



Flagright's Approach to

Rule Performance and Enhanced Detection



Financial institutions today face a dual challenge: they must **detect increasingly sophisticated illicit activity** while **minimizing false alarms** that drain compliance resources. Traditional rule-based monitoring systems often generate an overwhelming volume of false positives, by some estimates, roughly 90% of alerts in legacy systems are false positives. This alert fatigue not only wastes investigators' time but can also let true threats slip through. Flagright addresses these issues with a modern, AI-native platform that emphasizes rule performance, observability, and analytics. Its approach empowers compliance leads, risk officers, fraud/AML analysts, and product teams to fine-tune detection logic continuously without coding or engineering support.




In this guide, we explore in detail five key capabilities of Flagright's platform and how they work together to reduce noise, improve detection quality, and ensure adaptable, auditable compliance programs:

01	02
Rule Simulations Test detection logic and thresholds against historical data before going live.	Shadow Rules Safely trial new rules in real-time (production) without triggering real alerts.
03	04
Rule Analytics & Observability Built-in metrics and insights (trigger frequencies, false positive rates, trends) to monitor and refine rule performance over time.	Rule Threshold Recommender AI-driven suggestions for optimizing rule thresholds based on historical performance.
05	06
Rule Recommender AI guidance to identify coverage gaps and suggest new rule ideas based on typologies and patterns observed.	Dynamic rules Expand detection beyond fixed rules with adaptive, context-driven signals.

Each of these features contributes to a more precise and efficient compliance operation. Together, they enable safe, continuous tuning of detection systems, reducing false positives while bolstering confidence that real risks will be caught. Below, we describe each capability and its practical use in the Flagright platform.

Rule Simulations

Backtesting with Historical Data




One of the pillars of Flagright’s approach is Rule Simulation, which allows teams to test detection logic on historical data before deploying it in production. In practice, this means compliance analysts can create or modify a rule and then run it against past transaction datasets to see how it would have performed. This back-testing capability is invaluable for understanding a rule’s potential impact: it reveals how many alerts the rule would generate, how many transactions or accounts would be flagged, and even the likely false positive volume. Flagright’s platform makes this process fast and code-free; users can simulate and test rules in seconds without engineers.

How it works	Using the no-code Rule Builder, a user defines a new rule or adjustment (for example, a threshold change or added condition). They can then execute a simulation which processes historical transactions through the new logic. The platform returns a detailed summary of results, typically including metrics such as:
	✔ Alerts (Cases) that would have been generated How many alert cases the rule would have triggered on past data.
	✔ Transactions hit by the rule The count (and percentage) of historical transactions that meet the rule criteria.
	✔ Unique users/accounts affected How many distinct customers would have been flagged by the rule.




By examining these outputs, teams can gauge the rule's sensitivity and effectiveness. For instance, if a simulation shows that a proposed rule would have created hundreds of alerts in the last month (with most being low-risk transactions), the team might recognize a potential false-positive problem and refine the rule before it ever goes live. Conversely, if the rule barely fires at all on historical data, it might be too narrow to be useful. The Simulator lets users adjust parameters (e.g. threshold values) and immediately see the impact on alert volume, false positives, and investigator workload. This iterative tuning builds confidence that when the rule is deployed, it will catch genuine issues without flooding analysts with noise.



Importantly, simulation in Flagright is not a one-and-done process but part of an iterative rule development cycle. Teams often test multiple iterations of a scenario by tweaking conditions and rerunning the back-test, comparing outcomes each time. The platform supports side-by-side comparison of scenarios so compliance officers can **“compare multiple iterations showing cases created, transactions hit, and users affected to calibrate thresholds with confidence”**. This systematic approach ensures that rules are optimized for accuracy and efficiency before they ever impact live monitoring. It's a proactive strategy to prevent misconfigurations and oversights. Using historical data for back-testing makes it much easier to understand how new rules will affect the detection of unusual activity and the frequency of false positives/negatives. In short, Rule Simulation gives teams a safe sandbox to validate that a rule will perform as intended, thereby preventing compliance blind spots and false alarm floods in production.

