INCIDENT RESPONSE PLAN

Information Security Program

Version 1.0

2026-05-22

Prepared for Walker-Miller Energy Services

Project Baseline, Inc. | Colorado | EIN 27-0639457 | todd@project-baseline.com | project-baseline.com

Entity: Project Baseline, Inc.

Effective date: 2026-05-22

Version: 1.0

Owner: Todd Walton, Incident Commander

Review cadence: Annual, plus after any actual incident

Framework alignment: NIST CSF 2.0 (Respond function), NIST SP 800-61 Rev. 2

1. PURPOSE

This Incident Response Plan (IRP) defines how Project Baseline, Inc. (PB) detects, responds to, contains, eradicates, recovers from, and learns from security incidents. It exists to (a) limit damage when something goes wrong, (b) meet client and regulator notification SLAs, and (c) demonstrate to clients and auditors that PB has a documented, testable response capability.

2. SCOPE

This plan covers any security event affecting:

- PB-owned systems (VPS, SaaS accounts, endpoints)
- Client data PB processes
- Prospect data submitted to PB-controlled forms
- PB-owned intellectual property and source code
- Sub-processor compromises that may affect PB or PB's clients (sub-processor incident notification feeds into PB's response process)

3. INCIDENT SEVERITY LEVELS

LEVEL	DEFINITION	EXAMPLES	RESPONSE TIME
Critical (Sev-1)	Active compromise of client data, payment systems, or production services with confirmed exposure	Confirmed unauthorized access to client documents in nonprofit-portal; Stripe webhook signature bypass; ransomware on VPS	Immediate; Todd drops other work
High (Sev-2)	Active or imminent threat without confirmed exposure; or sub-processor breach notification affecting PB	Brute-force attack succeeding on nonprofit-portal admin; GitHub repo accidentally made public with secrets; Anthropic API key compromise notification	Within 4 hours
Medium (Sev-3)	Suspicious activity without immediate threat; policy violations	Sustained 429s from a single IP indicating abuse; failed MFA pattern; sub-processor security advisory not yet patched	Within 24 hours
Low (Sev-4)	Operational anomaly with security implications	Single failed login from unusual location; deprecated TLS version detected on internal endpoint; cosmetic security finding	Within 72 hours

4. RESPONSE PHASES (NIST SP 800-61)

4.1 PREPARATION

Already complete per the Information Security Program:

- This IRP document
- Contact list (Section 8 below)
- Tooling inventory (system logs, application logs, monitoring tools)
- Backup and recovery procedures (per Backup Policy)

- Communication templates (Section 9 below)

4.2 DETECTION AND ANALYSIS

Detection sources:

- VPS monitoring: nginx access/error logs, application logs, fail2ban, system journal
- Application alerts: cos-webhook alerts to Todd's Telegram, structured logs in PB Report Engine and nonprofit-portal
- SaaS notifications: GitHub security alerts, Cloudflare security events, Stripe Radar, Google Workspace alerts, DigitalOcean monitoring
- Sub-processor notifications: vendor breach notifications to todd@project-baseline.com
- External report: client, prospect, security researcher, or law enforcement reporting an issue to todd@project-baseline.com or via the website contact form

Analysis steps (within 2 hours of detection for Sev-1/Sev-2):

1. Verify the report. Is this a real incident or a false positive?
2. Classify severity (use table in Section 3)
3. Determine initial scope: which systems, which data, which clients, which timeframe?
4. Document everything in the incident log (Section 6 below)
5. Decide if external help is needed (counsel, forensic firm, breach coach)

4.3 CONTAINMENT

Short-term containment (within 1 hour for Sev-1):

- Isolate affected systems (revoke compromised credentials, disable affected accounts, block attacker IPs at Cloudflare, take down compromised services)
- Preserve evidence (capture logs, file hashes, memory dumps if technically feasible before remediation)
- Stop the bleed: rotate credentials, push emergency patches, take affected endpoints offline

Long-term containment (within 24 hours for Sev-1):

- Apply temporary fixes that allow business to resume while permanent remediation is developed
- For example: WAF rule blocking attack pattern while patched code is deployed

- For example: rate-limit further restriction on a compromised endpoint until investigation completes

4.4 ERADICATION

- Identify root cause (vulnerability, misconfiguration, credential theft, social engineering, sub-processor breach)
- Remove malicious artifacts (malware, backdoors, attacker accounts)
- Patch the underlying vulnerability
- Rotate ALL credentials with any possibility of compromise (not just the confirmed-compromised ones)
- Update detection rules so the same attack would be caught next time

