

# END USER LICENSE AGREEMENT FOR THE AUTONOMA TECHNOLOGIES, INC.

This End User License Agreement (“Agreement”) is entered into as of the subscription start date (“Effective Date”) by and between The Autonoma Technologies, Inc. (“Autonoma”, “we”, “us”, or “our”) and the entity or individual subscribing via AWS Marketplace (“Customer”, “you”, or “your”).

## RECITALS.

**WHEREAS**, Autonoma has developed a comprehensive AI agent platform for autonomous software development, security, operations, compliance, and governance (“Autonoma Platform” or “Service”);

**WHEREAS**, Customer desires to access and use one or more Autonoma capabilities via AWS Marketplace under the terms of this Agreement;

**NOW, THEREFORE**, in consideration of the mutual covenants contained herein, the parties hereto agree as follows:

## 1. DEFINITIONS.

**1.1 “Autonoma Capabilities”** means the autonomous AI agent platforms offered by Autonoma, including:

### **Core Development and Delivery**

- **BUILD**: Autonomous development platform (Requirements AI, Planner AI, Architect AI, Coder AI, Tester AI, Review AI, Deploy AI, Orchestrator AI)

### **Operations & Reliability**

- **MAINTAIN**: Autonomous operations platform (Maintain AI, Debug AI, Incident AI, Observe AI, Billing AI, Tenant AI, DBA AI, Backup AI, Success AI)

### **Security & Compliance**

- **SECURE:** Autonomous security platform (Security AI, Compliance AI, Data Governance AI)

### **Governance & Policy**

- **GOVERN:** Autonomous governance platform (Govern AI, Policy AI)

### **Optimization & Evolution**

- **OPTIMIZE:** Autonomous performance and cost optimization platform (Optimize AI, Cost AI, Capacity AI)

- **EVOLVE:** Autonomous improvement platform (Evolve AI, Adapt AI)

### **Bundled Offerings**

- **PLATFORM:** Combined BUILD + MAINTAIN + SECURE capabilities

- **ENTERPRISE:** All capabilities (BUILD + MAINTAIN + SECURE + OPERATE + GOVERN + OPTIMIZE + EVOLVE)

**1.2 “Subscribed Capabilities”** means the specific Autonomia Capabilities Customer has purchased via AWS Marketplace.

**1.3 “AI Agents”** means the specialized autonomous agents included in Customer’s Subscribed Capabilities.

**1.4 “API Keys”** means the authentication credentials (format: au\_live\_\* or au\_test\_\*) issued to Customer for programmatic access to the Service.

**1.5 “Customer Data”** means all the code, documentation, specifications, infrastructure configurations, security policies, and other materials uploaded to or processed by the Service.

**1.6 “Generated Outputs”** means the source code, configurations, documentation, security reports, policies, recommendations, and other content autonomously created by AI agents.

**1.7 “RIGOR Framework”** means Autonomia’s proprietary Research → Inspect → Generate → Optimize → Review autonomous reasoning methodology used across all AI agents.

**1.8 “Usage Dimensions”** means the metered consumption metrics specific to each capability (e.g., builds, deployments, security scans, incidents resolved, policy evaluations).

**1.9 “Agent Capability Matrix Framework”** means Autonomia’s proprietary 4-level autonomy progression system that governs how AI agents operate:

- **Level 1 - Learning:** Agents observe Customer workflows and build understanding without taking action

- **Level 2 - Recommendation:** Agents suggest actions but require explicit Customer approval before execution

- **Level 3 - Assisted (Human-in-the-loop):** Agents execute actions with real-time Customer oversight and intervention capability

- **Level 4 - Autonomous:** Agents execute actions independently within Customer-defined guardrails and policies

**1.10 “Autonomy Level”** means the current capability level at which a specific AI Agent or workflow operates, as configured by Customer.

**1.11 “Cross-Agent Learning”** means Autonomia’s shared knowledge graph system where insights, patterns, and learnings from one AI agent are immediately propagated to all other Customer-authorized agents within Customer’s subscription.