Shadow Rules

Safe Real-Time Testing in Production



Even with thorough offline simulation, there’s no substitute for observing a rule’s behavior on live data. That’s where Shadow Rules come in. A Shadow Rule in Flagright is essentially a rule run in “ghost mode”; it evaluates real-time production transactions just like a normal rule, but without triggering any actual alerts or operational actions. The new rule’s hits are logged for analysis, but they do not create cases in the live queue or block transactions. This allows compliance teams to safely test new rules in real-time, in parallel with active rules, without any risk of disrupting business or overwhelming analysts.

How it works


When a new rule (or an adjusted version of an existing rule) is set as a Shadow Rule, it is deployed to the production monitoring engine but marked as “shadow rule.” The system runs the rule logic on every incoming transaction (or event) just like a live rule, and it records whenever the rule condition is met (a “hit”). However, these hits remain invisible to the main alert workflow; no alerts are raised to investigators and no customer impact occurs. Instead, the platform compiles hit statistics and detailed logs for the Shadow Rule, which the team can review through the analytics interface. In essence, the compliance team gets to see exactly how the rule would have performed if it were live, how many alerts it would have generated, which transactions/customers would have been flagged, etc., all without any operational consequence.

Flagright designed Shadow Rules to eliminate the fear of deploying untested logic. It provides a risk-free environment to test and refine your configuration, running parallel to your live rules but not generating actual cases or alerts, thereby ensuring that compliance strategies can be tuned with precision without operational risks. This has enormous practical benefits:



✓ Detecting misconfigurations early

Misconfigured rules can be costly, leading to floods of false positives or missed alerts. Shadow mode catches these issues early.



In fact, since launching this feature, Flagright observed an **80% reduction in false positive alerts** due to rule misconfigurations, because errors are caught and fixed in shadow before going live. Clients typically test ~3 versions of a rule in shadow mode, fine-tuning thresholds, velocity limits, and target segments, before finalizing a live rule. This iterative testing greatly improves rule quality.

✓ Continuous optimization of existing rules

Shadow Rules aren't only for brand new scenarios, they can also be used to safely trial adjustments to existing live rules. For example, if analysts suspect a certain threshold is too low (causing too many false positives), they can clone the live rule, raise the threshold in the clone, and run it as a shadow rule for a week. During that time, they can observe how alert volume would change with the higher threshold without actually missing any alerts, since the original live rule is still running in parallel. This way, teams empirically validate that a change will reduce noise (or conversely, won't omit true hits) before applying it. The shadow mode is ideal for adjusting thresholds to reduce false positives and seeing the effect before committing the change.

✓ Adapting to regulatory changes and new typologies

When new regulations or emerging typologies require additional rules or updated logic, Shadow Rules provide a safety net for compliance assurance. Teams can implement the required rule changes in shadow first, confirming that they meet regulatory standards and work as expected on the institution's data. Only after validating do they promote the rules to active status, thereby seamlessly integrating new compliance measures without hiccups.

By running rules “live” but invisible, compliance officers gain total confidence in new detection logic. They can say goodbye to the fear of misconfigurations and never misconfigure a rule again thanks to this capability.

Some practical applications of Shadow Rules include



New rule testing

Deploy any new AML scenario as a shadow rule initially. This lets the team observe its performance on real transactions without triggering alerts, ideal for vetting novel typologies or complex logic.



Rule optimization

Use shadow mode to trial tweaks to existing rules (e.g. higher monetary thresholds, adding a geographic filter) to see if false positives drop. All adjustments can be evaluated in shadow first, ensuring the changes truly improve performance before affecting live monitoring.




Regulatory updates

When a regulation mandates changes (say a new sanction list or behavior to flag), run the updated rule in shadow to confirm it catches the intended cases and doesn't conflict with other rules. This proactive testing in real-time guarantees compliance updates can be rolled out confidently and smoothly.

By leveraging Shadow Rules in these ways, compliance teams massively reduce risk. They no longer have to “flip the switch” on an unproven rule and hope for the best. Shadow mode provides unprecedented assurance that when a rule finally goes live, it's already a well-oiled component of the monitoring strategy.

Rule Analytics and Observability

Performance Metrics & Monitoring



Flagright not only provides powerful tools to test and deploy rules, but also to continuously monitor and analyze rule performance once in production. Built-in Rule Analytics give compliance teams full observability into how each rule is behaving: how often it triggers, how many of those alerts turn out to be false positives, and how these metrics trend over time. This level of transparency is crucial for data-driven tuning and for demonstrating the efficacy of the monitoring program to auditors or management.

