Agent Midas Data Security & Privacy Whitepaper v1.2

Contents

AGENT MIDAS

Data Security & Privacy Whitepaper

How Agent Midas Protects Your Business Data, Financial Information, and Personal Privacy

Buji Development Corporation

1712 Pioneer Ave. Ste. 500, Cheyenne, WY 82001

www.agentmidas.xyz

April 2026 | Version 1.2

Document Classification: Public

Prepared for Agent Midas subscribers, prospective subscribers, and security reviewers

Table of Contents

1. Our Commitment to Your Security

At Buji Development Corporation, we built Agent Midas with a foundational belief: your data is yours. Period. It is not our product. It is not our inventory. It is not a revenue stream. We make money one way: your subscription. We have zero incentive to sell, share, or exploit your data because you are our customer, not our commodity.

This whitepaper provides a comprehensive, transparent overview of every security measure we have implemented to protect your business data, financial information, and personal privacy. We welcome scrutiny. We welcome tough questions. Security is not a feature we market; it is an architectural decision embedded in every layer of the platform.

Our security posture is built on five foundational principles:

- **Data Sovereignty:** Each subscriber's data is architecturally isolated at the database level through Row-Level Security policies, with dedicated container isolation for paid subscribers.

- **Minimum Necessary Access:** No internal system, administrator, or AI agent accesses data beyond what is required for its function. Financial data in the subscriber's silo is inaccessible to administrative systems.
- **Verified Identity:** Every individual who earns commissions through the platform undergoes identity verification and must connect a verified payout method before participating in the compensation program.
- **OAuth 2.0 for All Integrations:** No integration uses static credential sharing. All third-party connections use token-based authentication with scoped permissions, encrypted storage, and subscriber-controlled revocation.
- **Continuous Monitoring:** Append-only audit logging, session tracking, login monitoring, and automated health checks provide real-time and historical visibility into all system activity.

2. Architectural Overview: How Your Data Is Isolated

Agent Midas is a multi-tenant platform serving subscribers across eight subscription tiers. Despite serving multiple subscribers from shared infrastructure, your data is completely isolated from every other subscriber. This isolation is not implemented as an application-level filter that could theoretically be bypassed. It is enforced at the database engine level by PostgreSQL's Row-Level Security (RLS) system.

This means that even if an attacker somehow gained access to the database directly, RLS policies would prevent them from seeing any records belonging to a different subscriber. The isolation is enforced by the database engine itself, not by application code. This is the same approach used by financial institutions and healthcare systems that must guarantee data separation between clients.

For paid subscribers, we provide an additional layer of isolation through dedicated Docker containers. Each paid subscriber's Agent Midas instance runs in its own isolated container with separate compute, storage, and process space. No container can communicate with, access, or even detect the existence of any other subscriber's container.

3. Data Classification: Four Levels of Protection

All data within the Agent Midas platform is classified into four levels, each with specific handling requirements, access controls, and encryption standards.

Level	Classification	Examples	Access	Encryption
1	Public	Marketing materials, pricing, white papers	No restrictions	TLS in transit
2	Internal	Aggregated analytics, system config	Authenticated administrators only	TLS in transit
3	Confidential (Dossier)	Onboarding data, subscription info, brand assets	Administrative systems for billing/support	AES-256-GCM at rest + TLS

Level	Classification	Examples	Access	Encryption
4	Restricted (Subscriber Silo)	CRM contacts, emails, financials, meeting transcripts, banking credentials	ONLY subscriber's own instance. Not accessible to administrators.	AES-256-GCM at rest + TLS. RLS enforced.

Critical Design Decision: The separation between Level 3 (Dossier) and Level 4 (Silo) data is an architectural constraint, not an application-level filter. Row-Level Security policies enforce that queries against silo tables physically cannot return rows belonging to a different subscriber, regardless of the role or permissions of the querying user.

