

# Proofpoint NexusAI for Compliance

The digital communications revolution is spawning more content—and more types of content—than ever before. This information deluge is driving a revolution of its own. Artificial intelligence (AI) and machine learning (ML) are set to transform compliance and e-discovery review. But while compliance, surveillance and legal teams are inspired by the promise of AI, they are struggling with the reality. Many of these technologies are inefficient. Some just aren't effective. And most don't scale.

Proofpoint NexusAI for Compliance can help. It augments the rich detection and analytics features of Proofpoint Intelligent Supervision to further streamline and improve your review process.

## Introducing Proofpoint NexusAI for Compliance

Monitoring communications, as required in financial services and other regulated sectors, has always been a balancing act. Companies must cast a wide net to find all relevant compliance issues. But casting the net too widely results in more false positives—and heavier workloads.

NexusAI for Compliance is an add-on for Intelligent Supervision that uses advanced ML to greatly reduce your reviewers' workload. You can easily deploy, validate and implement NexusAI models (the code used to evaluate communications of all types) to improve supervision decision-making and automate key parts of your workflow.

Unlike standalone machine-learning tools, NexusAI for Compliance builds upon the already-rich workflows and reporting you get from Intelligent Supervision. By adding NexusAI for Compliance, your supervision and surveillance efforts can become even more efficient and effective, further lowering costs.

## Achieving real-world outcomes with Proofpoint

Naturally, results will be different for every customer; AI models, data sets, and training can vary widely by company. But early NexusAI for Compliance customers have reported the following:

- A 25-50% reduction in false positives
- Ongoing ROI of 125% per year

NexusAI models have already helped identify the following improvements for early customers:

- **Disclaimer detection.** Suggests disclaimer text to add to Supervision's configuration for pre-approved content
- **Exclusion detection.** Suggests senders and subjects, such as newsletter senders, to add to Exclusion Rule configuration
- **Auto review.** Secondary evaluation of flagged messages with the ability to intelligently clear or escalate the messages as required

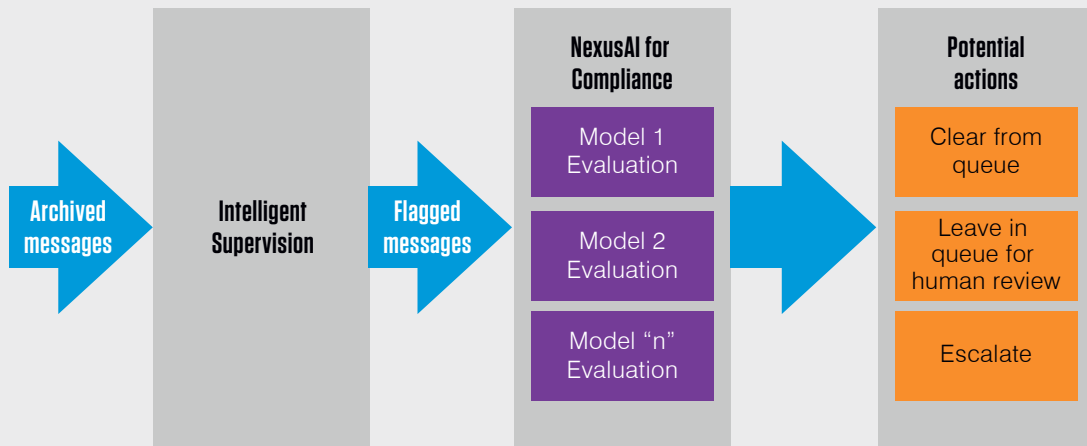


Figure 1: Workflow of Intelligent Supervision with NexusAI for Compliance

And these are just the start. NexusAI's open, scalable AI engine can tackle more complex and varied classes of communications-related decisions. These include detecting fraud, harassment, discrimination, coercion and other policy breaches. With NexusAI, your Proofpoint Archiving and Compliance deployment can make an even bigger business impact.

## Using models to automate and refine supervisory

With traditional compliance tools and static rules, supervision review queues can quickly fill up with low-risk content. Boilerplate text, standard disclaimers and automated notifications are just a few examples. Reviewing it all takes time and resources. For already-stretched compliance teams, that means higher costs and longer time to value.

NexusAI refines your review queues, intelligently winnowing out low-risk content to prevent false positives in future reviews. With other compliance review tools, identifying new disclaimers and newsletter senders can take weeks or even months. NexusAI shortens this cycle, thanks to ML models that enable much faster review decision-making.

NexusAI can even automate the process of clearing or escalating secondary content reviews. Reviewers are freed up to focus on the most business-critical content without worrying about missing out on the entire scope of data.

## A technical view: models and review automation

NexusAI for Compliance uses ML models to evaluate messages flagged for review by Proofpoint Intelligent Supervision. (See Figure 1.) These can be any message type supported by Proofpoint Enterprise Archive, such as:

- Email
- Social media (such as Twitter or LinkedIn)
- Collaboration platforms (such as Slack or Microsoft Teams)
- Mobile (such WhatsApp or voice)

Models are configured for items flagged for one or many rules. Each model evaluates the relevant items and assigns a score. That score is then added to the messages within Intelligent Supervision.

In evaluation mode, reviewers still review messages; reporting enables you to compare between model-driven and human-driven review decisions. Once you have high enough confidence, you can activate the model to take the same actions as a human reviewer would. In practice, it's automated, autonomous review.