**1.12 “Shared Knowledge Graph”** means the distributed learning system that enables real-time knowledge sharing across all AI agents within Customer’s subscription, improving performance and consistency across all Subscribed Capabilities.

**1.13 “Agent-Capability Maturity”** means the progression of an AI Agent’s autonomy level based on demonstrated performance, Customer trust settings, and domain-specific success metrics.

**1.14 “Emergent Behavior”** means unexpected but legitimate patterns or actions that arise from cross-agent learning and autonomous operations, not explicitly programmed but resulting from agent interactions with Shared Knowledge Graph.

## **2. LICENSE GRANT & SCOPE.**

### **2.1 SaaS License Grant**

Subject to Customer’s compliance with this Agreement and payment of applicable fees, Autonoma grants Customer a non-exclusive, non-transferable, worldwide license to:

- (a) Access and use Subscribed Capabilities via web interface and API
- (b) Integrate API calls into Customer’s workflows
- (c) Utilize AI Agents for autonomous tasks within Subscribed Capabilities
- (d) Export and use Generated Outputs in Customer’s projects

### **2.2 Capability-Specific Access**

Customer’s access is limited to Subscribed Capabilities purchased via AWS Marketplace. Access to additional capabilities requires separate subscription.

### **2.3 License Restrictions**

Customer shall NOT:

- (a) Reverse engineer, decompile, or disassemble the Service, AI Agents, or RIGOR Framework
- (b) Use the Service for competitive analysis or to build competing AI agent platforms
- (c) Share API Keys with unauthorized third parties or across legal entities
- (d) Exceed contracted developer seats or resource limits without upgrading subscription
- (e) Circumvent usage metering or attempt to manipulate dimension tracking
- (f) Use AI Agents for illegal, harmful, or unethical purposes
- (g) Extract, replicate, or reverse-engineer the RIGOR Framework methodology

(h) Resell or sublicense the Service to third parties

## **2.4 Developer Seats & Resource Allocation**

Each developer seat grants one individual access to AI Agents within Subscribed Capabilities. Seats are non-transferrable and must be assigned to specific individuals. Shared credentials or seat sharing is prohibited.

## **3. AI-GENERATED OUTPUTS & INTELLECTUAL PROPERTY**

### **3.1 Ownership of Generated Outputs**

Customer owns all Generated Outputs produced by AI Agents using Customer Data as input, including:

- Source code and infrastructure configurations (BUILD)
- Security remediation scripts and policies (SECURE)
- Runbooks and automation workflows (MAINTAIN)
- Governance policies and compliance reports (GOVERN)
- Improvement strategies and analytics (EVOLVE)

Autonoma retains no ownership rights to Generated Outputs.

### **3.2 AI Output Disclaimer**

**IMPORTANT:** AI Agents generate outputs autonomously. Customer acknowledges:

(a) **No Warranty:** Generated Outputs are provided “AS IS” without warranty of correctness, security, completeness, or suitability

(b) **Review Obligation:** Customer must review all Generated Outputs before production deployment or implementation

(c) **Customer Responsibility:** Customer is solely responsible for:

- Testing and validating generated code (BUILD)
- Verifying security recommendations (SECURE)
- Confirming operational changes (MAINTAIN)

- Reviewing governance policies (GOVERN)
- Validating improvement strategies (EVOLVE)

(d) **No Liability:** Autonomia is not liable for bugs, vulnerabilities, operational failures, compliance gaps, or business impacts from Generated Outputs

### **3.3 Autonomia Intellectual Property**

Customer acknowledges that Autonomia exclusively owns:

- (a) The Service, all AI Agents, and underlying technology
- (b) RIGOR Framework and autonomous reasoning methodologies
- (c) Usage analytics, performance metrics, and aggregated data across all capabilities
- (d) All improvements, updates, and derivative works of the Service
- (e) Training data, model architectures, and agent coordination logic

### **3.4 Customer Data Ownership**

Customer retains all ownership rights to Customer Data. Autonomia receives a limited license to process Customer Data solely to provide Subscribed Capabilities.