4. Encryption Standards

All data in transit is encrypted using TLS 1.3, which exceeds the industry-standard TLS 1.2 requirement. HTTP Strict Transport Security (HSTS) is enabled on all domains, and no plaintext HTTP connections are accepted under any circumstances.

All sensitive data at rest is encrypted using AES-256-GCM (Galois/Counter Mode), the same encryption standard used by the U.S. government for classified information. This applies to all stored credentials, API keys, OAuth tokens, Plaid access tokens, tax identification numbers, and other sensitive subscriber data.

Encryption keys are per-subscriber and derived from a master key using a key derivation function. Decrypted values exist only in memory at the moment of use for authorized API calls. They are never logged, cached, or written to any secondary storage.

Data Type	Encryption Method	Key Management
API keys and OAuth tokens	AES-256-GCM at rest	Per-subscriber derived keys
Plaid access tokens	AES-256-GCM at rest	Isolated in subscriber silo
Tax identification numbers	SHA-256 one-way hash	Cannot be reversed or displayed
Session tokens	httpOnly Secure cookies	SameSite=Lax, never in localStorage
All traffic	TLS 1.3 in transit	Cloudflare-managed certificates
Passwords	bcrypt with salt	Never stored in plaintext

5. Authentication and Access Control

Every interaction with the Agent Midas platform requires authentication. There are no anonymous data access paths. The platform uses cookie-based server-side rendering (SSR) authentication through Supabase Auth, which means your session is validated server-side on every single request. There are no exposed tokens in URLs and no client-side authentication shortcuts.

Authentication is implemented with JWT (JSON Web Tokens) containing custom claims including the subscriber's role, subscription tier, and subscriber ID. These claims are verified on every API request, and the subscriber ID used for all database queries is derived exclusively from the authenticated session, never from request parameters.

This architectural decision was validated during an internal security audit in February 2026 when five API routes were identified that previously accepted subscriber ID from the request body. All five were immediately corrected to derive the subscriber ID exclusively from the authenticated session. The vulnerability was discovered, documented, and remediated within 24 hours.

Production Access Controls

- **SSH Key Authentication:** All server access requires SSH key authentication. Password-based SSH is disabled. Keys are unique per authorized individual and rotated semi-annually.
- **Named Accounts:** No shared or generic accounts exist for production access. Every action is attributable to a named individual.
- **Role-Based Access:** Production access is tiered by role: CEO/CTO have full infrastructure access with MFA; senior developers have deployment and database read access with MFA; support staff have read-only dashboard access with no database or server access.
- **No Administrative Override for Silo Data:** Even administrators with full production access cannot read subscriber silo data (Level 4). The RLS policies have no administrative bypass.

6. Multi-Factor Authentication (MFA)

Multi-factor authentication is a critical component of our security architecture. All administrative accounts with access to production infrastructure, database management, deployment systems, or any system that stores or processes consumer financial data are required to use multi-factor authentication. No exceptions.

MFA is required on:

- Supabase dashboard access (database administration)
- DigitalOcean infrastructure access (server management)
- Cloudflare dashboard (WAF, DNS, security configuration)
- Stripe dashboard (payment processing, affiliate payouts)
- GitHub repository access (source code, deployment pipelines)
- Domain registrar and DNS management

MFA Before Banking Access: Before a subscriber is presented with the Plaid Link interface to connect their bank account, they must complete a step-up MFA verification. Even if the subscriber is already logged in, they must re-verify their identity via a second factor before the Plaid Link interface is displayed. If MFA verification fails, no Plaid Link token is generated and no banking interaction occurs.

7. Third-Party Integration Security (OAuth 2.0)

All third-party service integrations use OAuth 2.0 as the authentication mechanism. No integration uses static credential sharing, API key embedding in client-side code, or unencrypted token storage.

