

TECHNICAL WHITE PAPER



AgenticOps AI

AI Security and Communication Intelligence Platform

A production-ready technical overview of the AgenticOps AI platform, implemented capabilities, safety model, OKX.AI service, and planned roadmap.

Version 1.0

July 2026

agenticopsai.xyz

Executive Summary

AgenticOps AI is an AI Security and Communication Intelligence Platform for teams that need to understand, prioritize, and safely respond to modern communications. It combines deterministic security rules, AI-assisted analysis, explainable risk scoring, confidence scoring, and human approval workflows.

Version 1.0 focuses on two implemented communication contexts: Web3 Community Security and Business Communication Intelligence. The platform also includes document analysis, connected ingestion foundations, a Communication Inbox, a Human Approval Center, event logs, reports, and durable workflow storage.

The product is intentionally safety-first. It analyzes messages and documents, explains risks, recommends actions, and keeps humans in control. It does not autonomously send emails, reply to communities, moderate users, modify inboxes, spend money, or launch campaigns.

Problem Statement

Organizations now operate across fragmented communication surfaces: community chat, email, social messaging, business documents, spreadsheets, and support workflows. Important risks and opportunities are spread across these channels.

- Web3 scams, phishing, fake administrators, and social engineering move quickly through community channels.
- Manual moderation is slow, inconsistent, and difficult to audit.
- Business communication overload hides urgent requests, complaints, leads, and operational risks.
- Documents and spreadsheets contain valuable business signals but are often reviewed separately from communication records.
- Teams lack a single explainable system for message normalization, risk analysis, recommended actions, and review history.

AgenticOps AI Solution

AgenticOps AI provides one shared intelligence engine for communication and document review. Inputs are validated, normalized, assigned a communication context, analyzed by deterministic rules where appropriate, evaluated by AI-assisted reasoning, and presented as structured outputs.

The system supports safe decision support rather than autonomous external execution. Reviewers receive summaries, risk signals, intent, sentiment, priority, recommended actions, suggested reply outlines, confidence, and explanation.

Implemented: Web3 Community Security, Business Communication Intelligence, supported file analysis, connected ingestion foundations, event log, reports, and human approval records.

Future: broader email actions, customer support workflows, marketing intelligence, CRM integrations, advertising intelligence, and AI business operator automation.

Security and Communication Intelligence Engine

The engine is provider-neutral. Source channels do not automatically determine risk. A Telegram message, Discord message, Gmail message, Facebook Messenger event, Instagram foundation event, manual input, or document upload is normalized before analysis.

Context-aware analysis then determines whether the message should be treated as Web3 security, business communication, email intelligence, social communication, or a general communication review.

- Risk: scams, phishing, impersonation, credential requests, malicious links, abuse, and operational risk.
- Intent: support request, complaint, purchase intent, lead, FAQ, feedback, abuse, scam, phishing, or unknown.
- Priority: low, medium, high, or critical based on urgency, risk, and requested action.
- Review: sensitive outputs remain approval-required and non-executing by default.

Deterministic Security Rules

The deterministic rule catalogue originated with the Web3 Community Security MVP. The active catalogue includes SEC-001 through SEC-015 and detects high-confidence security patterns such as seed phrase requests, private-key requests, fake administrator behavior, suspicious links, remote access requests, OTP/password requests, prompt injection, spam, and abuse.

Deterministic rules are stable, reproducible, and auditable. They provide rule IDs, evidence, severity, and recommended action. Business Communication Intelligence also uses structured explainable analysis, but it is not falsely described as applying every Web3 security rule in every business context.

AI-Assisted Analysis

OpenRouter-compatible LLM reasoning adds summarization, sentiment, intent classification, priority assessment, missing-context detection, confidence scoring, suggested next steps, and reply outlines.

AI reasoning is useful for ambiguous communications where deterministic rules alone are not enough. AgenticOps AI uses AI to support human judgment, not replace it.

- Outputs are structured and schema-oriented.
- Errors are sanitized and bounded.
- Risk and confidence are explained in plain language.
- AI recommendations remain subject to human review.

Human Approval Workflow

AgenticOps AI separates received messages, AI analysis, suggested responses, human approval, approved execution, and audit history. Version 1.0 generates suggested replies and recommended actions, but external execution remains unavailable or approval-required.

- No automatic email sending or Gmail modification.
- No automatic Telegram, Discord, Facebook, or Instagram replies.
- No automatic moderation, bans, deletions, ad actions, payment actions, or CRM writeback.
- Approval updates internal review state unless a future provider-specific execution system is explicitly authorized and audited.