### **3.5 Aggregated Analytics**

Autonomia may use anonymized, aggregated usage data for:

- Service improvements
- Industry benchmarking
- Research and development

Anonymized data contains no Customer-identifying information or proprietary business details.

## **4. DATA PROCESSING & PRIVACY**

### **4.1 Data Processing by Capability**

Autonomia processes Customer Data using AI Agents for:

- BUILD: Code analysis, generation, testing, deployment automation

- MAINTAIN: Infrastructure monitoring, incident detection, debugging, optimization
- SECURE: Vulnerability scanning, threat detection, compliance checking, data governance
- OPERATE: Tenant provisioning, database management, capacity planning, backup automation, billing metering, customer success analytics
- GOVERN: Risk assessment, policy generation, audit automation, compliance tracking, OPA integration
- OPTIMIZE: Performance tuning, cost analysis and reduction, capacity forecasting, resource optimization
- EVOLVE: System evolution planning, adaptive infrastructure, continuous improvement recommendations

#### **4.2 Data Retention**

- Active Subscription: Customer Data accessible in real-time
- Logs & Telemetry: Retained for 30 days
- Usage Analytics: Retained for 90 days (anonymized)
- Security Audit Logs: Retained for 1 year (compliance requirement)
- Post-Termination: Customer Data deleted within 30 days

#### **4.3 Data Export**

Customer may export Customer Data and Generated Outputs at any time via API or web interface. Pre-termination export is Customer's responsibility.

#### **4.4 Privacy Compliance**

Autonoma complies with:

- GDPR (EU General Data Protection Regulation)
- CCPA (California Consumer Privacy Act)
- SOC 2 Type II requirements
- ISO 27001 security standards

See Privacy Policy at <https://www.theautonoma.io/legal/privacy>

#### 4.5 No Training on Customer Data

Customer Data is not used to train AI models or improve AI Agents for other customers without explicit written consent.

#### 4.6 Data Residency

Customer Data processed with AWS us-east-1 region unless otherwise specified in subscription. Multi-region support available for PLATFORM and GOVERN subscriptions.

### 5. SERVICE LEVEL AGREEMENT (SLA)

#### 5.1 Uptime Commitment

Autonoma commits to 99.5% monthly uptime for AI Agent Availability across all Subscribed Capabilities.

#### 5.2 Service Credits

Uptime Achievement	Service Credit
99.0% - 99.5%	10% monthly fee
98.0% - 99.0%	15% monthly fee
Below 98.0%	25% monthly fee

#### 5.3 Capability-Specific SLAs

- BUILD: 99.5% agent availability, <5min build queue time (P95), <30min deployment time (P95)
- MAINTAIN: 99.9% monitoring uptime, <1min incident detection (P95), <10min debugging time (P95)
- SECURE: 99.5% scanning availability, <24hr vulnerability report
- OPERATE: <15min tenant creation and infrastructure provisioning, 99.9% billing metering accuracy (AWS Marketplace integration), <1hr database optimization recommendations, <15min rollback completion for standard datasets (P95)

- GOVERN: 99.5% policy evaluation uptime, <1hr compliance report, real-time policy enforcement
- EVOLVE: 99.0% analytics availability, monthly evolution plans, adaptive scaling response <5min
- OPTIMIZE: <4hr performance bottleneck identification, weekly optimization recommendations, monthly capacity forecasts, <1hr cost saving opportunity alerts

#### **5.4 Planned Maintenance**

Scheduled maintenance windows (announced 7 days in advanced) are excluded from SLA calculations. Maintenance limited to 4 hours/month.

#### **5.5 SLA Exclusions**

SLAs does not cover:

- Customer's internet connectivity issues
- Third-party API failures (AWS, GitHub, etc.)
- Force majeure events
- Customer-caused outages or misconfigurations
- Beta/experimental features clearly marked as such

#### **5.6 SLA Credit Claims**

Customer must request SLA credits within 30 days of the incident via support@theautonoma.io with supporting evidence.