Integration	Auth Method	Scopes	Token Storage	Revocation
Plaid (Banking)	OAuth 2.0 + Link Token	Balances, transactions, identity	AES-256-GCM in subscriber silo	Subscriber-controlled via Settings
Stripe (Payments)	OAuth 2.0 + API Keys	Subscriptions, payouts	Encrypted at rest	Key rotation on schedule
Google (Auth)	OAuth 2.0	Auth, Calendar, Gmail, Drive	Encrypted in subscriber silo	Subscriber-revocable
Slack, Salesforce, etc.	OAuth 2.0 per provider	Varies, subscriber-initiated	AES-256-GCM in container	Subscriber-revocable

8. Plaid Banking Integration: Bank-Grade Security

The Plaid integration is available for Tier 7 (Finance and Legal, $1,500/month) and Tier 8 (The Oracle, $2,500/month) subscribers who require banking connectivity for automated bookkeeping, cash flow monitoring, and financial reconciliation. Plaid serves as the secure intermediary, and Agent Midas never sees, handles, or stores the subscriber's bank credentials.

How the Plaid Integration Works

- **Step 1:** Subscriber navigates to Finance section and clicks 'Connect Bank Account.'
- **Step 2:** Step-up MFA verification is required before proceeding (see Section 6).
- **Step 3:** Agent Midas creates a Plaid Link token via a server-side API call (never client-side).
- **Step 4:** Subscriber completes the Plaid Link flow in a Plaid-hosted modal. Agent Midas never sees the subscriber's bank credentials.
- **Step 5:** Plaid returns a public token. Agent Midas exchanges it server-side for an access token.
- **Step 6:** The access token is encrypted with AES-256-GCM and stored in the subscriber's isolated data silo.
- **Step 7:** The access token is used exclusively for authorized operations: balance checks, transaction retrieval, and identity verification.
- **Step 8:** The subscriber can revoke Plaid access at any time through Settings. Revocation immediately deletes the token and calls Plaid's removal endpoint.
- **Step 9:** No Plaid access token is ever accessible to administrators, other subscribers, or any cross-tenant system.

Buji Development Corporation has been approved by Plaid following a comprehensive security review. As part of this approval, we are required to attest to eight specific security requirements by August 2026. Our current compliance status is detailed in Section 18 of this document.

9. Row-Level Security: Database-Enforced Isolation

Row-Level Security (RLS) is the cornerstone of our data isolation architecture. RLS is a feature of PostgreSQL that enforces access policies at the database engine level. When a query is executed, the database engine automatically filters results to only include rows that the authenticated user is authorized to access. This happens before the application code ever sees the data.

In practical terms, this means that even if a bug existed in our application code that attempted to query another subscriber's data, the database would return zero results. The isolation cannot be circumvented by application logic because it is enforced by the database itself. Every subscriber-facing table in our database (38+ tables) has RLS enabled with policies that scope all queries to the authenticated subscriber's ID.

10. Infrastructure and Network Security

Agent Midas is protected by multiple layers of network, application, and infrastructure security. All public-facing traffic passes through Cloudflare before reaching our origin servers, providing defense-in-depth from the network edge through to the database layer.

10.1 Cloudflare Pro Plan Configuration

All traffic to agentmidas.xyz is routed through Cloudflare's global network, which operates in over 300 cities across 100+ countries. We operate on the Cloudflare Pro plan, which provides enterprise-grade protection including WAF managed rulesets, bot management, image optimization, and DDoS mitigation.

