nylas

SECURITY & COMPLIANCE · BUYER'S CHECKLIST

Enterprise AI vendor evaluation

10 questions for security, privacy, and procurement teams

Version 1.0 · July 2026 · Published by Nylas. For legal obligations specific to your organization, consult legal counsel.

HOW TO USE THIS CHECKLIST

Run this checklist during security reviews, RFP evaluations, or due diligence for any AI-enabled enterprise software vendor. Each question maps to a risk area. Red flags are rated by severity: **Minor** (note and monitor), **Significant** (escalate for review), or **Disqualifying** (reject or require contractual remediation before proceeding).

MINOR SIGNIFICANT DISQUALIFYING

WHO SHOULD OWN EACH SECTION

SECTION	PRIMARY OWNER	SECONDARY
Q1–3 · Data handling	Privacy / Legal	Security
Q4, Q8 · Security controls	Security	IT
Q5–6 · Governance	Procurement	Legal
Q7 · Output controls	Business owner	Security
Q9 · Abuse / misuse	Security	IT
Q10 · Responsibility	Legal	Procurement

01 Is customer data used to train AI models?

Why it matters. Undisclosed training use is both a regulatory exposure (GDPR, EU AI Act) and a trust issue. If your email, calendar, or contact data improves a vendor's general model, that benefit accrues to the vendor — and potentially to competitors using the same model.

ASK

- [ ] Is customer data used to train general-purpose AI models?
- [ ] Is customer-specific data used for model tuning or evaluation?
- [ ] Is explicit customer consent or opt-in required before any training use?
- [ ] Is training use documented in the data processing agreement or MSA?
- [ ] Is there a documented opt-out process?

RED FLAGS

FLAG	SEVERITY
Vague answer: "We take data privacy seriously" without a policy statement	SIGNIFICANT
Marketing language substituted for contractual commitment	SIGNIFICANT
"Subject to change" language in training policy	SIGNIFICANT
Vendor cannot confirm or deny training use	DISQUALIFYING
No contractual clause addressing training	DISQUALIFYING

> **How Nylas answers this.** Nylas does not use customer data to train general-purpose machine learning models. Nylas does not develop or operate proprietary foundation models. This commitment is documented in Nylas's data handling policies and can be confirmed in data processing agreements.

02 Which AI providers or subprocessors are involved?

Why it matters. The vendor you contract with may not be the entity operating the AI that processes your data. Every subprocessor is an additional data flow, an additional breach surface, and a potential compliance obligation under GDPR Article 28 and the EU AI Act's deployer requirements.

ASK

- [ ] Which third-party AI providers process customer data as part of this product?
- [ ] Which subprocessors are involved in AI-enabled features specifically?
- [ ] How and when are customers notified when subprocessors change?
- [ ] What contractual safeguards govern subprocessor data handling?
- [ ] Is a current subprocessor list available?

RED FLAGS

FLAG	SEVERITY
Vendor cannot name its AI providers	DISQUALIFYING
No subprocessor disclosure in documentation	SIGNIFICANT
No customer notification process for subprocessor changes	SIGNIFICANT
Subprocessors not covered by DPA or equivalent contractual terms	DISQUALIFYING

> **How Nylas answers this.** Nylas reviews third-party providers supporting platform operations for security, privacy, operational, and contractual risks. Customers requiring specific subprocessor information should contact Nylas directly. Subprocessor relationships are governed by contractual terms.

03 What data does the AI system actually process?

Why it matters. "We protect all your data" is not an answer. Buyers need to know specifically what categories of data AI features can access — message content, metadata, calendar events, contact records, authentication tokens — and whether processing can be scoped or limited.

ASK

- [ ] What specific data types does the AI system access or process?
- [ ] Does AI access message content, message metadata, calendar data, or contact records?
- [ ] Can data sharing with AI features be limited or scoped by the customer?
- [ ] How long is data retained after AI processing?
- [ ] How is data minimization applied — does the AI access only what a specific function requires?
- [ ] Is metadata treated with the same protections as message content?