Every alert generated in Flagright is tied back to the rule(s) that triggered it, and the platform aggregates rich statistics at the rule level. Some of the key performance metrics tracked per rule include:




Trigger frequency

How frequently the rule fires alerts in a given period (e.g. alerts per day/week). Consistently high trigger counts may indicate an overly sensitive rule or a surge in the underlying activity being monitored. By tracking alert volumes over time, the team can spot trends, for instance, if a rule's alerts are rising week over week, or if a seasonal pattern emerges. This helps in deciding if a rule needs adjustment or further investigation.



False positive rate

What percentage of the rule's alerts are ultimately deemed "false alarms." Flagright's case management and QA integrations allow teams to label or resolve alerts, and thus measure how many were true issues vs. cleared as benign. A high false-positive rate (e.g. 80%+) for a rule is a red flag that the rule's logic or threshold may be too broad. With rule-level false positive metrics, compliance officers can prioritize which scenarios to refine first (focusing on those causing the most noise).



Alert outcome trends over time

The platform can display how a rule's performance changes over time, for example, a time series of monthly alert counts and the percentage that were escalated or SAR-filed vs. closed as false. This trend analysis is vital for ongoing tuning. If a rule's false positive rate was 30% last quarter but has crept up to 50% this quarter, it signals a degradation in effectiveness that warrants investigation. Trends can also reflect external changes (maybe a new fraud pattern causing more hits on a rule) or the impact of tuning changes made by the team. Having this historical lens supports proactive management rather than set-and-forget.

Flagright's approach to observability means these insights are readily accessible via dashboards and reports, without manual data crunching. The platform essentially provides a feedback loop: it not only executes rules but measures their outcomes and lets you drill into each rule's track record. Compliance teams can thus answer questions like: "How many alerts has Rule X generated this month versus last? What percentage were false positives?"

Is Rule X catching any truly suspicious cases or just noise?” Armed with this information, they can then make informed decisions to tweak or retire rules. In many legacy systems, such analysis might require exporting data to spreadsheets or involve significant engineering support. Flagright bakes it into the UI as part of “comprehensive observability” for the rules engine.

Crucially, these analytics are tied into the continuous tuning workflow. Detecting that a rule is firing too often on benign events or trending in the wrong direction is only half the battle, Flagright makes it easy to act on those findings. If analytics show a rule’s alert volume spiked due to false positives, a compliance officer can quickly adjust the rule’s parameters (no code needed) and immediately monitor the effect. Thanks to the **version control and comment logging** in the platform, every change is tracked and auditable, and if needed the team can even run the adjusted rule in Shadow mode to validate improvements before fully deploying the change. **The result is a virtuous cycle of measure → adjust → measure, which keeps the detection program lean and effective.** In practice, a team might notice from the dashboard that “Rule 15” has an unusually high hit count with trivial cases; they decide to refine its scope (e.g. exclude low-risk customer segments) and document this change. The platform logs that update (with the user’s note on why it was made) and the subsequent metrics will reflect whether the false positive rate drops as expected. This kind of data-driven rule management greatly reduces guesswork and enables safe, continuous tuning of scenarios without waiting on engineers or external modelers.

Rules > My rules

JSON ☐ Simulator ☐

My rules Templates AI detection

Live rules Shadow rules

Download Refresh Settings


ID	Name	Description	Hit rate	Queue	Creator
P1 R-1 (0afc5eba)	First payment of a Customer	First transaction of a user	0.00% No hits	Team 2	12:37 P Jun 26,
P1 R-1 (26d580b8)	First payment of a Customer	First transaction of a user	4.55%	default	1:41 PM Jun 26,
P1 R-1 (2d619895)	First transaction of a user.	First transaction of a user.	0.00% Rule not run	default	6:44 AM Dec 7, 2
P1 R-1 (c89ff105)	First payment of a Customer	First transaction of a user	0.00% Rule not run	default	11:46 A Oct 4, 2

Another aspect of Flagright’s rule analytics is support for auditability and reporting. All rule configurations and changes are stored with a historical audit trail. If an auditor or internal compliance committee asks “why was the threshold for Rule Y increased last month?”, the system can show the change record and the rationale that was entered. Moreover, testing activities like simulations or shadow deployments can be documented and even exported. Flagright encourages users to treat rule management as an iterative, well-documented process, for example, “ensure that all testing procedures are properly recorded and compliant with internal and external auditing requirements,” emphasizing that such documentation is critical for regulatory purposes and continuous improvement.