Feature	Status	Purpose
Page Shield	Active	Monitors third-party scripts (Stripe, LiveKit, Resend) for malicious injection or tampering
Super Bot Fight Mode	Active	Blocks AI scrapers, credential stuffing bots, and automated attacks
Cloudflare Managed Ruleset (WAF)	Active	Continuously updated ruleset maintained by Cloudflare's threat intelligence team
OWASP Core Ruleset (WAF)	Active	Blocks SQL injection, XSS, and OWASP Top 10 attacks across all forms and API endpoints

Feature	Status	Purpose
Leaked Credentials Mitigation	Active	Detects login attempts using known compromised passwords; rate-limits IP for 10 seconds
Block AI Training Bots	Active	Prevents AI crawlers from scraping subscriber content, blog posts, and landing pages
Browser Integrity Check	Active	Evaluates HTTP headers for known threat signatures
DDoS Protection	Active	Always-on volumetric and application-layer DDoS mitigation at the edge
Brotli Compression	Active	Compresses text-based responses for faster delivery
Image Optimization (Mirage + Polish)	Active	Mobile optimization and lossless compression at the edge
Always Online	Active	Serves cached versions if origin becomes unreachable
Rocket Loader	Disabled	Deliberately disabled — breaks React/Next.js SSR hydration
Under Attack Mode	Disabled	Reserved for active DDoS incidents only
Development Mode	Disabled	Bypasses caching — used only during active development

10.2 Web Application Firewall (WAF)

The Agent Midas WAF operates two managed rulesets simultaneously:

- **Cloudflare Managed Ruleset:** Maintained and updated by Cloudflare's threat intelligence team. Protects against zero-day exploits and emerging threats without requiring manual intervention.
- **OWASP Core Ruleset:** Implements the OWASP Top 10 attack prevention rules, providing defense against SQL injection, XSS, CSRF, SSRF, and other common web application attacks. Every form submission, API call, and data input on the platform passes through these rules.

The WAF processes all traffic before it reaches our origin servers. Malicious requests are blocked at Cloudflare's edge, ensuring they never consume application server resources or reach the database layer.

10.3 Bot Management and AI Crawler Protection

- **Super Bot Fight Mode:** Identifies and blocks automated bots attempting credential stuffing, content scraping, or denial-of-service attacks using machine learning models trained on traffic patterns from millions of websites.

- **AI Training Bot Blocking:** A dedicated managed rule blocks AI training crawlers from all major AI companies. This prevents subscriber content, blog posts, and landing pages from being ingested into third-party AI training datasets. Your intellectual property remains yours.

10.4 SSL/TLS Configuration

Setting	Configuration
SSL Mode	Full (Strict) — Cloudflare validates the origin certificate before connecting
Origin Certificate	Cloudflare Origin Certificate (RSA 2048, 15-year validity). Covers agentmidas.xyz and *.agentmidas.xyz.
TLS Version	TLS 1.3 enforced. TLS 1.0/1.1 disabled. TLS 1.2 accepted as minimum fallback.
HSTS	HTTP Strict Transport Security enabled on all domains
HTTP Redirect	All plaintext HTTP requests automatically redirected to HTTPS
Certificate Coverage	agentmidas.xyz (main application) and *.agentmidas.xyz (all subdomains)

10.5 Performance and Caching

Cloudflare's CDN caches static assets at edge locations worldwide, reducing load on our origin servers and delivering faster page loads to subscribers regardless of their geographic location.

10.6 Load Balancing and Automatic Failover

Agent Midas uses Cloudflare Load Balancing to provide automatic failover between our production and standby servers. If the primary server becomes unreachable, subscriber traffic is automatically redirected to a healthy server within approximately two minutes — with no manual intervention required.

Component	Configuration
Primary Origin	Atlanta, GA (DigitalOcean) — receives all traffic under normal operation
Failover Origin	San Francisco, CA (DigitalOcean) — standby server
Health Check Protocol	HTTP health check against the root path, every 60 seconds
Failure Threshold	2 consecutive failures before failover (approximately 2 minutes)

10.7 Application and Database Layer

Layer	Technology	Protection Provided
CDN / WAF	Cloudflare Pro	WAF, DDoS mitigation, SSL/TLS termination, bot filtering, Page Shield
Transport	TLS 1.3 + HSTS	All traffic encrypted. Full (Strict) origin validation.
DNS	Cloudflare DNS + DNSSEC	DNS-level security. No spoofing or hijacking.
Application	Next.js (SSR) + TypeScript	Server-side rendering for sensitive logic. Type safety.
Database	Supabase (PostgreSQL 15)	Row-Level Security on 38+ tables. JWT-authenticated connections.
Hosting	DigitalOcean (managed)	SSH key authentication only. Firewall rules.
Container Isolation	Docker (paid subscribers)	Dedicated container per subscriber.

