sublime

January 2026

# Sublime Security
# AI Overview

# Applying AI Thoughtfully

1

## Controllable

You define the guardrails – turn capabilities on or off instantly

2

## Transparent

Every decision is explainable in natural language

3

## Trustworthy

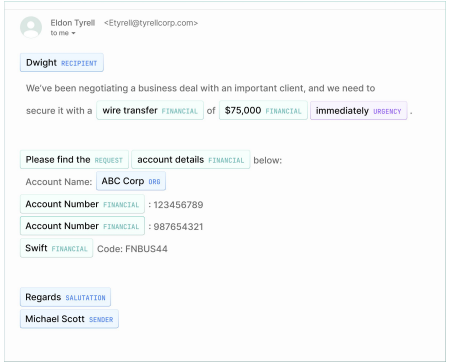
Auditable verification of agent actions and logic

# Sublime’s Layered AI Makes Detections Resilient

Sublime AI can coordinate extensive AI experience through MQL, providing our agents with a common language to collaborate in.

## ml.nlu_classifier

Sublime’s Natural Language Understanding leverages a local LLM model capable of detecting the intent of a message, key topics, and important context like urgency and tone.




```
(
  any(body.previous_threads,
      any(ml.nlu_classifier(.text).intents, .name == "bec")
  )
  or any(ml.nlu_classifier(sender.display_name).intents, .name == "bec")
  or any(body.previous_threads,
         (
           any(ml.nlu_classifier(.text).entities, .name == "urgency")
           and any(ml.nlu_classifier(.text).entities, .name == "request")
         )
  )
)
```

## ml.logo_detect

## profile.by_sender

© 2026 Sublime Security, Inc.

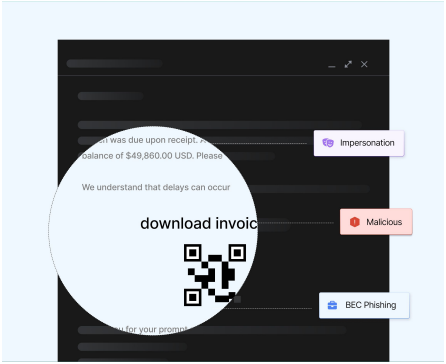
# Sublime's Layered AI Makes Detections Resilient

Sublime AI can coordinate extensive AI experience through MQL, providing our agents with a common language to collaborate in.

## ml.nlu_classifier

## ml.logo_detect

If the user sees a Microsoft logo or a QR Code, the rule sees it too. Sublime computer vision models analyzes messages for brand misuse and anomalous login pages.




```
// conditions for QR code via text
any(file.explode(.),
    any([.scan.strings.raw, .scan.ocr.raw],
        regex.icontains(., 'scan|camera')
        and regex.icontains(., '\bQR\b|Q\.R\.|barcode')
    )
)
or any(file.explode(.),
        .scan.qr.type == "url" and .scan.qr.url.domain.valid
)
// conditions for text
and any(file.explode(.),
        // review/change terms in file content
        any([.scan.strings.raw, .scan.ocr.raw],
            (
                regex.icontains(.,
```

## profile.by_sender

© 2026 Sublime Security, Inc.

# Sublime's Layered AI Makes Detections Resilient

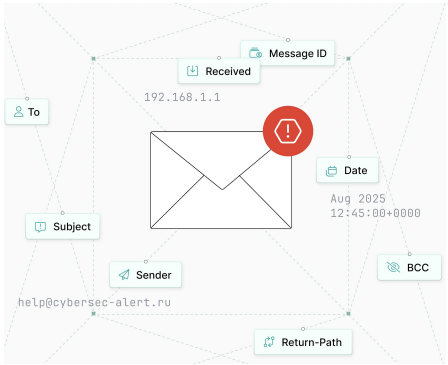
Sublime AI can coordinate extensive AI experience through MQL, providing our agents with a common language to collaborate in.