### **6. USAGE METERING & BILLING**

#### **6.1 Contract Dimensions**

Customer's base subscription includes:

- Developer Seats: As specified in AWS Marketplace contract
- Intelligence Tier: CORE, PRO, or ULTRA (per contract selection)
- Capability Access: Specific to Subscribed Capabilities

## **6.2 Included Usage Quotas (Per Developer Per Month)**

Usage quotas vary by capability:

### **BUILD:**

- 60 builds, 30 RIGOR cycles, 30 code reviews, 12 deployments, 240 agent hours

### **MAINTAIN:**

- 24/7 monitoring, 50 incidents, 30 debug sessions, 20 RIGOR cycles, 160 agent hours

### **SECURE:**

- 30 vulnerability scans, 10 compliance checks, 20 RIGOR cycles, 120 agent hours

### **OPERATE:**

- 20 tenant provisions, 100 database optimizations, monthly capacity plans, daily backups,  
120 agent hours

### **GOVERN:**

- 100 policy evaluations, 10 audit reports, 15 RIGOR cycles, 80 agent hours

### **EVOLVE:**

- Daily analytics, 5 evolution cycles, weekly improvement recommendations, 10 RIGOR  
cycles, 60 agent hours

### **OPTIMIZE:**

- Daily cost analysis, 10 performance optimization cycles, monthly capacity forecasts, 100  
agent hours

### **PLATFORM:**

- Combined quota from all three categories (BUILD, MAINTAIN, SECURE)

### **ENTERPRISE:**

- Combined quota from all seven categories (BUILD, MAINTAIN, SECURE, OPERATE,  
GOVERN, EVOLVE, OPTIMIZE)

## **6.3 Overage Charges**

Usage beyond quotas billed per Usage Dimensions specific to each capability. See AWS Marketplace listing for overage pricing by Intelligence Tier.

### 6.4 Metering Accuracy

Autonoma uses AWS Marketplace BatchMeterUsage API for consumption tracking. Customer may dispute charges within 30 days with supporting evidence.

### 6.5 Intelligence Tier Pricing

All Subscribed Capabilities use the selected Intelligence Tier (CORE/PRO/ULTRA). Tier changes apply to all capabilities and take effect on next billing cycle.

## 7. API USAGE & SECURITY

### 7.1 API Key Management

Customer must:

- Store API Keys securely (encrypted, secrets management systems)
- Rotate keys every 90 days (recommended best practice)
- Immediately revoke compromised keys via dashboard or API
- Not embed keys in public repositories or client-side code
- Use environment-specific keys (test vs production)

### 7.2 Rate Limiting by Intelligence Tier

API calls subject to rate limits:

Intelligence Tier	Requests/Minute	Concurrent Agents
CORE	100	3
PRO	250	5
ULTRA	500	10

### 7.3 API Authentication

All API requests must include valid API Key in Authorization header:

Authorization: Bearer au\_live\_XXXXXXXXXXXX

### 7.4 Customer Liability for Unauthorized Use

Customer liable for all usage via their API Keys, including unauthorized access resulting from Customer's security failures or credential leakage.

## **7.5 Webhook Security**

For integrations with MAINTAIN and GOVERN capabilities, Customer must:

- Validate webhook signatures
- Use HTTPS endpoints only
- Implement proper authentication

## **8. CAPABILITY-SPECIFIC TERMS**

### **8.1 BUILD Capability**

- (a) Code Generation: AI-generated code provided without warranty of fitness, correctness, or security
- (b) Deployment Automation: Customer responsible for production deployment validation
- (c) Testing: AI-generated tests supplement but do not replace Customer's testing obligations

### **8.2 MAINTAIN Capability**

- (a) Monitoring Scope: Limited to infrastructure and applications Customer configures
- (b) Incident Response: AI recommendations require Customer approval before execution
- (c) Performance Impact: Autonomia not liable for performance degradation from AI recommendations