10.8 Subscriber Container Isolation

For paid subscribers, Agent Midas provides dedicated Docker containers. Each container runs its own isolated instance of the platform with separate compute, storage, and process space. No container can communicate with, access, or even detect the existence of any other subscriber's container. This provides defense-in-depth beyond the database-level RLS isolation that applies to all subscribers.

11. Fraud Detection: SENTINEL-PAY and Hermes

Agent Midas operates two independent fraud detection systems that work in tandem to protect the platform and its subscribers.

SENTINEL-PAY is an 8-signal fraud detection engine that runs every 6 hours. It monitors for suspicious patterns across affiliate signups, commission activity, chargeback rates, referral patterns, and payout behavior. When SENTINEL-PAY identifies a high-risk pattern, it can automatically freeze an account for manual review.

Hermes is our AI-powered fraud scoring engine (powered by DeepSeek R1) that runs daily at 2:00 AM EST. Hermes performs deep analysis of affiliate behavior patterns, assigns risk scores on a 0-100 scale, and flags accounts that exhibit characteristics associated with fraudulent activity.

Together, these systems provide layered protection: SENTINEL-PAY catches known fraud patterns in real-time, while Hermes identifies novel threats through AI analysis. Three or more chargebacks trigger an automatic account freeze per Visa/Mastercard compliance requirements.

12. AI Data Handling: What Our AI Does and Does Not Do

Agent Midas is powered by the Supra Intelligence Engine, a multi-LLM system coordinating specialized AI agents. Given the sensitivity of the data our AI agents process, we maintain strict

boundaries around how AI interacts with subscriber data.

What our AI does NOT do:

- **It does not train on your data.** Your business information is never used to improve models for other subscribers or for any external purpose.
- **It does not share your data across subscribers.** Each subscriber's AI agents operate exclusively within their data silo. No cross-subscriber data leakage is possible.
- **It does not store raw audio from meetings.** Our meeting intelligence system (Nemo) transcribes audio and immediately deletes the audio file. Only the text transcript is retained.
- **It does not make financial or legal decisions without your approval.** All financially consequential actions are presented to the subscriber for explicit confirmation.
- **It does not send data to unauthorized third parties.** External AI model providers receive only the minimum context necessary for the specific task, and no subscriber-identifying information is included in API calls.

12.1 Harpocrates: Your Private AI Engine

Harpocrates — named after the Greek god of silence and secrets — is the private AI engine at the heart of Agent Midas. When you chat with MidasBot, generate content, ask your AI agents to analyze your business, or use any AI-powered feature on the platform, the majority of your queries are served by Harpocrates on dedicated AI infrastructure that Agent Midas operates directly.

This is a deliberate architectural choice with a clear privacy benefit: for Harpocrates-served interactions, your data does not travel to OpenAI, Google, Anthropic, or any third-party AI provider. It stays inside Agent Midas infrastructure.

What this means for you:

- **No training on your data.** Your queries, documents, CRM entries, financial records, meeting transcripts, and brand assets are never used as training examples — not for us, not for another subscriber, not for any outside company.
- **Complete subscriber isolation.** No other subscriber can ever see, access, or benefit from your information. Harpocrates operates within the same Row-Level Security and container isolation model that governs the rest of the platform.
- **Transparent routing.** For specialized tasks that require a different model, Agent Midas may route a query to a trusted external provider. When this happens, no subscriber-identifying information is included, and only the minimum context needed for the task is sent.
- **Dedicated AI infrastructure.** Harpocrates runs on dedicated AI infrastructure that Buji Development Corporation provisions and manages. We are not renting a multi-tenant slice of someone else's AI service — we operate the hardware that serves your queries.