## ml.nlu_classifier

## ml.logo_detect

## profile.by_sender

Sender Profiling surfaces Sender/Recipient relationship anomalies, without an overreliance on domain reputation, asking questions like: *Is this normal behavior?* Or *Have I seen this sender before?*




```
type.inbound
and any(attachments,
        .file_extension in~ ('htm', 'html') or .file_type == "html"
)
and (
  (
    profile.by_sender().prevalence in ("new", "outlier")
    and not profile.by_sender().solicited
  )
  or profile.by_sender().any_messages_malicious_or_spam
)
and not profile.by_sender().any_messages_benign

// negate highly trusted sender domains unless they fail DMARC authentication
and not (
  sender.email.domain.root_domain in $high_trust_sender_root_domains
  and coalesce(headers.auth_summary.dmarc.pass, false)
)

```

© 2026 Sublime Security, Inc.

# Agentic AI

# Meet the Sublime AI Agents

## ASA

### Autonomous Security Analyst

**Mission:** Automate phishing triage.

- Multimodal analysis of complete message
- Access to Sublime Knowledge Base for analysis best practices and labeling
- Operates 24/7 without fatigue
- Instantly remediates threats
- Provides detailed reporting on verdict reasoning

## ADÉ

### Autonomous Detection Engineer

**Mission:** Proactively close detection gaps.

- Identifies attacks that bypassed existing defenses.
- Drafts custom MQL rules.
- Backtests rules against historical data to ensure accuracy.
- Proposes new coverage for approval in hours, not weeks.

COMING SOON

## MIRA

### Metrics & Insights Reporting Assistant

**Mission:** Answer business & security questions with natural language queries.

- Provides direct access to the underlying data that populates the Sublime Reporting Dashboard
- Ad-Hoc queries to understand threat landscape "*Show me threats users being targeted by BEC attacks*"
- Generate custom reports containing dynamic visualizations & human-readable summaries

© 2026 Sublime Security, Inc.

Message Query Language (MQL)

# The Lingua Franca of Sublime's Agentic AI

## 1. The Collaboration Protocol

Agents like ASA and ADÉ don't just output text; they exchange deterministic MQL. This ensures precise handoffs with zero ambiguity.

## 2. Transparency is a Feature, Not a Bug

MQL turns probabilistic intent into deterministic action that you can see, audit, and trust.