### **8.3 SECURE Capability**

- (a) Vulnerability Detection: AI scans identify potential vulnerabilities; Customer responsible for remediation
- (b) False Positives: Customer acknowledges AI may flag non-issues; manual review required
- (c) Compliance: AI-generated compliance reports do not constitute legal or audit advice
- (d) Zero-Day Threats: Autonomia not liable for undetected zero-day vulnerabilities

#### **8.4 GOVERN Capability**

- (a) Policy Generation: AI-generated policies require legal/compliance review before implementation
- (b) Risk Assessment: AI risk scores are recommendations, not guarantees
- (c) Regulator Compliance: Customer responsible for final compliance determinations

#### **8.5 EVOLVE Capability**

- (a) Strategic Recommendations: AI suggestions are not business, financial, or legal advice
- (b) Predictive Analytics: Historical performance does not guarantee future results
- (c) Learning Models: Optimization recommendations based on Customer's usage patterns

#### **8.6 OPERATE Capability**

- (a) Tenant Provisioning: AI-automated infrastructure provisioning requires Customer approval for production environments
- (b) Database Management: DBA AI recommendations supplement but do not replace Customer's database administration responsibilities
- (c) Billing Accuracy: Billing AI metering guaranteed 99.9% accurate; discrepancies resolved within 48 hours
- (d) Backup Integrity: Backup AI follows 3-2-1 rule; Customer responsible for validating restore procedures
- (e) Capacity Planning: Capacity forecasts are statistical projections, not guarantees

#### **8.7 OPTIMIZE Capability**

- (a) Cost Savings: Cost AI recommendations; actual savings may vary
- (b) Performance Impact: Performance optimizations may require testing before production deployment
- (c) Capacity Forecasts: Based on historical usage patterns; unexpected load spikes may exceed forecasts

(d) No Guarantee: Optimization recommendations do not guarantee specific cost savings or performance improvements

## **9. WARRANTIES & DISCLAIMERS**

### **9.1 Service Warranties**

Autonoma warrants that:

- (a) Service will perform substantially as described in documentation
- (b) Service will comply with applicable laws
- (c) Autonoma has rights to provide the Service
- (d) Service protected by industry-standard security measures

### **9.2 DISCLAIMER OF WARRANTIES**

EXCEPT AS EXPRESSLY STATED IN SECTION 9.1, THE SERVICE IS PROVIDED “AS IS” AND “AS AVAILABLE”. AUTONOMA DISCLAIMS ALL OTHER WARRANTIES, EXPRESS OR IMPLIED, INCLUDING:

- WARRANTY OR MERCHANTABILITY
- FITNESS FOR A PARTICULAR PURPOSE
- NON-INFRINGEMENT
- ERROR-FREE OR UNINTERRUPTED OPERATION
- ACCURACY OR RELIABILITY OF AI-GENERATED OUTPUTS
- DETECTION OF ALL VULNERABILITIES, THREATS, OR RISKS
- COMPLIANCE WITH ALL APPLICABLE REGULATIONS

### **9.3 AI-Specific Disclaimers**

AI AGENTS OPERATE AUTONOMOUSLY. AUTONOMA DOES NOT WARRANT THAT:

#### **BUILD Capability:**

- Generated code will be bug-free, secure, or optimal
- Architecture designs will meet all performance requirements

- Test coverage will identify all defects
- Deployments will succeed without errors
- Code reviews will catch all issues
- Sprint plans will accurately estimate effort
- RIGOR reasoning will identify all technical risks

**MAINTAIN Capability:**

- All incidents will be detected or correctly diagnosed
- Root cause analysis will identify actual causes
- Debugging recommendations will resolve all issues
- System monitoring will detect all anomalies
- Incident response will prevent all downtime
- Observability insights will reveal all performance bottlenecks

