThreatCanaryAI, agencies stand to save hundreds of thousands of hours as a function of productivity gains

Enhanced Performance

ThreatCanaryAI's advanced continuous monitoring capabilities, provide threat detection and contextual analysis. Its sophisticated pattern recognition can identify subtle threat indicators across vast amounts of online data, delivering critical alerts with speed and precision. In direct comparison with third-party OSINT tools, ThreatCanaryAI's predictions were 42% more accurate than its closest competitor.

ThreatCanaryAI vs Third-Party Competition: Threat Detection Performance

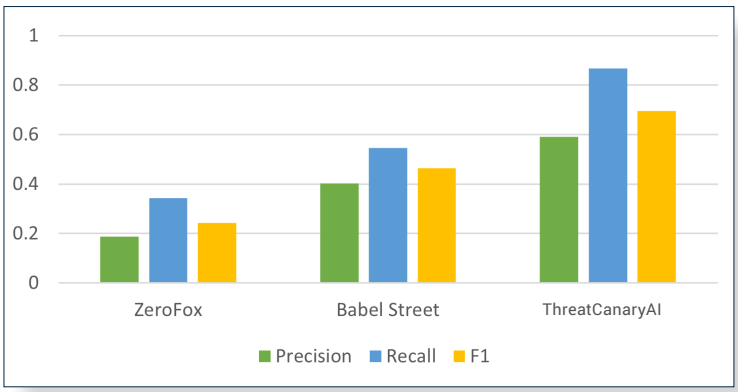

ThreatCanaryAI vs ZeroFox

	Precision	Recall	F1
ThreatCanaryAI	0.671	0.797	0.729
ZeroFox	0.188	0.344	0.243

ThreatCanaryAI vs Babel Street

	Precision	Recall	F1
ThreatCanaryAI	0.513	0.937	0.663
Babel Street	0.403	0.547	0.464

Unit of Measure	Definition
Recall & Precision	Recall and precision are classic metrics for evaluating AI classification models, such as ThreatCanaryAI. Classification metrics are relevant because we are classifying whether a comment is, or is not, a potential threat.
Recall	Recall is how comprehensive the model performs in capturing all potential threats. Optimizing for recall will minimize false negatives.
Precision	Precision measures accuracy of the classification model when it makes a positive classification that something is a potential threat. Optimizing for precision will minimize false positives.
F1	F1 is a balanced score combining both precision and recall.
Score Interpretation	A score closer to 1 signifies better performance.

Collaborating with QinetiQ

QinetiQ welcomes the opportunity to discuss how our experiences and customized solutions can help partners in the homeland security space meet mission goals.

www.QinetiQ.com/en-US

For further information please contact:

Andrés Meneses
Data & Technology Group Lead,
QinetiQ US
(202) 714-9870
andres.meneses@us.qinetiq.com