EXTERNAL USE

OWL (On-call Watch List)
SECURITY QUESTIONNAIRE RESPONSES

Date:	December 30, 2025
Prepared by:	Kevin Koplin, CEO/Founder
Application:	OWL (On-call Watch List)
URL:	https://vip.foundopportunity.com

SCOPE NOTICE: This questionnaire applies to OWL (On-call Watch List) only — a rule-based VIP alert service that monitors newly received Inbox messages. Found Opportunity ("FO") is a separate service with spam/junk-folder-only access and has its own security questionnaire.

1. PRODUCT BEHAVIOR & DATA ACCESS

1. What OAuth/Graph scopes do we request?

Google (Gmail):

- https://www.googleapis.com/auth/gmail.readonly — Read-only access to email
- https://www.googleapis.com/auth/userinfo.email — Access to user's email address

Microsoft (Outlook):

- https://graph.microsoft.com/Mail.Read — Read mail
- https://graph.microsoft.com/User.Read — Read user profile
- offline_access — Refresh tokens for continued access

2. Which folders/labels do we actually read from?

- Gmail: ONLY the INBOX label (labelIds: "INBOX").
- Outlook: ONLY the Inbox folder (displayName.lower() == "inbox").
- Other folders: No code paths exist that read Spam/Junk, Sent, Drafts, Trash/Deleted Items, or any other folders.

Additionally: OWL is designed to process **newly received Inbox messages only** (no historical Inbox scanning to generate alerts).

3. For alerts (matches), what fields do we read and store?

Field	Read from Provider?	Stored in Database?
Sender email	Yes	Yes (encrypted)
Sender display name	Yes	Yes
Recipient(s) (To/Cc)	Yes (as needed)	No

Subject	Yes	Yes
Alert Details (up to 10,000 chars)	Yes (for matches)	Yes (7-day retention)
Full body	Yes (as needed)	No (not stored as complete body)
Attachments	No	No
Message ID	Yes	Yes (for duplicate tracking)
Folder/label	Yes (for validation)	No
Timestamp	Yes	Yes
Thread/conversation identifiers	Yes	Yes (for thread rules & dedup)

4. Do we ever store complete email bodies, attachments, or non-Inbox emails?

- **Complete email bodies:** No. OWL stores Alert Details excerpts for matching messages (up to 10,000 characters) with 7-day retention. OWL does not store complete email bodies.
- **Attachments:** No. OWL does not read or store attachments.
- **Non-Inbox emails:** No. OWL is restricted to Inbox-only access.

5. For non-matches (messages that do not match any OWL rule), do we store anything?

- We store the Message ID hash in a processed_emails table to prevent duplicate processing.
- We do NOT store content, metadata, sender information, or any other details about non-matches.
- Aggregate counts may appear in application logs (e.g., "Checked 50 inbox emails, generated 3 alerts").

6. Do we ever modify the user's mailbox?

Action	Do We Do This?
Mark as read/unread	No
Move messages	No
Delete messages	No
Send email	No
Draft messages	No

We have read-only OAuth scopes only. No write operations are possible.

2. INBOX-ONLY ENFORCEMENT & FAILSAFES

7. How do we restrict reads to Inbox folder in code?

- Gmail API calls are hard-coded to only request the INBOX label (labelIds: "INBOX").
- Outlook API calls only access folders named "Inbox".
- OWL processing is designed to evaluate newly received Inbox messages only.

8. Are there any queries that don't specify Inbox folder/label?

No. All Gmail and Outlook email retrieval functions explicitly require the Inbox folder/label.

9. Do we have a failsafe for non-Inbox access attempts?

Yes. A dedicated security module provides circuit breaker protection:

- Custom validation functions block any non-Inbox access
- Optional guardrails block historical Inbox scanning for alert generation (newly received only)

Behavior when triggered:

- Raises a custom security exception (blocks processing)
- Logs CRITICAL alert
- Returns no data to calling function

10. Is there unit/integration testing for Inbox-only access?

Yes:

- Diagnostic scripts verify Inbox-only access behavior
- Automated tests verify Inbox-only access and run on deployment
- OWL additionally tests "newly received only" behavior (no historical Inbox scans to generate alerts)

3. DATA STORAGE, RETENTION & DELETION

11. What database do we use?

DigitalOcean Managed PostgreSQL (version 14.x)

- Automatic encryption at rest (AES-256)
- SSL/TLS connections required
- Daily automated backups

12. What fields are stored for OWL alert records?

Alert records include: sender information (encrypted), subject, Alert Details excerpt (up to 10,000 characters), matching rule identifier, timestamps, thread/conversation identifiers, and user feedback fields. Full schema available upon request under NDA.