## 3. Deterministic Control

If you disagree with an agent's logic, you can tweak the MQL code directly. You aren't fighting a black box; you are collaborating with the AI.



```
type.inbound

// SECTION 1: TrustWallet impersonation with homoglyph detection
and (
  // homoglyph-normalized matching for character substitution attacks
  (
    strings.ilike(strings.replace_confusables(sender.display_name), "*trust*wallet*")
    or strings.ilike(strings.replace_confusables(subject.subject), "*trust*wallet*")
  )
  // OR character substitution patterns (AlRDROP vs AIRDROP, TrustWaIIet vs TrustWallet)
  or (
    regex.icontains(sender.display_name, "a[il]rdrop", "trustwa[il]+et")
    or regex.icontains(subject.subject, "sh[il]b", "a[il]rdrop")
  )
)

// SECTION 2: Not from legitimate TrustWallet domains
and sender.email.domain.root_domain not in ("trustwallet.com", "trustwalletapp.com")

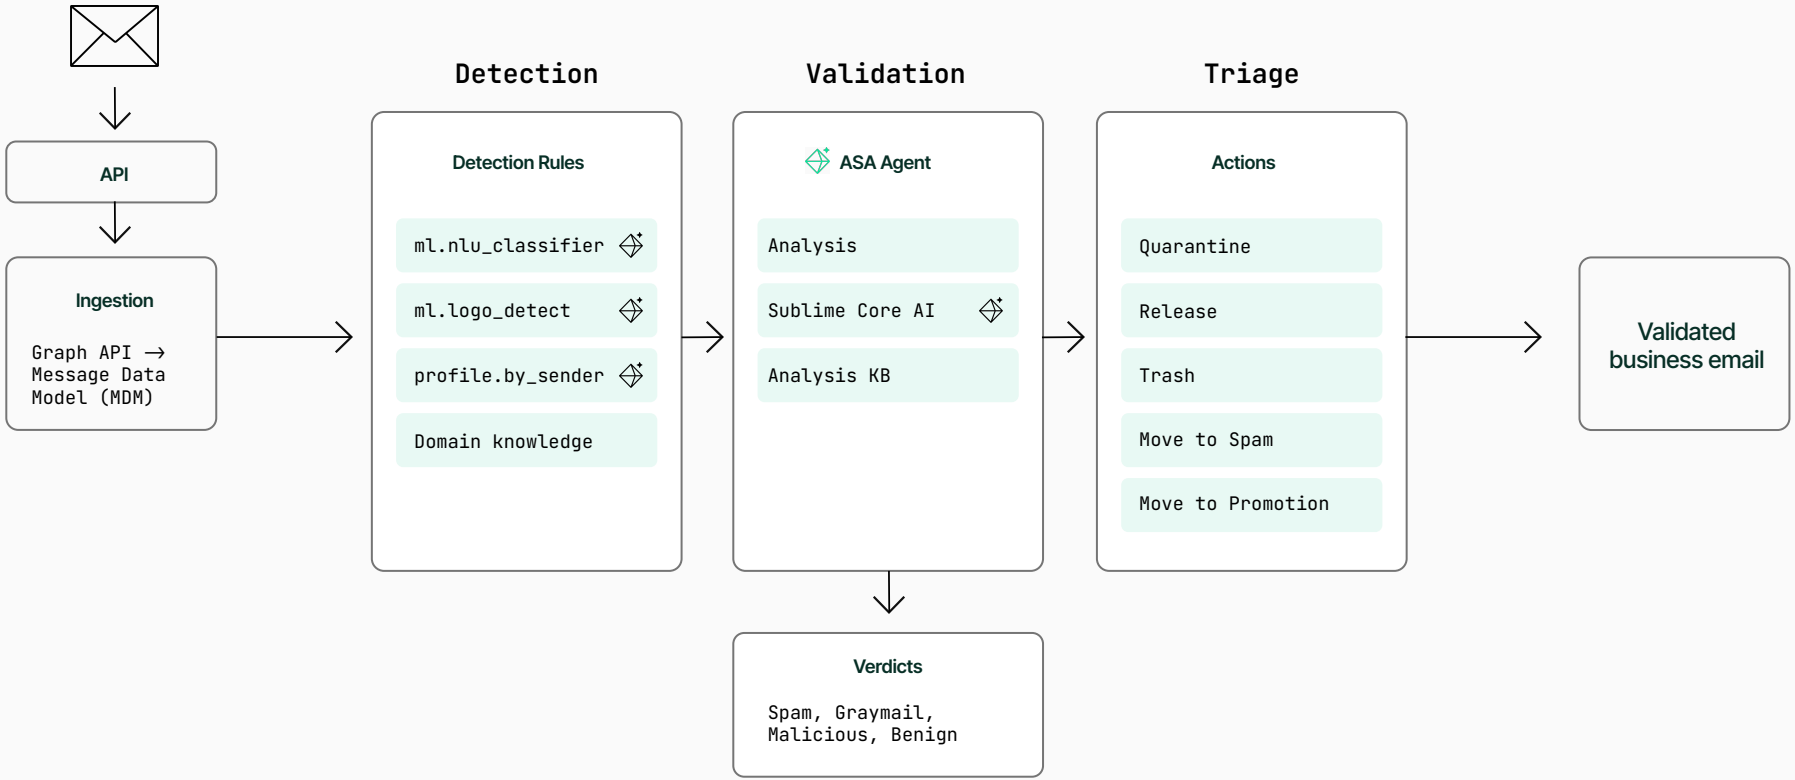
// SECTION 3: Malicious call-to-action links (service abuse compatible)
and any(body.links,
        // Airdrop/wallet-related call-to-action text
        strings.ilike(.display_text,"*get*airdrop*","*connect*wallet*","*get*token*")
        // OR suspicious redirect patterns in legitimate infrastructure
        or (.href_url.domain.root_domain in ("exacttarget.com","salesforce.com","exct.net")
          and strings.ilike(.href_url.url, "*airdrop*","*claim*","*wallet*","*token*")
        )
)

// SECTION 4: Social engineering and urgency indicators
and
// OR urgency/financial entities detected by NLU
any(ml.nlu_classifier(body.current_thread.text).entities,
    .name in ("urgency", "financial")
)

// SECTION 5: Expensive ML operations LAST (conditional execution)
and (
  length(body.current_thread.text) > 50
  and (
    // TrustWallet logo detection
    any(ml.logo_detect(beta.message_screenshot()).brands,
        .name == "TrustWallet" and .confidence in ("medium", "high")
    )
    // OR credential theft / financial fraud intent
    or any(ml.nlu_classifier(body.current_thread.text).intents,
           .name in ("cred_theft", "steal_pii")
           and .confidence in ("medium", "high")
    )
  )
)
```