Harpocrates is the reason we can make a promise that almost no other AI platform can make: when Agent Midas says your data is yours, we mean it at the infrastructure level — not just as a policy we could change tomorrow.

13. Consumer Consent Management

Buji Development Corporation obtains explicit, informed consent from every subscriber at every stage of the customer lifecycle. We do not assume consent. We do not pre-check boxes. We do not bury consent in terms of service. Every data collection event is accompanied by a clear, visible, and actionable consent mechanism.

- **Account Creation:** Subscribers must actively check a required checkbox confirming agreement to Terms of Service and Privacy Policy. The timestamp, IP address, and user agent of consent are recorded.
- **Onboarding:** Every question in the 89-question onboarding questionnaire is optional. Subscribers may skip any question without penalty. The AI generates a summary dossier which the subscriber reviews and approves before it becomes active.
- **Third-Party Integrations:** Every integration (Google, Slack, Plaid, Stripe) requires a separate, explicit subscriber-initiated action. No integrations are pre-connected or assumed.
- **Affiliate Enrollment:** Affiliates must check two separate checkboxes: Terms of Service and Affiliate Agreement with Income Disclaimer. Both are timestamped separately.
- **Communications:** Marketing communications require separate opt-in. All emails include unsubscribe mechanisms compliant with CAN-SPAM.

All consent records are stored with: consent type, timestamp, IP address, user agent, and the version of the document accepted. Consent records are retained for the lifetime of the account plus 7 years for regulatory compliance.

14. Data Retention and Deletion

Data Category	Retention Period	Deletion Trigger
Active subscriber data	Duration of subscription + 90 days	Cancellation + grace period expiry
Subscriber silo data	Duration of subscription + 90 days	Account deletion or retention expiry
Plaid tokens and banking data	Until subscriber revokes or cancels	Revocation, cancellation, or disconnection
OAuth tokens	Until subscriber revokes or cancels	Revocation or account cancellation
Affiliate commission history	Duration + 7 years	IRS 1099 reporting requirements
Audit logs and security events	7 years minimum	Automated purge after retention period
Consent records	Account lifetime + 7 years	Legal proof of consent

Subscribers may request full account deletion at any time through Settings or by contacting midas@agentmidas.xyz. Upon receiving a deletion request, all subscriber silo data is permanently deleted within 30 days. All OAuth tokens and Plaid access tokens are immediately revoked and deleted. Only records required for regulatory retention remain.

- **GDPR:** EU data subjects may exercise their right to erasure. We honor all verified requests within 30 days.
- **CCPA:** California residents may exercise their right to deletion. We honor all verified requests within 45 days.

15. Vulnerability Management, Penetration Testing & CASA Tier 2 Certification

Agent Midas maintains a continuous vulnerability management program combining automated scanning, third-party assessment, and formal certification. Our program includes the following controls:

- **Cloudflare WAF:** Continuous web application firewall protection filtering malicious traffic, SQL injection attempts, cross-site scripting (XSS), and other common attack vectors in real-time.
- **Cloudflare Page Shield:** Continuous monitoring of all third-party JavaScript loaded in subscriber browsers. Detects compromised or tampered scripts from payment processors, video conferencing, and email providers.
- **GitHub Dependabot:** Automated dependency vulnerability scanning. Security alerts are generated when known vulnerabilities are discovered in our dependencies.
- **npm Audit:** Package vulnerability scanning integrated into our CI/CD pipeline. Builds fail if critical vulnerabilities are detected.
- **Endpoint Protection:** All development machines require full-disk encryption (FileVault/BitLocker), real-time malware scanning, and OS security patches within 7 days of release.
- **Internal Security Audits:** Conducted on an ongoing basis. Our February 2026 audit identified and remediated five API route vulnerabilities within 24 hours of discovery.