13. How long is each field retained?

Data Type	Retention Period	Cleanup Method
Alert records	7 days	Automated cleanup job
Processed email IDs	Indefinite (while active)	For duplicate prevention
User accounts	Until deletion request	Manual or user-initiated

Application logs	Disk rotation	No formal day limit

Cleanup job runs frequently throughout the day.

14. Do we have a documented retention policy?

Yes:

- OWL alert data: 7 days
- Logs: Disk-based rotation (no formal day limit)
- Backups: 7 days (DigitalOcean managed)

Documented in Privacy Policy at https://foundopportunity.com/privacy

15. When a user disconnects their email account:

Action	What Happens
OAuth tokens	Deleted immediately from database
Connection timestamps	Cleared
Existing alerts	Remain until 7-day retention expires
Processed email hashes	Remain (for duplicate prevention if they reconnect)

16. When a user fully deletes their account:

Action	Timeline
User record	Deleted immediately
OAuth tokens	Deleted immediately
All alerts	Deleted immediately
Processed email hashes	Deleted immediately
Data in backups	Purged after 7 days (backup retention period)

Timeline: Deletion completes within seconds of request; backups purge within 7 days.

4. AUTHENTICATION, AUTHORIZATION & ADMIN ACCESS

17. How do end-users log in?

Method	Available?
Magic link (passwordless)	Yes — primary method
Email + password	Yes — bcrypt hashed (12 rounds)
SSO (Google/Microsoft)	No
2FA	Yes — via magic link verification codes

18. How do internal admins access systems?

System	Access Method
Production database	PostgreSQL SSL connection via application or direct query
Admin dashboard	Web-based, protected by authentication
DigitalOcean console	Account login with 2FA enabled
Server SSH	SSH key-based only (password authentication disabled)

19. Is 2FA required for internal/admin access?

System	2FA Status
DigitalOcean account	Yes — enabled
GitHub	Yes — enabled
SSH access	Key-based only (no passwords)
Production server	SSH keys required; password login disabled

20. Who can access what (by role)?

Access Level	Who Has Access
Production database queries	Founder only
Logs with user identifiers	Founder only
Environment variables/secrets	Founder only (permissions 600, root-only)
Admin dashboard	Founder only
Server SSH	Founder only

21. Do we have least privilege implemented?

Currently a single-person operation (Founder only). As the team grows:

- Role-based access control will be implemented
- Support staff will have limited read-only access
- Engineering will have scoped permissions
- Only founder/CTO will have full database access

5. INFRASTRUCTURE, ENCRYPTION & BACKUPS

22. Which cloud providers do we use?

Service	Provider
Application hosting	DigitalOcean (Droplet — Ubuntu 24 LTS)
Database	DigitalOcean Managed PostgreSQL
Object storage	Not currently used
Email delivery	SendGrid
AI/LLM	None (OWL does not use LLM providers)
Payments	Stripe
Domain registrar	Namecheap
DNS provider	Namecheap
Push notifications	Apple APNs, Google FCM

23. In which regions is data stored?

All application data: **DigitalOcean NYC3 (New York, United States)** — No data is stored outside the United States in our primary infrastructure.

Namecheap may store website visitor logs (IP addresses, user agents, etc.) in the regions where their web and DNS infrastructure operates.

Apple/Google push delivery is handled by APNs/FCM infrastructure.

24. How is data encrypted?

In Transit:

- TLS 1.2+ for all web traffic (Let's Encrypt certificate)
- HSTS enabled
- PostgreSQL SSL/TLS connections required
- All API calls (Gmail, Outlook, SendGrid, Stripe, APNs/FCM endpoints) use HTTPS

At Rest:

- DigitalOcean Managed PostgreSQL: AES-256 automatic encryption
- OAuth tokens: application-level encryption (Fernet) before database storage
- Passwords: bcrypt hashing (12 rounds)
- Backups: AES-256 encrypted by DigitalOcean

25. How are OAuth tokens and secrets stored?

Item	Storage Method
OAuth access tokens	Encrypted before database storage
OAuth refresh tokens	Encrypted before database storage
Encryption keys	Secure environment file with restricted permissions (600, root-only)
API keys	Environment variables, never in code or Git

26. How often are backups taken?

Backup Type	Schedule	Location	Retention
Database (PostgreSQL)	Daily automated	DigitalOcean NYC region	7 days
Droplet snapshots	Enabled	DigitalOcean	7 days