The platform itself facilitates this by logging each test and change. At any point, compliance leads can generate reports of key rule KPIs (like total alerts, clearance rates, etc.) and show the history of modifications. This level of governance and transparency helps make the compliance program adaptable (able to change quickly as needed) but also auditable (every change can be explained and justified), a balance that regulators increasingly expect. In summary, Rule Analytics & Observability in Flagright transform raw monitoring data into actionable intelligence and give teams the tools to track performance, adjust proactively, and prove the effectiveness of their controls over time.

AI-Powered Rule Threshold Recommender

Optimizing Thresholds



Selecting the right threshold values for rules (e.g. transaction amount limits, velocity counts, percentage deviations, etc.) is often one of the trickiest parts of rule tuning. Too low of a threshold can trigger excessive false positives, while too high of a threshold might let suspicious behavior go unflagged. Traditionally, finding the “sweet spot” is a manual, iterative effort relying on analyst intuition or lengthy trial-and-error. Flagright streamlines this process with an AI-driven Rule Threshold Recommender, essentially an intelligent assistant that analyzes historical data and prior alert outcomes to suggest optimized threshold settings for each rule.

What it does

The Rule Threshold Recommender uses machine learning models (trained on both the institution’s data and broader patterns seen across Flagright’s client base) to evaluate how different threshold values would impact a rule’s performance. It then provides the user with data-backed suggestions for threshold adjustments that could improve efficiency without sacrificing effectiveness. In practice, a recommendation might look like: “Rule Z is currently triggering 100 alerts/day with a false-positive rate of 85%. Based on historical alert data, increasing the threshold from \$1,000 to \$1,500 is projected to cut false positives by ~40% while still capturing 98% of true hits.” The platform would present this kind of insight, allowing the analyst to accept the suggestion or further investigate.



How it works under the hood

The AI considers a variety of factors when generating a threshold recommendation. These include each **rule's historical metrics**, for example, the false positive rate of the rule, how many transactions and users it typically flags, and the distribution of values around the current threshold. By looking at where clusters of false vs. true alerts occur relative to the threshold, the AI can infer if a higher or lower cutoff might be better. It might also **weigh the rule's purpose and context**; e.g., if it's a rule meant to catch structuring (smurfing) transactions, the AI might **analyze patterns of deposits** just under the current limit and suggest a new limit accordingly. Essentially, it is performing a form of statistical analysis and simulation behind the scenes: testing various threshold scenarios on the historical data to see which configuration yields the optimal balance of catching true suspicious events while cutting out noise. All of this is done automatically, saving the team from manually slicing and dicing data.



Manual threshold tuning

- Alert noise increases over time
- Manual analysis of historical alerts
- Threshold set through guesswork
- Changes tested and results awaited
- Revisited weeks later



AI-powered threshold recommendation

- Rule performance monitored continuously
- Automated analysis of historical outcomes
- Data-backed threshold recommendation generated
- Analyst reviews and approves changes
- Continuously optimised

Flagright's threshold recommender goes a step further by incorporating feedback and outcomes. For instance, if analysts consistently mark certain alerts as false (or true), that information feeds the model. Over time, the recommendations can become more refined as the system learns from the institution's specific environment. The AI suggestions are presented with explainable analytics, the platform doesn't just say "change this to X", it provides the rationale (e.g. showing the data distribution or trend that led to the suggestion). This transparency is key for analyst trust and for audit purposes. It's worth noting that the compliance team remains in control: recommendations are just that, recommendations. Analysts can review and decide whether to apply the new threshold, perhaps after a quick simulation or shadow run to double-confirm the impact.

The benefit of the Threshold Recommender is **significant time and efficiency gains**. Instead of periodically guessing and checking threshold values, teams have a data-science assistant continuously watching rule performance and nudging them with improvement ideas. This helps minimize false positives proactively. For example, if a certain rule's optimal threshold drifts over time due to changing customer behavior, the AI can catch that drift and alert the team before the rule becomes grossly inefficient. By leveraging AI on historical data, the platform essentially automates a lot of the heavy analytical lifting that a quant or risk analyst might otherwise do in Excel, and it does it continuously and at scale.