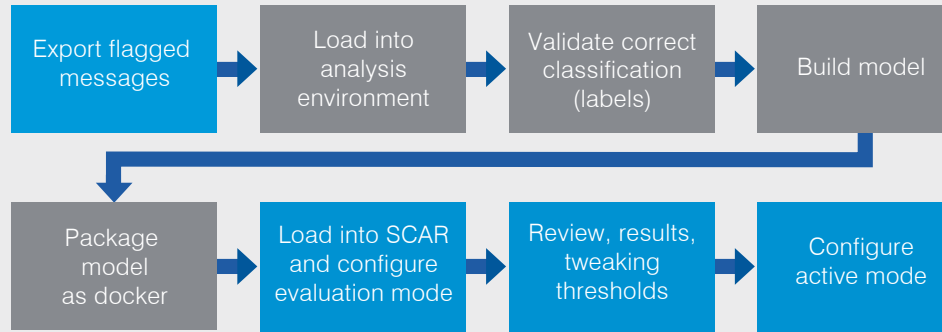


Figure 2: NexusAI for Compliance training and deploying ML models

For flagged messages with corresponding NexusAI models, the scoring ranges and actions may be set up as follows within Intelligent Supervision:

- 0.0 to 0.15: High confidence that the flagged item is not an issue. NexusAI clears the issue.
- >0.15 to <0.8: Low confidence in the message evaluation. NexusAI leaves the message in the queue for human review.
- 0.8 – 1.0: High confidence that the flagged item is an issue. NexusAI escalates the message for further investigation.

When the models are activated, evaluations, actions and related data are all recorded to the same standard as if a human reviewer is making the decision.

## Key attributes of NexusAI models

Like most ML models, NexusAI's are developed using data to train ML algorithms.<sup>1</sup> Before any model is produced, the data first needs to be prepared. (This step is sometimes called data wrangling.) The data is gathered, prepared and divided into two sets: a training set and a testing set. At this stage, the data must be classified and labeled to help describe the meaning a human would derive from it.

Intelligent Supervision provides a natural mechanism for labeling data—simply train your staff to properly categorize issues as they review them. The message data (including metadata about which flags matched and why) and review decisions can easily be exported from Intelligent Supervision. (The format is well suited for building NexusAI models.)

With the data ready, you have a baseline model. You can measure the effectiveness of the model by testing it on the test data set. Effectiveness is typically measured using two scores: precision and recall.<sup>2</sup>

In simple terms, precision is the percentage of the results correctly classified as being of interest. Recall is how many messages were correctly identified (that is, not missed).

## Training and deploying models

With the baseline model complete, you can train and compare new models to see which ones offer better precision and recall scores. Sometimes, the two scores may be a tradeoff. Recall may be more important in some situations, while precision may matter more in others.

To be used in the NexusAI for Compliance framework, a model is packaged as a web service that implements our REST API specification and is delivered in a Docker container.<sup>3</sup>

The model can then be set up to operate one of these ways in production:

- **Evaluate.** The model evaluates every message in a particular class but takes no action. This mode allows the model to be tested in production on real-world data flows without performing actual reviews.
- **Enabled.** The model evaluates communications in a class and acts on them. You can also specify a sampling rate to evaluate a sample of messages but take no action on them. This approach enables discrepancy reports that compare model performance against that of human reviewers.

<sup>1</sup> Javatpoint. "Machine learning life cycle." Accessed April 2021.

<sup>2</sup> Wikipedia. "Precision and recall." Updated March 2021.

<sup>3</sup> Docker. "What is a Container?" Accessed April 2021.

This support for running new and updated models offers two main benefits. First, it lets you evaluate them effectively. Second, it provides the evidence of successful testing you may need under validation mandates. You can compare updated models to both human reviewer decisions and those made by the model.

## Using a flexible ecosystem for model execution

NexusAI for Compliance is a highly flexible platform. It runs models written in any language, using virtually any analytics framework that fits your needs.

You can use the out-of-the-box models directly. Your data science team can also build new models that run within the framework. This flexibility makes use of your team's unique knowledge of the business—without the work of building and supporting a scalable platform from the ground up. If you need more support, Proofpoint professional services can also help you create custom NexusAI models.

## Conclusion and next steps

NexusAI for Compliance achieves the promise of AI and ML at scale. Our unique framework helps streamline and improve the way you supervise digital communications to comply with regulations and corporate mandates. It deploys quickly, makes training and testing easy, and offers an unmatched return on investment. With NexusAI, you can safely modernize your supervision practices for new levels of efficacy and efficiency.

### LEARN MORE

For more information, visit [www.proofpoint.com/us/nofalsepositives](https://www.proofpoint.com/us/nofalsepositives).

---

#### ABOUT PROOFPOINT

Proofpoint, Inc. (NASDAQ: PFPT) is a leading cybersecurity and compliance company that protects organizations' greatest assets and biggest risks: their people. With an integrated suite of cloud-based solutions, Proofpoint helps companies around the world stop targeted threats, safeguard their data, and make their users more resilient against cyber attacks. Leading organizations of all sizes, including more than half of the Fortune 1000, rely on Proofpoint for people-centric security and compliance solutions that mitigate their most critical risks across email, the cloud, social media, and the web. More information is available at [www.proofpoint.com](https://www.proofpoint.com).

©Proofpoint, Inc. Proofpoint is a trademark of Proofpoint, Inc. in the United States and other countries. All other trademarks contained herein are property of their respective owners. [Proofpoint.com](https://www.proofpoint.com)