27. Have we tested restoring from backups?

- Last verification: November 2025
- Method: Confirmed backup availability via DigitalOcean console
- Documented process: Yes, in internal documentation
- RTO: 4 hours
- RPO: 24 hours

6. LOGGING, MONITORING & ALERTS

28. What events do we log?

Event Type	Logged?
User logins/authentication	Yes
Email connection/disconnection	Yes
Inbox rule-check jobs (start, completion, counts)	Yes
Alert generation results	Yes
Push delivery attempts/failures (where available)	Yes
Errors and exceptions	Yes
API call failures	Yes
Circuit breaker triggers	Yes (CRITICAL level)

29. Where are logs stored?

Log Type	Location
Application logs	Local disk
System logs	journalctl (systemd)
Nginx access/error logs	Local disk
DNS / website logs	Namecheap (as part of hosting/DNS services)
External logging service	Not currently used

30. How long are logs retained?

- Application logs: Rotate based on file size (no formal day limit)
- System logs: journalctl defaults (typically 4GB or ~30 days)
- Target: ~30 days (documented in internal log retention documentation)

31. Do we have alerting configured?

Alert Type	Configured?	Method
High error rates	Yes	Health check emails every 20 minutes
System down	Yes	UptimeRobot (1-minute checks)
No Inbox checks in 15 min	Yes	Health endpoint returns 503, triggers UptimeRobot
Circuit breaker triggers	Yes	CRITICAL log entries
Suspicious access patterns	Yes	Immediate email alert to support@

32. Which uptime/monitoring tools do we use?

UptimeRobot (paid subscription):

- Health monitoring endpoints checked every 1 minute
- Alert recipients: Founder (email)

Internal health monitoring:

- Cron job runs automated health check script every 20 minutes
- Emails Founder with system status and auto-diagnostics if CRITICAL

7. SECURE DEVELOPMENT & CHANGE MANAGEMENT

33. How do we manage source code?

GitHub (private repository)

- Branch: main
- Access: Founder only

34. Are all code changes made through pull requests?

No. Currently a single-developer operation. Code changes are committed directly to main branch.

- No formal PR review process
- No required reviewers before merge
- Will implement PR workflow as team grows

35. Do we use security scanning tools?

Tool Type	Currently Used?
Static analysis (SAST)	Planned when migrating to Github Org
Dynamic analysis (DAST)	Covered via CASA Tier 2 assessment approach
Dependency vulnerability scanning	Yes — Intruder.io (continuous) + Dependabot alerts

36. How are changes deployed to production?

Manual deployment via Git with automated security tests that block deploy on failure.

37. Who can approve/deploy changes?

- Deploy permissions: Founder only (sole person with SSH access)
- GitHub push access: Founder only
- Production server access: Founder only

8. VULNERABILITY MANAGEMENT & THIRD-PARTY ASSESSMENTS

38. Do we use a vulnerability scanning service?

Intruder.io:

- Environments scanned: Production only
- Frequency: Continuous/scheduled scans
- Scope: External attack surface

39. How do we handle vulnerability scan findings?

Severity	Response Time
Critical	Immediate / same day
High	Within 7 days
Medium	Within 30 days

Low/Informational	Evaluated case-by-case

40. Have we undergone third-party security assessments?

CASA Tier 2 Assessment:

- Assessor: TAC Security
- Completed: October 2025
- Status: Passed after remediation
- Letter of Validation: Submitted to Google for OAuth verification

Intruder.io provides continuous automated vulnerability scanning.

41. Do we track and patch OS/library vulnerabilities regularly?

- OS updates: Applied as needed (no formal schedule)
- Python dependencies: Updated during deployments
- Monthly patching scheduled for 1st Saturday per internal patching policy

9. INCIDENT RESPONSE & BREACH HANDLING

42. Do we have a documented incident response plan?

Yes.