In summary, the Rule Threshold Recommender improves detection quality by making rules more precise. It reduces noise (fewer trivial alerts once thresholds are optimized) and raises confidence that thresholds are set at the right level to catch illicit activity. This feature exemplifies Flagright's AI-native philosophy: using machine intelligence to augment human decision-making in compliance. The outcome is a monitoring system that dynamically adapts to data, stays efficient in minimizing false positives, yet remains effective at catching true risks, all with minimal manual effort from the compliance team.

By comparing your institution's rule set and risk profile to this collective intelligence, the AI can highlight gaps. For example, it might notice: "You don't have any rule addressing rapid peer-to-peer payments after account opening, but similar institutions often monitor that pattern for mule account activity." It would then suggest a rule (or even automatically generate a draft rule) to cover this pattern, tailored to your data. Another example: if regulatory alerts have come out about a new scam, the system might proactively recommend a rule to catch hallmarks of that scam.

The recommendations are typically presented in a user-friendly way; e.g. a dashboard might list “Suggested New Rule: Crypto exchange deposits followed by immediate withdrawals to different account” with an explanation of why it’s relevant (perhaps referencing an observed instance in your data or a known typology). The compliance team can review these suggestions and choose to implement them, modify them, or ignore them. This saves teams from having to constantly research emerging threats or guess what they might be missing. It’s like having a knowledgeable consultant who continuously reviews your program against a library of risks.

Rules > Recommendations

1

Basic details

Provide your business & transaction details to help us suggest relevant rules and recommendations

✓

Recommended rules

Select the rules from the recommendations that you want to configure.

✓

Rule

Description

Typology

Type

R-4

✓

Crypto exchange deposits followed by rapid withdrawals

Deposits to a crypto exchange followed by immediate withdrawals to a different account or destination.

Layering behaviour

New activity

View

From a compliance adaptability perspective, the Rule Recommender is extremely powerful. It helps ensure your monitoring program is **not static or lagging behind new schemes**. Instead, it keeps evolving. Particularly for smaller teams or those without in-house data science, this AI feature is a force multiplier, it brings in cutting-edge knowledge and pattern recognition that would be hard to develop alone. It also aligns with regulatory expectations around risk-based monitoring: regulators expect institutions to adjust their controls as risks evolve, and here the system is actively highlighting where an adjustment (new rule) might be warranted.

Flagright emphasizes that these recommendations are **grounded in data and context, not generic one-size-fits-all rules**. Because it considers your transaction patterns, alert history, and even SAR filings, the suggestions are relevant to your specific risk landscape. For example, if the AI sees that analysts have filed several SARs related to a particular novel fraud case, it might suggest a new rule to automatically flag similar cases in the future. Or if your alert closure data shows many false negatives of a certain type (incidents caught outside the system), the AI will learn from that and propose a rule to catch them next time. The inclusion of “regulatory updates” in the input means if, say, FinCEN or FATF issues guidance on a new money laundering method, that knowledge feeds the AI’s suggestion engine. This way, compliance teams can feel confident they are covering all bases and not missing major typologies that they should be monitoring.

In practice, adopting a suggested rule can be as simple as **one-click import** followed by any customization needed. Since Flagright's Rule Builder is no-code, even the suggested rules can often be tweaked quickly (e.g. adjusting a risk threshold or adding an exclusion) and then tested via simulation or shadow mode. Many users leverage the Rule Recommender as a creative brainstorming tool as well, it might surface ideas that the team hadn't considered. Even if not every suggestion is taken, it stimulates a review of the coverage. Over time, this leads to a more robust and refined rule library, with fewer gaps for criminals to exploit.

What rule recommendations delivers



New rule ideas surfaced

Relevant rule suggestions from observed patterns and typologies.



Industry-informed insights

Rules informed by internal data and ecosystem risks.



Confident rule adoption

Clear rationale enables safe review, testing, and approval.



Less manual research

Reduced need to track emerging risks manually.



Coverage gaps identified early

Detection gaps flagged before they are exploited.



Proactive risk coverage

Rules suggested ahead of new threat patterns.