RED FLAGS

FLAG	SEVERITY
"We only process what's necessary" without specifics on what that includes	SIGNIFICANT
No retention policy for AI-processed data	SIGNIFICANT
No data minimization principle documented	SIGNIFICANT
Customer cannot limit AI access to specific data types	SIGNIFICANT
Metadata not addressed in data handling documentation	MINOR

> **How Nylas answers this.** Where AI-enabled functionality is present in Nylas products, it may process customer-provided content for purposes such as summarizing information, structuring data, or assisting workflow automation. Nylas applies data minimization principles to AI-related processing. Not every customer uses AI-enabled functionality.

04 How does AI fit into the vendor's existing security program?

Why it matters. AI features that exist outside a vendor's standard security program — separately assessed, separately monitored, separately governed — create gaps. An AI feature with sophisticated model architecture but weak surrounding controls is not a safe enterprise choice.

ASK

- [ ] Is AI-related data processing covered by the vendor's existing encryption standards?
- [ ] Are AI features included in the vendor's access control policies?
- [ ] Is AI activity included in the vendor's security monitoring and logging?
- [ ] Are AI vendors and subprocessors subject to the same vendor security review process as other providers?
- [ ] Is AI included in the vendor's incident response plan?
- [ ] Is the vendor's security certification (SOC 2, ISO 27001) scoped to include AI features?

RED FLAGS

FLAG	SEVERITY
AI treated as a separate product line outside the security program	SIGNIFICANT
No AI-specific monitoring or logging	SIGNIFICANT
Vendor security certifications explicitly exclude AI features	SIGNIFICANT
AI vendors not subject to vendor security review	DISQUALIFYING

How Nylas answers this. Nylas's AI governance is integrated with its broader security program. AI-enabled features are subject to the same access controls, encryption standards, monitoring, and vendor review processes as other platform functionality.

05 Are AI features optional or configurable?

Why it matters. Enterprise organizations have different risk tolerances, internal AI policies, and regulatory obligations. A vendor whose AI features cannot be disabled or scoped forces a binary choice: accept the risk or don't use the product.

ASK

- [ ] Can AI-enabled features be disabled at the account or tenant level?
- [ ] Can AI features be enabled only for specific user roles or groups?
- [ ] Can specific AI-enabled workflows be disabled independently of others?
- [ ] Are there configuration options for data sharing or retention specific to AI features?
- [ ] Can AI features be separated from non-AI platform functionality?

RED FLAGS

FLAG	SEVERITY
AI features cannot be disabled — all-or-nothing deployment	DISQUALIFYING
No tenant-level or role-level configuration for AI	SIGNIFICANT
AI-enabled and non-AI functionality cannot be separated	SIGNIFICANT

How Nylas answers this. Nylas designs AI-enabled features to be configurable. Not every customer uses AI-enabled functionality, and organizations can evaluate which features apply to their use case before enabling them.

06 How does the vendor make AI governance decisions?

Why it matters. "We have an AI ethics policy" is not governance. Governance is a documented process for how decisions about AI systems are made, reviewed, and updated. Buyers should distinguish between vendors with operational governance and those with aspirational statements.

ASK

- [ ] What is the process for reviewing a new AI feature before it is built or enabled?
- [ ] How are AI vendors and subprocessors evaluated before onboarding?
- [ ] How does the vendor test security and privacy risks in AI features before launch?
- [ ] How does the vendor document data flows involving AI systems?
- [ ] How are AI features monitored for misuse or unexpected behavior in production?
- [ ] How does the vendor update its AI governance practices as regulations change?

RED FLAGS

FLAG	SEVERITY
Only an AI policy document exists — no described process	SIGNIFICANT
No pre-launch review for AI features	SIGNIFICANT
No ongoing monitoring of AI feature behavior	SIGNIFICANT
Vendor cannot describe how AI regulatory changes are incorporated	MINOR

> **How Nylas answers this.** Nylas treats AI governance as an ongoing operational process, not a one-time certification. Nylas continues to evaluate and refine its governance practices as technologies, customer expectations, and regulatory obligations evolve.

07 How are AI outputs reviewed or controlled?

Why it matters. An AI summary carries a different risk profile than an AI system that can send an email, update a CRM record, or escalate a support ticket without human approval. Buyers need to know what the system can do, not just what it generates.