- Internal documentation
- Incident team: Founder (sole responder)
- Classification: P1 (Critical) through P4 (Low)
- Escalation procedures documented
- Communication templates included

43. How do we detect potential incidents?

Detection Method	What It Catches
UptimeRobot monitoring	Site down, health endpoint failures
Health check emails (20 min)	Inbox-check failures, database issues, stuck queues
Application error logs	Exceptions, API failures, circuit breaker triggers
DMARC reports	Email spoofing attempts (for our domains)
Intruder.io	New vulnerabilities

44. Process for confirmed incidents affecting customer data:

1. **Contain:** Immediately isolate affected systems
2. **Assess:** Determine scope and affected users
3. **Remediate:** Fix vulnerability/breach vector
4. **Notify:** Inform affected users within 72 hours (GDPR requirement where applicable)
5. **Document:** Complete incident report
6. **Review:** Post-mortem and preventive measures

45. Do we have notification templates?

Yes. Templates included in incident response documentation for:

- Initial customer notification
- Follow-up with remediation details
- Regulatory notification (if required)

10. PRIVACY, DATA SUBJECT RIGHTS & LEGAL

46. Where are our policies stored?

Document	Location
Privacy Policy	https://foundopportunity.com/privacy
Terms of Service	https://foundopportunity.com/terms
Data Processing Addendum	https://foundopportunity.com/dpa

47. How can users exercise their data rights?

Right	How to Exercise
Access data	Account dashboard or email privacy@foundopportunity.com
Delete data	Account settings "Delete Account" or email request
Export data	Email request to privacy@foundopportunity.com
Rectify data	Email request to privacy@foundopportunity.com

48. Do we have internal guidelines for data requests?

Yes:

- Data subject requests: Respond within 30 days
- Law enforcement requests: Require valid legal process (subpoena, warrant)
- Document all requests and responses

49. What regulatory frameworks do we align with?

GDPR:

- 7-day data retention for OWL alerts
- Encryption at rest and in transit
- Right to erasure implemented
- Right to data portability (export)
- Data breach notification within 72 hours (where applicable)

CCPA/CPRA:

- Aligned with GDPR practices
- Do not sell personal information
- Deletion rights honored

50. Do we maintain records of data flows and processors?

Yes:

- Internal data classification documentation — categories of personal data and protection levels
- Internal data flow documentation — system architecture and data flows
- Third-party processor list maintained and documented in Privacy Policy

11. THIRD-PARTY SERVICES & PROCESSORS

51. List of all third-party providers:

Provider	Purpose	Data Received
DigitalOcean	Hosting & Database	OWL application data (alerts, rules, metadata)
SendGrid	Email delivery	User emails; transactional email content
Stripe	Payment processing	Payment info (via Stripe.js; no card data on our servers)
Google/Microsoft	OAuth + email provider APIs	OAuth tokens/authorization; read-only mail access
UptimeRobot	Monitoring	URL endpoints only (no user email data)
Intruder.io	Security scanning	External attack surface only
Namecheap	Domain/DNS + hosting	Account/billing/domain config; visitor/DNS metadata
Apple (APNs)	Push notifications (iOS)	Device push tokens; notification payload metadata
Google (FCM)	Push notifications (Android)	Device push tokens; notification payload metadata

52. Provider certifications and data handling:

Provider	Certifications	Sub-processor?	Data Retention
DigitalOcean	SOC 2 Type II, ISO 27001	Yes	Per our configuration
SendGrid	SOC 2, ISO 27001	Yes	Transient only
Stripe	PCI-DSS Level 1, SOC 2	Yes	Per Stripe policies
Google	SOC 2, ISO 27001	No (user provider)	N/A
Microsoft	SOC 2, ISO 27001	No (user provider)	N/A
Namecheap	ISO 27001	Yes	Per Namecheap policies
Apple (APNs)	(Apple security practices)	No (delivery infra)	Per Apple policies
Google (FCM)	(Google security practices)	No (delivery infra)	Per Google policies

53. Do we have contracts with all processors?

Provider	Contract Type	Data Protection Clauses
DigitalOcean	Terms of Service + DPA	Yes
SendGrid	Terms of Service + DPA	Yes
Stripe	Stripe Services Agreement + DPA	Yes
Namecheap	Terms of Service + DPA	Yes

12. INSURANCE & BUSINESS CONTINUITY

54. Do we have cyber liability insurance?

Status: $1M policy with Hiscox Insurance Company

Coverage:

- Data breach response costs
- Cyber extortion and crime
- Business interruption
- Third-party liability
- Regulatory defense

55. Do we have a business continuity/disaster recovery plan?

Yes.

Metric	Target
Recovery Time Objective (RTO)	4 hours
Recovery Point Objective (RPO)	24 hours

Recovery capabilities:

- Database: Restore from DigitalOcean daily backups
- Application: Redeploy from GitHub within 2 hours
- Alternate infrastructure: AWS account prepared (not active)

Documented in internal documentation.

Document Version:	1.0
Last Updated:	December 30, 2025
Next Review:	March 30, 2026