**OPERATE Capability:**

- Tenant provisioning will complete without errors or meet all requirements
- Database optimizations will improve performance or not cause regressions
- Capacity forecasts will accurately predict future resource needs
- Backups will be complete, uncorrupted, or successfully restorable
- Billing metering will be 100% accurate for all usage dimensions
- Customer success predictions will be correct or actionable
- Infrastructure-as-code will deploy without conflicts
- DBA recommendations will not cause database downtime or data loss

**SECURE Capability:**

- All vulnerabilities or threats will be identified
- Vulnerability scans will have zero false positives or false negatives
- Compliance checks will ensure regulatory adherence

- Data governance policies will prevent all data breaches
- Security recommendations will eliminate all security risks

**GOVERN Capability:**

- Governance policies will ensure complete compliance
- Risk assessments will identify all business risks
- Policy-as-code will enforce all organizational requirements
- Audit reports will be comprehensive and legally sufficient
- OPA integrations will prevent all policy violations

**OPTIMIZE Capability:**

- Cost optimizations will achieve projected savings
- Performance tuning will improve system performance
- Capacity forecasts will accurately predict resource requirements
- Rightsizing recommendations will not cause performance degradation
- Cost analysis will identify all savings opportunities
- Optimization cycles will not introduce new bottlenecks or inefficiencies
- AWS Cost Explorer integrations will capture all cost anomalies

**EVOLVE Capability:**

- Recommendations will achieve desired business outcomes
- Strategic improvements will succeed as planned
- System evolution will not introduce breaking changes
- Adaptive scaling will handle all traffic patterns
- Predictive analytics will accurately forecast trends
- Learning models will continuously improve without regressions

**Cross-Capability Disclaimers:**

- RIGOR reasoning will identify all issues, risks, or opportunities across any capability

- Multi-agent coordination will never produce conflicting recommendations
- Agent-capability maturity progression will always improve performance
- Cross-agent learning will propagate only correct patterns
- Emergent behaviors will always be beneficial or expected

### **CUSTOMER ACKNOWLEDGMENT:**

Customer explicitly acknowledges and accepts that AI Agents:

1. Make mistakes at all autonomy levels and across all capabilities
2. Provide recommendations, not guarantees - all outputs require Customer validation
3. Learn from experience - early recommendations may be less accurate than later ones
4. Operate probabilistically - no AI system achieves 100% accuracy
5. Require human oversight - particularly at Level 3 (Assisted) and Level 4 (Autonomous)
6. Cannot replace human expertise - AI augments but does not replace Customer's responsibilities
7. May produce unexpected results - especially with cross-agent learning and emergent behaviors

### **9.3.1 Third-Party Integration Disclaimers**

Autonoma is NOT responsible for:

- AWS service outages affecting OPERATE, OPTIMIZE, or any capability
- Large Language Model API limitations affecting AI Agent quality or availability
- Incident Management outages affecting MAINTAIN incident management
- GitHub/GitLab availability affecting BUILD deployments
- Velero backup failures in OPERATE capability
- AWS Marketplace metering delays affecting OPERATE billing
- Cost Explorer data lag affecting OPTIMIZE recommendations
- Auto Scaling service issues affecting EVOLVE adaptive capabilities

### **9.3.2 Data Quality Disclaimers**

AI Agent accuracy depends on Customer Data quality:

- Incomplete or inaccurate Customer Data may produce incorrect outputs
- Missing context may result in suboptimal recommendations
- Outdated infrastructure configurations may cause OPERATE provisioning errors
- Historical data gaps may reduce OPTIMIZE and EVOLVE forecast accuracy
- Insufficient telemetry may limit MAINTAIN incident detection

### **9.3.3 Tiered Intelligence Disclaimers**

Intelligence tier affects capability performance:

- CORE tier: Basic AI reasoning; higher error rates expected
- PRO tier: Advanced reasoning; improved accuracy but not guaranteed
- ULTRA tier: Cutting-edge reasoning; best available accuracy but still probabilistic

Lower tiers may produce:

- Less accurate capacity forecasts (OPTIMIZE, OPERATE)
- Slower incident detection (MAINTAIN)
- More false positives in security scans (SECURE)
- Less sophisticated code generation (BUILD)
- Simpler governance policies (GOVERN)
- Limited evolution strategies (EVOLVE)

### **9.3.4 Autonomy Level Risk Disclosure**

Risk increases with autonomy level across ALL capabilities:

Highest Risk Autonomous Operations:

- OPERATE tenant provisioning (Level 4): May provision incorrect infrastructure
- MAINTAIN incident response (Level 4): May execute incorrect remediation
- BUILD deployments (Level 4): May deploy broken code

- OPTIMIZE cost changes (Level 4): May cause service degradation

## **10. LIMITATION OF LIABILITY**

### **10.1 Liability Cap**

AUTONOMA'S TOTAL LIABILITY UNDER THIS AGREEMENT SHALL NOT EXCEED THE AMOUNT PAID BY CUSTOMER IN THE 12 MONTHS PRECEDING THE CLAIM.

### **10.2 Excluded Damages**

IN NO EVENT SHALL AUTONOMA BE LIABLE FOR:

- INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES
- LOST PROFITS, REVENUE, OR BUSINESS OPPORTUNITIES
- COST OF SUBSTITUTE SERVICES OR PROCUREMENT
- DAMAGES FROM GENERATED OUTPUTS (code, policies, configurations)
- FAILURES IN CUSTOMER'S PRODUCTION SYSTEMS
- UNDETECTED SECURITY VULNERABILITIES OR THREATS
- REGULATORY NON-COMPLIANCE OR AUDIT FAILURES
- BUSINESS DECISIONS BASED ON AI RECOMMENDATIONS

### **10.3 Capability-Specific Liability Exclusions**

- BUILD: Not liable for production bugs, deployment failures, or code vulnerabilities in  
Generated Outputs
- MAINTAIN: Not liable for undetected incidents, incorrect diagnosis, or operational  
disruptions
- SECURE: Not liable for undetected threats, compliance gaps, or security breaches
- GOVERN: Not liable for regulatory violations, audit failures, or policy inadequacies
- EVOLVE: Not liable for business outcomes, strategic failures, or predictive inaccuracies
- OPERATE: Not liable for tenant provisioning errors, database performance issues, backup  
failures, capacity planning inaccuracies, or billing disputes

- OPTIMIZE: Not liable for cost overruns, performance degradation from optimizations, capacity forecasting errors, or unachieved savings targets

#### **10.4 Exceptions to Limitations**

Liability limitations do NOT apply to:

- Customer's violation of Section 2.3 (License Restrictions)
- Customer's unauthorized sharing of API Keys
- Gross negligence or willful misconduct by either party
- Indemnification obligations under Section 10

### **11. INDEMNIFICATION**

#### **11.1 Autonomia's Indemnification**

Autonomia will indemnify Customer against third party claims that the Service infringes intellectual property rights, provided Customer:

- (a) Promptly notifies Autonomia of the claim
- (b) Grants Autonomia sole control of defense and settlement
- (c) Cooperates with Autonomia's defense efforts

#### **11.2 Customer's Indemnification**

Customer will indemnify Autonomia against claims arising from:

- (a) Customer Data content or use
- (b) Customer's violation of this Agreement
- (c) Generated Outputs deployed by Customer
- (d) Customer's breach of API usage terms
- (e) Customer's regulatory non-compliance
- (f) Third-party claims related to Customer's use of AI recommendations

### **12. TERM & TERMINATION**

#### **12.1 Subscription Term**

Agreement commences on Effective Date and continues per AWS Marketplace subscription term (monthly or annual) for each Subscribed Capability.

### **12.2 Termination for Breach**

Either party may terminate on 30 days' written notice if the other party materially breaches this Agreement and fails to cure within 30 days of notification.

### **12.3 Termination of Convenience**

Customer may cancel via AWS Marketplace per the Refund Policy (see Section 13).

Autonoma may terminate on 90 days' notice for business reasons.