4.5 RECOVERY

- Restore from clean backups if necessary
- Verify system integrity before returning to production
- Monitor closely for re-occurrence (elevated monitoring for 30 days post-incident)
- Document recovery completion in incident log

4.6 LESSONS LEARNED

Within 14 days of incident closure:

- Write a post-incident review (timeline, root cause, what worked, what failed, what changes)
- Update affected policies, procedures, and controls
- Update Risk Register
- Schedule a follow-up control verification at 30, 60, and 90 days

5. NOTIFICATION SLAS

AUDIENCE	TRIGGER	NOTIFICATION WINDOW	NOTIFICATION CHANNEL
Affected clients	Confirmed exposure of client data	**Within 72 hours of confirmed exposure**	Direct email from todd@project-baseline.com, follow-up call within 24 hours of email
Walker-Miller (if WM data affected)	Confirmed exposure of WM trade ally data or WM client data	Within 72 hours of confirmed exposure	Email to Nana Ellis (WM IT) and Crystal Davis (WM contract owner), with phone follow-up
Regulators (if applicable)	Notifiable breach under CCPA, GDPR, state breach notification laws, or sectoral regulation (HIPAA, GLBA, etc.)	Per applicable law (CCPA: "without unreasonable delay"; GDPR: 72 hours to supervisory authority)	Per regulator's prescribed channel; counsel consulted before submission
Affected individuals (if applicable)	Per state breach notification law thresholds	Per applicable law	Direct notification per applicable law; counsel consulted on form
PB internal (Todd)	Any Sev-1 or Sev-2	Immediate (Todd is both detector and responder in solo-firm state)	Telegram alert from cos-webhook, plus Todd self-aware
Law enforcement	If criminal activity suspected	Promptly, after counsel consultation	Local FBI field office (Chicago: 312-421-6700) or local police
Cyber insurance	Any incident potentially triggering coverage	Per policy terms (usually immediately or within 24 hours)	Per insurer's claims process; check active policy for current insurer
Sub-processors	If incident involves data shared with them	Promptly	Per sub-processor's incident response contact

72-hour Walker-Miller notification is a hard commitment built into the Vendor / Sub-processor Management Policy and the executed subcontractor agreement. Missing it has business and reputational consequences and is explicitly tracked.

6. INCIDENT DOCUMENTATION

Every incident, regardless of severity, is logged in the incident register at:

`output/project-baseline/_agency/security/incident-register.md` (to be created within 30 days of policy effective date)

Each incident entry includes:

- Incident ID (format: `INC-YYYYMMDD-NNN`)
- Detection timestamp
- Severity level
- Affected systems and data
- Affected clients
- Description and timeline
- Containment, eradication, and recovery actions taken
- Notifications sent (when, to whom, by what channel)
- Root cause analysis
- Lessons learned and follow-up actions
- Closure date

7. TABLETOP EXERCISES

Annual tabletop exercise on or before May 22 each year. Even as a solo firm, Todd runs through a scenario (for example: "the nonprofit-portal SQLite database has been exfiltrated; what do you do in the next 4 hours") and documents the response. Tabletop log entry recorded in the incident register and Security Awareness Training Record.

8. CONTACT LIST

Maintained at `output/project-baseline/_agency/security/contact-list.md` (to be created within 30 days). Includes:

- Todd Walton (Incident Commander): todd@project-baseline.com, mobile (in brain vault)
- Cyber insurance carrier: TBD (pending COI quote selection)
- Legal counsel: TBD (placeholder)
- Walker-Miller IT (Nana Ellis): nana@wmenergy.com (via Marco Diaz introduction)

- Walker-Miller contract owner (Crystal Davis): crystal@wmenergy.com
- Cloudflare emergency contact: per Cloudflare dashboard
- DigitalOcean support: per DigitalOcean dashboard
- Stripe incident response: per Stripe Dashboard
- Google Workspace admin: Todd (self)
- Anthropic security: security@anthropic.com (per public security policy)
- FBI Chicago field office: 312-421-6700
- Illinois Attorney General Consumer Protection: 312-814-3000