To sum up, the Rule Recommender ensures that a compliance program is dynamic and forward-looking. It identifies gaps by analyzing both patterns in your data and known threats in the wider financial ecosystem, and then suggests concrete rules to address those gaps. This AI-driven guidance means your monitoring isn't limited to what you already know or have experienced, it can preemptively adapt to new risks. In conjunction with simulations and shadow testing, teams can safely implement these suggestions and stay one step ahead of financial criminals. The end result is greater detection coverage (you're more likely to catch novel schemes) with confidence that you are following industry best practices and regulatory insights in your rule design.

Dynamic Rules

Unsupervised Detection for Emerging and Hidden Risks

»

Rules > Templates

JSON

⌵ Simulator

My rules

Templates

🔍

Search for any rule or use case using natural language

⌵

Create scenario

⌵ ⚙️

ID	Rule type	Description	Type	Default nature	Typology	Default action	Action		
Name									
R-25	Dynamic rules	Average transaction value spike	Transaction	Check for spikes in average transaction amounts	Anomaly detection	FRAUD	Unusual behaviour	Flag	Configure
R-26	Dynamic rules	Round-amount anomaly	Transaction	Same user ID performs an abnormally high percentage of round-value transactions ending in 00.00 (without cents) compared to their usual historical behavior, indicating anomalous transaction patterns.	Anomaly detection	FRAUD	Structuring, Unusual behaviour	Flag	Configure
R-27		Total transaction volume exceed past period total transaction volume	Transaction	The total volume of transactions of a user in the last t1 days, is >= X times higher than total volume of transactions in t2 days	Volume comparison	FRAUD	Money mules, Acquiring fraud, Layering	Flag	Configure
R-30	Recommended	High velocity user	Transaction	If a user makes >= x transactions within time t	Velocity	FRAUD	Money mules, Acquiring fraud, Layering	Flag	Configure
R-33		User inactivity	User ongoing screening	Check if a user has been inactive for x days	Anomaly detection/Screening	FRAUD	Unusual behaviour	Flag	Configure
R-41									

Rules, simulations, observability, and threshold optimization form the foundation of an effective monitoring program. However, not all suspicious activity fits neatly into predefined scenarios or static thresholds. Some risks emerge gradually, some are designed to resemble normal behavior, and others fall outside patterns that have historically been labeled or investigated. Dynamic Rules address this gap by adding an unsupervised layer of detection.

Dynamic Rules are Flagright’s unsupervised detection capabilities. Rather than relying on manually defined logic or historical labels, they identify anomalous behavior by learning what is normal within a given population and surfacing meaningful deviations from that baseline. This enables teams to detect emerging risks, previously unseen patterns, and subtle outliers that may not be captured by traditional monitoring logic alone.

Dynamic Rules are not a replacement for static rules, threshold-based scenarios, or supervised models where those are available. They are designed to complement them. In practice, this means institutions can continue using highly explainable rules for known risks while also introducing adaptive detection for unknown or fast-evolving patterns. The result is broader coverage, fewer blind spots, and a more resilient monitoring framework.

What it does

Dynamic Rules continuously assess behavior across relevant dimensions and identify patterns that deviate from expected norms. Depending on the use case, this may include unusual transaction activity, abrupt shifts in customer behavior, atypical combinations of profile attributes, unexpected changes in account usage, or activity that appears inconsistent with a user’s historical baseline.

Unlike traditional rules, Dynamic Rules do not depend on a single fixed cutoff. They evaluate behavior in context. An activity level that is normal for one segment may be highly unusual for another. A sudden increase in volume may be expected for one type of customer but anomalous for a different one. This context-aware approach makes unsupervised detection especially useful for identifying genuine outliers without requiring an ever-growing library of rigid rules.

The output remains operationally practical and explainable. When activity is flagged, the system surfaces the underlying signals and contributing factors so that investigators can understand why the case stands out and where to focus their review.

How it works under the hood

Dynamic Rules use unsupervised models to identify outliers across selected groups of features. These models are designed to detect meaningful deviations without requiring prior investigation outcomes. In practice, Flagright can evaluate multiple anomaly lenses independently, for example across transactional behavior, customer profile characteristics, account activity patterns, or other available signals.