Web3 Security Use Cases

- Detect seed phrase, private key, wallet recovery, and credential requests.
- Identify fake administrator and impersonation messages.
- Flag malicious or suspicious links.
- Detect phishing, scams, abuse, spam, and prompt injection attempts.
- Generate safe moderator reply guidance for human review.
- Run batch analysis and produce auditable security reports.

Business Communication Use Cases

- Analyze pasted business messages and supported document uploads.
- Classify customer support, business email, sales conversation, internal team, general communication, business audit, and budget review purposes.
- Extract and analyze TXT, text-based PDF, DOCX, CSV, and XLSX content.
- Generate executive summaries, risk indicators, missing context, recommended actions, and report-ready findings.
- Support preliminary business review, not certified accounting, statutory audit, legal advice, or financial advice.

Integrations

Version 1.0 includes implemented and foundation-level communication integrations. These integrations feed the shared normalized message pipeline and preserve analyze-only or approval-required boundaries.

- Gmail: readonly OAuth connection, manual sync, message normalization, analysis, and suggestions. No send, modify, archive, label, or delete.
- Telegram: webhook ingestion, secret validation, text-message normalization, analysis, and approval-required suggestions.
- Discord: persistent Gateway worker, normalized message handoff, shared analysis pipeline, and durable event workflow.
- Facebook Messenger: webhook verification, signature validation, supported message-event normalization, analysis, and persistence where events are delivered.
- Instagram: verified webhook and supported payload foundation. Broader production event coverage depends on provider configuration, permissions, and delivery.

A2MCP / OKX.AI Service

AgenticOps AI exposes an A2MCP-compatible Communication Intelligence service for external AI systems. The callable tool is `analyze_communication_risk`.

Production endpoints are <https://agenticopsai.xyz/api/mcp>, <https://agenticopsai.xyz/api/okx/analyze>, and <https://agenticopsai.xyz/api/v1/health>.

The service accepts bounded message content, communication context, and source channel. It returns summary, intent, sentiment, priority, risk level, risk signals, recommended action, suggested reply outline, human-review requirement, confidence, and explanation.

The OKX service is analyze-only and does not expose OAuth operations, provider credentials, approval mutation, Gmail synchronization, private integration events, or autonomous execution.

Architecture

The architecture uses a shared validation, normalization, analysis, output, and audit pattern. Provider adapters remain separate from the core analysis engine.



Privacy and Safety

- Provider credentials stay server-side.
- Google OAuth tokens are encrypted at rest when durable storage is configured.
- Gmail uses readonly scope only.
- Inputs and outputs are bounded.
- Webhook signatures and secrets are validated where configured.
- Raw provider payloads and unbounded private message bodies are not stored by default.
- Human review remains the boundary for sensitive recommendations and future external actions.

Roadmap

The following capabilities are planned future work and must not be interpreted as implemented Version 1.0 features.

- Communication channels: X, LinkedIn, Slack, Microsoft Teams, WhatsApp Business, and live chat.
- Social intelligence: broader social comment intelligence, direct-message intelligence, brand monitoring, and reputation analytics.
- Marketing intelligence: social ad management with human approval and AI marketing intelligence.
- Business operations: Gmail send/modify actions with user approval, calendar intelligence, workflow automation, and CRM integrations.
- Enterprise: BigQuery, Google Cloud integrations, business intelligence expansion, organization management, workspaces, billing, and role-based access.
- AI platform: multiple LLM providers, AI agent orchestration, and AI business operator workflows.

Conclusion

AgenticOps AI Version 1.0 is the first production release of an AI Security and Communication Intelligence Platform. It provides practical Web3 security analysis, business communication intelligence, document analysis, integrations, explainable outputs, reports, durable event workflows, and human approval boundaries.

The long-term vision is a broader AI operations layer for communication, business intelligence, audit, compliance, marketing, and workflow automation. The Version 1.0 boundary remains clear: analyze communications, explain risk, recommend actions, and keep humans in control.

- Implemented today: Web3 Community Security, Business Communication Intelligence, supported file analysis, connected ingestion foundations, reports, event logs, and approval-required recommendations.
- Not autonomous: AgenticOps AI does not send messages, modify inboxes, moderate users, launch campaigns, spend money, or change external systems without future explicit authorization and audit controls.
- Roadmap direction: expand communication intelligence into broader email, customer support, marketing, CRM, business data, audit, compliance, and AI operations workflows.
- Production service: agenticopsai.xyz hosts the website, MCP endpoint, analysis endpoint, and health endpoint used by the callable service.
- Operating model: deterministic checks, AI-assisted reasoning, structured outputs, durable event records, and human approval work together to support auditable decisions.
- Review boundary: Version 1.0 is decision-support software and does not claim certified audit, legal advice, financial advice, or autonomous business management.