ASK

- [ ] Which AI-generated actions require explicit human review before execution?
- [ ] Can the AI system send messages, modify records, or trigger workflows autonomously?

- [ ] What happens when an AI output is incorrect or out-of-scope — is there an escalation path?
- [ ] Are AI-generated actions distinguishable from user-generated actions in logs and audit trails?
- [ ] Can customers configure the level of automation vs. human review for AI-assisted workflows?

RED FLAGS

FLAG	SEVERITY
AI can take consequential actions (send, modify, escalate) with no human review option	DISQUALIFYING
No approval workflow available for high-risk automated actions	SIGNIFICANT
AI-generated and user-generated actions not distinguishable in logs	SIGNIFICANT
No escalation path for incorrect AI outputs	SIGNIFICANT

> **How Nylas answers this.** Nylas designs AI-enabled features to support customer-controlled configuration, including the ability to insert review, approval, or escalation steps into automated workflows. The appropriate level of oversight depends on the specific workflow and the customer's configuration choices.

08 Can customers audit AI activity?

Why it matters. Auditability is a baseline expectation in enterprise security and an increasingly explicit requirement under AI governance frameworks including the EU AI Act. Without an audit trail specific to AI activity, security teams cannot investigate incidents, compliance teams cannot demonstrate controls, and operations teams cannot diagnose unexpected behavior.

ASK

- [ ] Is AI activity logged separately from standard user actions?
- [ ] Can administrators see when and how AI features were invoked?
- [ ] Are AI-generated actions distinguishable from user-generated actions in audit logs?
- [ ] What retention period applies to AI activity logs?
- [ ] Can audit logs be exported for external review or SIEM integration?

RED FLAGS

FLAG	SEVERITY
No AI-specific audit trail — AI actions indistinguishable from user actions	DISQUALIFYING
Audit logs not exportable	SIGNIFICANT
Short or undefined retention period for AI activity logs	SIGNIFICANT
Administrators cannot see when AI features were used	SIGNIFICANT

> **How Nylas answers this.** Customers with specific auditability requirements should contact Nylas directly to understand what logging and audit capabilities apply to the AI features relevant to their deployment.

09 How does the vendor address AI abuse and misuse?

Why it matters. Communications platforms are high-value targets for AI-assisted attacks. Vendors need to have specifically assessed and mitigated the AI-specific risks that email, calendar, and contact systems face — not just general security controls.

ASK

- [ ] Has the vendor specifically tested for prompt injection vulnerabilities in AI features that process email or calendar content?
- [ ] How does the vendor limit phishing and impersonation risks from AI features that compose or route communications?
- [ ] How are excessive permissions prevented — do AI features operate under least-privilege principles?
- [ ] How does the vendor detect and respond to AI-related incidents or misuse patterns?
- [ ] Are AI-specific abuse scenarios documented in the vendor's threat model?

RED FLAGS

FLAG	SEVERITY
Generic security posture claims with no AI-specific abuse testing	SIGNIFICANT
Vendor has not evaluated prompt injection for email/calendar AI features	SIGNIFICANT
AI features not subject to least-privilege access principles	DISQUALIFYING
No AI-specific incident detection or response process	SIGNIFICANT

How Nylas answers this. Nylas evaluates prompt injection, phishing amplification, and excessive permissions as part of its AI security program. Nylas applies least-privilege principles to AI-enabled feature design and maintains incident response processes covering AI-related misuse and unexpected outputs.

10 Who is responsible when something goes wrong?

Why it matters. Shared responsibility models define what each party controls, what each party is liable for, and what each party must do when an incident occurs. Without a documented model, responsibility defaults to ambiguity — which almost always means the customer bears the operational consequence while the vendor disclaims liability.

ASK

- [ ] Is there a documented shared responsibility model for AI-enabled features?
- [ ] What incident response processes does the vendor maintain for AI-related events?
- [ ] If a third-party AI provider causes an incident, what is the vendor's obligation?
- [ ] How and within what timeframe are customers notified of AI-related incidents?
- [ ] What documentation does the vendor provide to support customer investigation of AI incidents?
- [ ] What contractual remedies exist if AI-related events cause data exposure or operational harm?