Each Dynamic Rule is treated as an independent detection signal rather than being forced into a single opaque decision. This is important because suspicious behavior in one dimension should still be surfaced even if another dimension appears normal. By preserving these signals independently, the platform gives compliance and risk teams more transparency, more control, and more useful investigative context.

As with any detection logic, calibration is essential. Unsupervised methods can be introduced quickly, but their effectiveness depends on tuning them against real data and validating whether the anomalies being surfaced are meaningful in practice. For that reason, Dynamic Rules fit naturally into Flagright's broader workflow of simulation, shadow testing, analytics, and continuous refinement.

Why it matters

Dynamic Rules help institutions identify the activity that predefined logic is most likely to miss. This includes emerging typologies, behavioral shifts, low-frequency outliers, and cases that have never previously been encoded into rules or reflected in historical labels. Because the models are not trained to reproduce past analyst decisions, they can also add a less biased perspective to the monitoring program and expand detection beyond the institution's current control set.

This is especially valuable in fast-changing environments where risk patterns evolve quickly and monitoring programs must adapt without waiting for every new pattern to be manually translated into a rule. Dynamic Rules provide an additional layer of coverage that can surface unusual activity earlier, giving teams the opportunity to review and respond before a pattern becomes widespread.

Used effectively, they also strengthen the broader control framework. Outcomes from Dynamic Rules can inform future threshold changes, inspire new rule creation, and highlight where existing monitoring may be too narrow or too dependent on known patterns.

Dynamic Rules and dynamic thresholds

Dynamic Rules are distinct from threshold optimization, but the two are highly complementary. Threshold-based rules remain essential because they are transparent, auditable, and easy to align to known regulatory and operational requirements. At the same time, fixed thresholds can become stale as behavior changes over time.

Flagright addresses this from two angles. First, the Rule Threshold Recommender helps teams optimize explicit thresholds based on historical performance and alert outcomes. Second, Dynamic Rules reduce reliance on hard-coded cutoffs by evaluating behavior in relative, context-aware terms. Together, these capabilities create a more adaptive monitoring framework: threshold-based controls continue to govern known risks, while unsupervised detection identifies anomalies that fall outside predefined limits. The outcome is a stronger hybrid approach, one that combines precision for known scenarios with adaptability for emerging and less obvious risks.

Operational safeguards and tuning

The objective of Dynamic Rules is not simply to identify as many outliers as possible, but to surface the right anomalies at an operationally manageable volume. For that reason, Dynamic Rules should be introduced with the same disciplined approach applied to other monitoring controls: tested on historical data where possible, observed in shadow mode, measured through analytics, and refined using investigation outcomes.



This staged rollout is critical for controlling false positives. An unsupervised model may detect many statistical outliers, but not all of them will be relevant from a compliance or risk perspective. By using simulation, shadow deployments, and performance monitoring, teams can calibrate Dynamic Rules to produce useful signal rather than excess noise.

In practice, institutions can start with a focused set of Dynamic Rules, evaluate hit rates and case outcomes, and gradually expand coverage as confidence increases. This allows unsupervised detection to strengthen the program without overwhelming investigators or disrupting existing workflows.

Practical applications of Dynamic Rules include



Emerging risk detection

Surface suspicious patterns that are not yet covered by existing rules, including new or evolving typologies that have not been formally encoded into the monitoring framework.



Coverage gap discovery

Reveal areas where traditional rules may be too narrow or where suspicious activity is occurring without triggering existing controls.



Behavior shift monitoring

Identify users or accounts whose activity changes materially over time in ways that are inconsistent with their historical baseline.



Independent signal generation

Provide additional, explainable signals alongside static rules and supervised models, allowing investigators to evaluate activity from multiple perspectives.

What Dynamic Rules deliver



Broader detection coverage

Captures behaviors that fall outside predefined rules and thresholds.



Earlier visibility into new risks

Surfaces unusual patterns before they are formally encoded into monitoring scenarios.



Reduced dependence on historical decisions

Adds detection capability without inheriting all of the bias of past decisioning.



More adaptive monitoring

Responds to changing patterns and evolving behavior over time.



Explainable anomaly signals

Gives investigators visibility into the factors that contributed to a flag.