9. COMMUNICATION TEMPLATES

9.1 CLIENT NOTIFICATION (SEV-1)

```
Subject: Important security notice regarding your Project Baseline engagement

Dear [Client Name],

We are writing to inform you of a security incident that occurred at Project Baseline on [date] and that

What happened: [brief, factual description]

What information was involved: [specific data types, do NOT speculate; if unknown, say so]

What we are doing: [containment, eradication, recovery actions taken; cooperation with law enforcement i

What you can do: [specific protective steps, for example monitor accounts, change passwords on shared sy

We take this seriously and we are sorry for any concern this causes. We are available to answer any ques

Todd Walton
Principal, Project Baseline, Inc.
```

9.2 WALKER-MILLER NOTIFICATION (72-HOUR SLA)

```
Subject: 72-hour security notice: incident affecting WM data

Nana, Crystal,

Per our subcontractor agreement and incident response SLA, I am notifying you within 72 hours of confirm

Incident summary: [factual]

WM data affected: [specific scope]

Actions taken so far: [containment, investigation]

Next steps: [eradication, recovery, follow-up reporting cadence]

I am available for a call today to walk through this in detail. Please advise on WM's preferred next ste

Todd Walton
Principal, Project Baseline, Inc.
```

9.3 INTERNAL INCIDENT LOG ENTRY

See Section 6 above for required fields.

10. PLAN MAINTENANCE

- Annual review on or before May 22 each year
- Post-incident review within 14 days of any actual incident
- Update when sub-processor list changes materially
- Update when notification law changes
- Update when business model changes (for example, launch of $99 audit report tier increases prospect data volume)

Signature: Todd Walton, Principal and Incident Commander, Project Baseline, Inc.

Date: 2026-05-22

INCIDENT LOG

INC-2026-05-24-001: STRIPE LIVE SECRET KEY EXPOSURE

Severity: Sev-3 (credential exposure, no malicious use detected)

Discovery date: 2026-05-24

Discovery method: Foreground automation agent self-detected immediately after action

Status: Mitigation pending (key rotation authorized for later by Todd)

What happened:

During a VPS .env audit script execution, a redaction filter used exact-string matching for credential field names (`ADMIN_PASSWORD`, `JWT_SECRET`, `STRIPE_SECRET`, `SESSION_SECRET`). The actual environment variable name is `STRIPE_SECRET_KEY` (with `_KEY` suffix). The exact-match filter did not catch it, and the full live Stripe key value (`sk_live_...`) was printed to the Claude Code conversation transcript.

Scope of exposure:

- Channel: Claude Code transcript history for Todd Walton's local session
- Audience: Todd Walton (sole conversation participant)
- Persistence: Conversation transcript stored locally on Todd's Windows 11 device
- External access: None known

Mitigation actions:

1. Immediate: Stop further automated reads of .env values; switch to Python `os.environ.get()` pattern
2. Authorized rotation: Todd will rotate the Stripe key via Stripe Dashboard (Developers, API Keys, Roll secret key)
3. VPS update: After rotation, agent updates `~/claudework/public/project-baseline/nonprofit-portal/.env` STRIPE_SECRET_KEY via SSH + PM2 restart
4. Audit: Stripe dashboard "Recent activity" + "Live mode logs" review for any anomalous API calls between exposure and rotation

Preventive actions implemented:

1. New PostToolUse hook at `.claude/hooks/redact-secrets.js` (2026-05-24) scans all Bash tool output for known credential patterns (Stripe, AWS, GitHub, OpenAI, Anthropic, Google, Slack, DigitalOcean, JWT, and generic env-var-secret patterns) and redacts before display. Audit log at `C:/ClaudeWork/private/credential-redaction-audit.log`
2. New memory rule: NEVER use exact-string comparison for credential field redaction. Use regex/pattern matching that captures any field containing PASSWORD/SECRET/KEY/TOKEN in its name

3. New memory rule: NEVER use `cat`, `grep`, `awk`, `sed` directly on .env files. Use a Python script that reads via `os.environ` or filters by allowlist of safe-to-display field names

WM client notification:

Not required. WM customer data was not affected (Stripe key controls Project Baseline's own Stripe account, not WM data). 72-hour SLA only applies when client data may be impacted.

Lessons:

- Exact-string matching for security filters fails open on field name variations. Pattern matching is mandatory.
- Real incidents will recur. This Incident Response Plan was used in practice within 30 hours of being authored. Tested.
- The detect-document-mitigate cycle worked as designed.

Closed: Pending key rotation by Todd, VPS update, and Stripe activity audit.

END OF DOCUMENT

Project Baseline, Inc. | Colorado | EIN 27-0639457 | todd@project-baseline.com | project-baseline.com