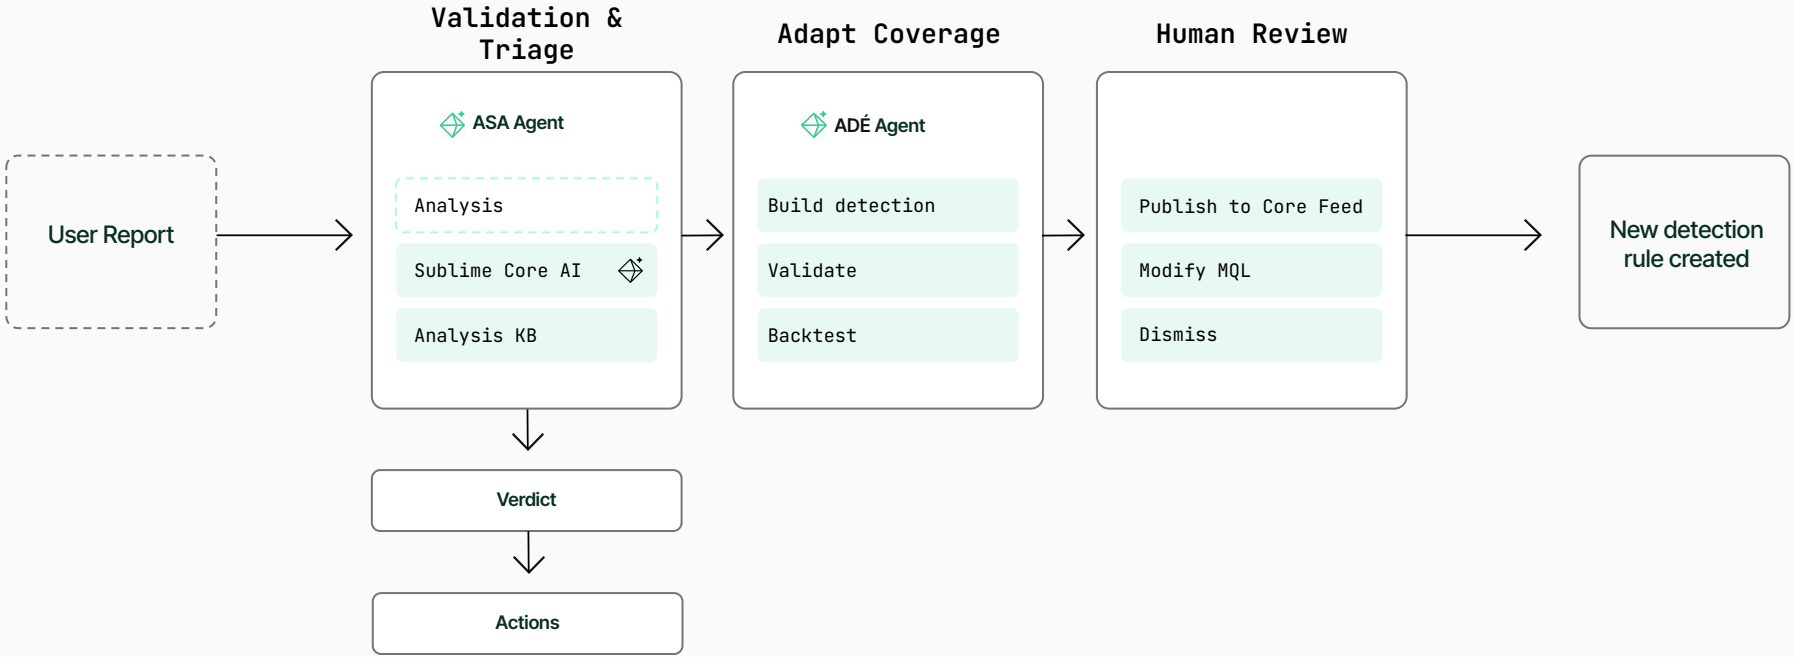
© 2025 Sublime Security, Inc.