CASA Tier 2 Certification (April 2026)

In April 2026, Agent Midas completed a Google CASA (Cloud Application Security Assessment) Tier 2 Static Application Security Testing (SAST) assessment conducted by TAC Security through the ESOF AppSec ADA framework. This is the same assessment framework Google requires of applications that handle sensitive user data through its OAuth platform.

Metric	Result
Assessment Framework	Google CASA Tier 2 — ESOF AppSec ADA (TAC Security)
ESOF Cyber Score	9.7 out of 10.0
Risk Classification	Low Risk (highest maturity tier)
Assessment Date	April 16, 2026
Scope	Complete Agent Midas source code (243,000 lines)

Metric	Result
Coverage	OWASP Top 10, injection flaws, authentication, cryptography, access control, data exposure
Critical Findings	0
High-Severity SAST Findings	0
Verified By	TAC Security — ADA CASA Assessment
Certification Status	Verified and Secured — ESOF Badge Issued

The ESOF Cyber Score classification scale operates as follows:

- 7.6 – 10.0: Low Risk (our score: 9.7)
- 5.6 – 7.5: Medium Risk
- 2.6 – 5.5: High Risk
- 0.1 – 2.5: Critical Risk

Our 9.7 score places Agent Midas in the top tier of the Low Risk classification. Zero critical findings and zero high-severity SAST findings were identified across a 243,000-line codebase. A small number of medium-severity findings (primarily standard React dangerouslySetInnerHTML usage for server-generated content) were documented and are being addressed through planned remediation.

All Security Assessment Questionnaire (SAQ) items were answered affirmatively, and three supporting policy documents were submitted alongside the code scan:

- Agent Midas Information Security Policy
- Agent Midas Incident Response Policy
- Agent Midas Vulnerability Disclosure Policy

Remediation of medium-severity findings is in progress. A Google CASA badge will be issued upon completion of the remediation phase. The CASA Tier 2 assessment will be reconducted annually to maintain certification status.

External Penetration Testing

External penetration testing is planned for Q2 2026, targeting the subscriber-facing application, all API routes, authentication flows, and the Plaid integration. Testing will be conducted annually thereafter. We are also preparing for SOC 2 Type II certification, targeted for completion within 12 months of this publication.

16. Incident Response

Severity	Description	Response Time
Critical (1)	Confirmed data breach affecting financial data or PII. Unauthorized access to banking tokens.	Immediate containment. CEO notification within 1 hour. Affected subscribers notified within 24 hours.
High (2)	Attempted unauthorized access detected and blocked. Vulnerability in production code.	Containment within 4 hours. Investigation within 24 hours. Remediation within 48 hours.
Medium (3)	Anomalous login patterns. Failed authentication spikes.	Investigation within 48 hours. Remediation per risk assessment.
Low (4)	Policy violation with no data impact. Configuration drift.	Documented. Addressed in next maintenance cycle.

17. Regulatory Compliance and Security Certifications

- **FTC Act:** All affiliate program income claims include explicit disclaimers. The Two-Up compensation structure is clearly disclosed in the Affiliate Agreement.
- **IRS Reporting:** Affiliates earning above $600/year receive 1099-NEC forms. Tax identification is collected through verified onboarding.
- **GLBA / Regulation P:** Subscriber financial data connected through Plaid is treated as Nonpublic Personal Information (NPI). We do not share, sell, or disclose financial data to any third party.
- **Google CASA Tier 2 (SAST) — Certified April 2026:** ESOF Cyber Score of 9.7/10.0 (Low Risk classification) achieved in April 2026 via TAC Security ADA CASA Assessment. Covers OWASP Top 10, injection flaws, authentication, and access control. See Section 15 for full details.
- **SOC 2 Reliance and Pursuit:** We rely on SOC 2 certifications of our infrastructure partners: Supabase (database), Stripe (payments), Plaid (banking), Cloudflare (network), and DigitalOcean (hosting). Agent Midas is pursuing its own SOC 2 Type II certification, targeted for completion within 12 months.
- **GDPR Preparedness:** Data processing agreements, data portability (subscriber data export), and right-to-deletion requests are supported for EU subscribers.
- **State Privacy Laws:** We monitor compliance with CCPA (California), VCDPA (Virginia), and CPA (Colorado) as our subscriber base expands.