### **12.4 Effect of Termination**

Upon termination:

- (a) Customer's access to Service immediately ceases
- (b) Customer must cease all API usage and revoke API keys
- (c) Customer may export data within 30 days
- (d) Autonoma will delete Customer Data within 30 days (except anonymized analytics)
- (e) Customer retains rights to previously Generated Outputs
- (f) Usage charges through termination date remain due

### **12.5 Survival**

Sections 3 (IP), 4 (Data), 8 (Disclaimers), 9 (Warranties), 10 (Liability), 11 (Indemnification), and 14 (Miscellaneous) survive termination.

## **13. REFUND POLICY**

### **13.1 Satisfaction Guarantee**

30-day satisfaction guarantee on initial purchases. See complete Refund Policy at

<https://www.theautonoma.io/legal/refund-policy>

### **13.2 Refund Requests**

Email: [support@theautonoma.io](mailto:support@theautonoma.io)

Include: AWS Marketplace Customer ID, subscription details, reason for request

### **13.3 Processing Timeline**

- Acknowledgement: 24 business hours
- Decision: 5-7 business days
- Refund processing: 10 business days via AWS Marketplace

## **14. MISCELLANEOUS**

### **14.1 Entire Agreement**

This Agreement, together with AWS Marketplace terms, Privacy Policy, and Refund Policy, constitutes the entire Agreement between parties.

### **14.2 Amendments**

Autonoma may modify this Agreement on 30 days' notice for material changes. Continued use after notice constitutes acceptance. Notification via email to AWS Marketplace account email.

### **14.3 Governing Law**

This Agreement is governed by the laws of the State of Delaware, excluding conflicts of law provisions.

### **14.4 Dispute Resolution**

Disputes will be resolved via binding arbitration under AAA Commercial Arbitration Rules in Wilmington, Delaware. Either party may seek injunctive relief in court for IP violations.

### **14.5 Assignments**

Customer may not assign this Agreement without Autonoma's prior written consent.

Autonoma may assign to affiliates or in connection with merger, acquisition, or sale of substantially all assets.

### **14.6 Force Majeure**

Neither party liable for delays caused by events beyond reasonable control (natural disasters, war, terrorism, pandemics, government actions, infrastructure failures, third-party service outages).

#### **14.7 Severability**

If any provision is found unenforceable, remaining provisions remain in effect and enforceable provision will be substituted to reflect original intent.

#### **14.8 Waiver**

Failure to enforce any provision does not constitute waiver of that provision or any other provision.

#### **14.9 Export Compliance**

Customer agrees to comply with all applicable export laws and regulations. Service may not be exported to embargoed countries or sanctioned entities.

#### **14.10 Government End Users**

If Customer is a U.S. Government entity, the Service constitutes “commercial computer software” subject to restricted rights under FAR 12.212 and DFARS 227.7202.

#### **14.11 Notices**

Legal notices sent to:

- Autonomia: [legal@theautonomia.io](mailto:legal@theautonomia.io)
- Customer: Email address on AWS Marketplace account

#### **14.12 Independent Contractors**

Parties are independent contractors. This Agreement does not create partnership, joint venture, or employment relationship.

### **15. CONTACT INFORMATION**

**The Autonomia Technologies, Inc.**

Website: <https://www.theautonomia.io>

Support and Refunds: support@theautonoma.io

Legal: legal@theautonoma.io

Security: security@theautonoma.io

AWS Marketplace:

<https://aws.amazon.com/marketplace/seller-profile?id=seller-smdnyfrrlqegu>

**BY SUBSCRIBING VIA AWS MARKETPLACE, CUSTOMER ACKNOWLEDGES HAVING**

**READ, UNDERSTOOD, AND AGREED TO BE BOUND BY THIS AGREEMENT.**

Document Version: 1.0

Last Updated: December 11, 2025

Applicable To: All Autonoma capabilities (BUILD, MAINTAIN, SECURE, PLATFORM, OPERATE, GOVERN, EVOLVE, OPTIMIZE)