# Inbound Detection Flow




© 2026 Sublime Security, Inc.

# User-Report Message Flow




© 2026 Sublime Security, Inc.

A Collaborative Defense

# How MQL Primitives & Agents work together to close gaps

## Detect



### Real-time Filtering with MQL

NLU, Computer Vision, and Anomaly Detection scan inbound messages, flagging threats using deterministic rules.

## Respond



### Automated Analysis & Triage with ASA

When users report suspicious emails, ASA investigates instantly, remediates threats, and labels the data.

## Adapt



### Evolve Security Posture with ADÉ

ADÉ analyzes what bypassed defenses and autonomously writes new MQL rules to close the gap.

© 2026 Sublime Security, Inc.

# Privacy & Trust by Design

1

## No Training on Data

Your sensitive data is yours. We do not train our core AI models on customer data, preventing model leakage.

2

## Foundation Models

Our agents leverage LLMs directly through Amazon Bedrock. No data is shared with frontier model providers.

3

## Opt-In Control

Agents are inactive by default. You decide the level of autonomy, activating specific agents/automations only when ready.

4

## Single-Tenant Isolation

Deploy directly to Sublime Cloud or your own AWS environment. Your data never leaves the controlled boundary, ensuring complete isolation.

5

## Ephemeral Processing

Zero-retention analysis. Transient payloads are purged immediately after inference.

Our agents do not store your raw email content.

6

## Responsible AI Guardrails

Designed for safety with human-in-the-loop validation, continuous efficacy testing, and strict governance controls like MFA and legal review.

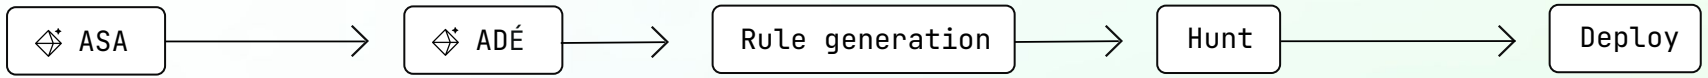
© 2026 Sublime Security, Inc.

sublime

@sublime-sec
https://sublime.security

# Thank you!

# Example: ASA ↔ ADÉ




## User-reported msg

- Message analyzed by ASA
- Verdict & IOCs presented in human-readable report
- All IOCs sourced with MDM locations for ADÉ consumption

## Synthesize Report

- Analyze the IOCs and compare them to existing Detection Rules
- Determines whether to update an existing Rule or create a new one from scratch

## Write & Validate

- ADÉ uses IOCs & a Sublime Detection Engineering KB to author a candidate rule
- Subagents perform analysis of rule brittleness
- Rule is validated by MQL subagent
- Kicks off L14D Retrohunt

## Continuous Testing

- ADÉ performs its own CI/CD for candidate Rules
- ADÉ validates the Rule against historical data
- Refines Rule over time to minimize false positives
- Continues cycle until efficacy thresholds are met

## Close-the-loop

- Presents the final Rule with a detailed explanation
- Require human to review and approve
- Send ADE rule to MQL Editor for final tuning
- Rule → Auto-Quarantine action
- Audit logging for full transparency

© 2026 Sublime Security, Inc.