RED FLAGS

FLAG	SEVERITY
No documented shared responsibility model	SIGNIFICANT
Vendor disclaims all liability for third-party AI provider incidents	DISQUALIFYING
Customer notification process not defined or documented	SIGNIFICANT
No documentation provided to support incident investigation	SIGNIFICANT
Contractual remedies for AI incidents absent or severely limited	DISQUALIFYING

How Nylas answers this. Nylas maintains incident response processes covering AI-related misuse and unexpected outputs. Organizations should also maintain their own incident response processes for AI-enabled workflows. Specific questions about contractual remedies and notification timelines should be directed to the Nylas team.

Overall assessment scoring rubric

After completing all 10 sections, score the evaluation. Count your red flags by severity.

SEVERITY	COUNT	ACTION
DISQUALIFYING	_____	Any disqualifying red flag = require remediation or reject
SIGNIFICANT	_____	3 or more significant = escalate to legal / CISO before proceeding
MINOR	_____	Note for ongoing monitoring

DATA GOVERNANCE · Q1–3

- [ ] Customer data is **not** used for model training
- [ ] Subprocessors are disclosed and contractually governed
- [ ] Specific data types processed are documented
- [ ] Data minimization principles apply to AI features

SECURITY & PRIVACY · Q4, Q8, Q9

- [ ] AI is within scope of the security program and certifications
- [ ] AI activity is logged and auditable
- [ ] AI-specific abuse scenarios have been assessed and mitigated

GOVERNANCE · Q5–6

- [ ] AI features are configurable at account or tenant level
- [ ] Governance is an operational process, not only a written policy

OUTPUT CONTROLS & RESPONSIBILITY · Q7, Q10

- [ ] Human oversight is available for consequential AI actions
- [ ] AI- and user-generated actions are distinguishable in logs
- [ ] Shared responsibility model is documented
- [ ] Incident response and notification processes are defined

OVERALL VERDICT

OUTCOME	CRITERIA
Proceed	No disqualifying flags; fewer than 3 significant flags; all data governance items confirmed
Conditional approval	No disqualifying flags; 3+ significant flags addressed through contractual remediation
Escalate	Any disqualifying flag or 5+ significant flags
Reject or re-evaluate	Multiple disqualifying flags; vendor unable to answer core questions

If a vendor fails this evaluation

A failed evaluation does not always mean the relationship ends. It means specific issues need resolution before the deployment proceeds. Common remediation paths:

- **Contractual.** Add AI-specific data handling terms, training use opt-out, subprocessor notification obligations, and incident response commitments to the MSA or DPA.
- **Technical.** Require the vendor to demonstrate configuration options, audit log access, and scope controls before go-live.
- **Operational.** Define internal governance controls on the customer side — limiting which workflows use AI features, establishing internal review checkpoints, and documenting the shared responsibility model.
- **Escalation.** For deployments involving regulated data or high-risk workflows, involve legal counsel and a data protection officer before proceeding.

Regulatory context

This checklist addresses governance considerations relevant to:

- **EU AI Act (2024).** Risk-based framework assigning obligations to AI system providers and deployers. Organizations using vendor AI features are generally acting as deployers and should evaluate their obligations based on how the AI system is used, what data it processes, and what decisions it influences.
- **GDPR Article 28.** Requires data processing agreements with subprocessors, including AI providers that process personal data.
- **US state privacy laws.** California CPRA, Virginia VCDPA, and similar frameworks may impose obligations on organizations using AI systems that process personal information.

Consult legal counsel regarding your organization's specific regulatory obligations.

nylas

About Nylas

Nylas provides email, calendar, contacts, and scheduling APIs used by software teams integrating communications workflows into their products. Nylas does not use customer data to train AI models, designs AI features to be configurable, and integrates AI governance into its broader security program.

For questions about Nylas's AI governance practices, data handling policies, or subprocessor relationships, contact the Nylas team or visit **nylas.com/security**.

RELATED RESOURCES

How Nylas Approaches AI Governance
Privacy by Design in Email and Calendar APIs
9 Questions to Ask AI Vendors About Data Usage
EU AI Act for Enterprise Communication Platforms