Conclusion

A Data-Driven, Confident Compliance Workflow

Flagright's approach to rule performance, observability, and analytics represents a modern, proactive paradigm for financial crime compliance. By combining no-code flexibility with AI intelligence and robust testing tools, the platform allows compliance professionals to design and maintain a detection program that is both highly effective and highly efficient. Let's recap how the pieces fit together and the tangible benefits they provide:

Simulation & Shadow

Test before you invest

With Rule Simulation and Shadow Rules, teams can fine-tune detection logic using real data (historical and live) in a safe environment. This dramatically reduces false positives and errors upfront. Misconfigured or poorly calibrated rules, a common source of noise and regulatory risk, are identified and fixed through iterative testing. The result is that when rules go live, they are already vetted to be precise, leading to fewer needless alerts and more focus on true threats. Additionally, this capability means no engineering deployment bottlenecks; compliance officers can adjust and validate scenarios within days, not months, which keeps the program agile.

Analytics

Always-on performance monitoring

Flagright builds observability into the rules engine, so the work doesn't stop at deployment. Each rule's performance is continuously measured (alerts triggered, false positive rate, trendlines, etc.), allowing the team to spot issues and improvements in real time. Rather than waiting for a quarterly review or an audit finding, compliance teams have live dashboards and reports to track what's working and what isn't. This leads to a culture of continuous improvement, rules can be tweaked the moment they start underperforming. Moreover, the audit trail and versioning mean the program is fully transparent and defensible. Changes are documented with rationale, and one can always trace back why a threshold was set a certain way or how a rule evolved. In an industry where regulators demand both effectiveness and accountability, this level of governance is invaluable.

Recommendations

Augment human expertise with AI

The Rule Threshold Recommender and Rule Recommender illustrate Flagright's AI-native advantage. They turn vast amounts of data and collective intelligence into actionable suggestions for optimization and risk coverage. This not only saves time (by automating complex analysis), but also often uncovers improvements a human might miss. By following AI-driven threshold tweaks, institutions have seen significant drops in false alerts without missing true cases. And by considering AI-suggested new rules, firms ensure they're not blind to emerging typologies or changing patterns. Essentially, AI acts as a force multiplier for the compliance team, doing the heavy analytics and cross-institution learning in the background, so that the team can make informed decisions quickly. All of this is done with the compliance officer in the loop, preserving the explainability and control needed in a regulated environment.

Dynamic Rules

Uncover hidden risks early

Dynamic Rules extend detection beyond predefined logic by identifying anomalies and behavioral shifts that fall outside static rules and thresholds. By learning what is normal across different segments, they surface meaningful deviations in context rather than relying on fixed cutoffs. This enables teams to detect emerging risks, subtle outliers, and previously unseen patterns that traditional monitoring may miss. Each signal remains explainable, giving investigators clear visibility into why an activity stands out. When combined with simulation, shadow testing, and analytics, Dynamic Rules can be introduced in a controlled and measurable way. Over time, they strengthen the monitoring program by reducing blind spots and improving overall coverage.

Empowerment and agility

Underpinning Flagright's approach is a fundamental empowerment of the compliance function. Rules can be created, tested, adjusted, or retired without writing code and without lengthy deployment cycles. This means compliance programs can adapt at the pace of evolving risk; a new fraud trend can be addressed with a rule the same day it's discovered, and that rule can be fine-tuned by the end of the week through simulations and shadow testing. The dependence on engineering or vendor support is minimized, freeing teams to be more creative and responsive. Flagright's platform essentially treats rule management as an ongoing, agile process rather than a static configuration. That is a huge shift from legacy systems, and it directly translates to better risk coverage and efficiency.

In a world where financial criminals constantly innovate and regulators raise the bar for compliance, having such a data-driven, agile rule management capability is a competitive advantage for financial institutions. Flagright's comprehensive suite; from simulation sandboxes to real-time shadowing, from rich analytics to AI recommendations, ensures that compliance teams can continuously calibrate their defenses with confidence. The outcome is a dramatically lower noise level (false positives cut by orders of magnitude in some cases), a higher detection rate of truly suspicious activity, and a well-documented, auditable program that can stand up to scrutiny.

By adopting Flagright's approach to rule performance and observability, financial institutions can turn what was once a cumbersome, error-prone process into a streamlined cycle of test, observe, improve; ultimately staying one step ahead of financial crime while operating with maximum efficiency and assurance.



**See what intelligent rule management
looks like with Flagright**



Scan to book
a personalized demo