18. Plaid Attestation Compliance (8 Requirements)

As part of our approved integration with Plaid, Buji Development Corporation is required to attest to eight specific security requirements by August 20, 2026. The following table summarizes our current compliance status and implementation timeline.

#	Attestation	Status	Target
1	Secure tokens and certificates for authentication	85% Complete	Apr 2026
2	Role-based access control (RBAC)	75% Complete	May 2026
3	Periodic access reviews and audits	40% Complete	May 2026
4	Vulnerability scanning	50% Complete	Jun 2026
5	End-of-life software monitoring	In Progress	Jun 2026
6	Automated de-provisioning	30% Complete	Jul 2026
7	Zero trust access architecture	45% Complete	Jul 2026
8	Centralized identity and access management	35% Complete	Aug 2026

All eight attestations are scheduled for completion by August 20, 2026. Our architecture already embodies most of these principles in production. The remaining work is primarily formalization, documentation, tooling, and filling specific gaps where we have the practice but not yet the formal process.

19. Continuous Monitoring and Audit

- **Append-Only Audit Log:** Every authentication event, data access event, administrative action, and security-relevant system event is recorded in an immutable audit log. No UPDATE or DELETE permissions exist on this table.
- **Login Tracking:** Every authentication records IP address, timestamp, user agent, and cumulative session count. Anomalous patterns trigger alerts.
- **System Health Checks:** Cloudflare Load Balancer health checks run every 60 seconds against the production server. Two consecutive failures trigger automatic failover.
- **Integration Monitoring:** Each third-party integration is monitored for connection health, token validity, and API response status. Failures trigger alerts and automatic retry.
- **Quarterly Audits:** RLS policy audit on all tables. API route authentication audit. Access control review.
- **Semi-Annual Reviews:** Full security policy review. Risk register update. Third-party integration security review.
- **Annual Assessment:** Comprehensive security assessment. Vendor security review. Policy document update and re-approval. CASA Tier 2 reassessment.

20. Your Rights as a Subscriber

As an Agent Midas subscriber, you have the following rights regarding your data:

- **Right to Access:** You can view all data we hold about you through your dashboard at any time.
- **Right to Export:** You can export your complete data set through Settings. Exports are delivered in standard formats (CSV, JSON).
- **Right to Correction:** You can edit any information in your profile, dossier, or onboarding answers at any time.
- **Right to Deletion:** You can request full account deletion through Settings or by emailing midas@agentmidas.xyz. Deletion is processed within 30 days.
- **Right to Revoke:** You can disconnect any third-party integration (Google, Plaid, Slack, etc.) at any time through Settings. Revocation is immediate.
- **Right to Transparency:** This whitepaper is our commitment to telling you exactly how we protect your data. If you have questions, email midas@agentmidas.xyz.

We live in a world where your phone listens to your conversations to serve you ads,

your email provider scans every message for keywords,

your social media platforms sell your behavioral data to the highest bidder,

and your free productivity apps monetize your documents.

Agent Midas does none of that.

We make money one way: your subscription. You are our customer, not our inventory.

For questions, concerns, or security inquiries, contact us at midas@agentmidas.xyz

Buji Development Corporation | 1712 Pioneer Ave. Ste. 500, Cheyenne, WY 82001 | www.agentmidas.xyz