LUNA BASE SERVICE LEVEL AGREEMENT

Effective Date: August 1, 2025
Version: 1.0
Last Updated: August 1, 2025

1. INTRODUCTION AND APPLICABILITY

1.1 Purpose and Scope. This Service Level Agreement ("SLA") defines the service level commitments, measurement methodologies, and remedies applicable to Luna Base Inc.'s ("Luna") provision of the Luna AI software development platform services (the "Services"), including Luna Base Cloud, Luna Base IDE, and Luna Autopilot. This SLA applies only to Customers who have explicitly selected SLA coverage in their Enterprise Agreement or have purchased a subscription tier that includes SLA commitments. Free tier users, trial users, and subscriptions without explicit SLA coverage receive no service level commitments, and the Services are provided on an "as-available" basis without any guarantees.

1.2 Default Service Tier. Unless explicitly specified otherwise in an executed Enterprise Agreement or subscription order, all SLA commitments default to the Standard Tier service levels. Customers seeking Premium or Elite tier service levels must have such tiers explicitly documented in their commercial agreement with Luna. In the absence of clear tier designation, Standard Tier commitments, response times, and credit conversion rates apply.

1.3 Precedence and Integration. This SLA is incorporated by reference into and forms part of the applicable Enterprise Agreement, Terms of Service, or other master agreement between Luna and Customer (the "Agreement"). In the event of any conflict between this SLA and the Agreement, the Agreement shall prevail except with respect to technical service level specifications explicitly defined herein. This SLA supersedes all prior service level commitments, whether written or oral.

1.4 Modification Rights. Luna reserves the right to modify this SLA at any time upon ninety (90) days' written notice to Customer. Modifications shall not apply retroactively to claimed SLA breaches occurring before the modification effective date. If Customer objects to any modification that materially degrades the service level commitments, Customer may terminate the affected Services upon thirty (30) days' written notice, provided such notice is given within thirty (30) days of receiving Luna's modification notice.

2. SERVICE AVAILABILITY COMMITMENTS

2.1 Availability Scope. Luna's availability commitments apply only to the Luna-controlled application layer, which includes:

- Luna's application servers and microservices

- API gateway and routing logic

- Authentication and authorization services

- Project management and storage logic (but not underlying cloud storage)

- Credit consumption and billing services

- **Luna's orchestration layer for AI agent coordination (but NOT the third-party AI models themselves)**

Explicitly excluded from availability measurements:

- Underlying cloud infrastructure (AWS, Azure, GCP)

- **Third-party AI model providers (OpenAI, Anthropic, Google, Meta, Mistral, Cohere, etc.)**

- **AI model API availability, performance, or quality**

- Content delivery networks (CDNs)

- Third-party monitoring or analytics services

- Internet backbone and ISP networks

Critical AI Service Clarification: Luna's "AI Agent Processing" availability metric in Section 2.2 refers solely to Luna's ability to receive, route, and process requests through its orchestration layer. It does NOT guarantee that third-party AI models will be available to fulfill those requests. If Luna's orchestration layer is functional but all AI providers are down, Luna has met its SLA obligations.

2.2 Availability Targets. Subject to the exclusions and conditions set forth herein, Luna commits to the following Monthly Uptime Percentages for Luna-controlled services during each calendar month. Unless Customer's Enterprise Agreement explicitly specifies Premium or Elite tier coverage, Standard Tier targets apply:

Service Component	Standard Tier (Default)	Premium Tier	Elite Tier
Luna Base Cloud Platform	99.0%	99.5%	99.9%
AI Agent Processing	98.5%	99.0%	99.5%
API Endpoints	99.0%	99.5%	99.9%
Authentication Services	99.5%	99.5%	99.9%
Project Storage & Retrieval	99.0%	99.5%	99.9%

2.3 Service-Specific Considerations. The Luna Base IDE (Visual Studio Code extension) availability is dependent on both Luna's cloud infrastructure and the local development environment. SLA commitments apply only to Luna's cloud services that the IDE connects to, not to the IDE software itself or local functionality. Luna Autopilot availability is measured separately from the core platform and follows the same tier-based targets as the Luna Base Cloud Platform.

2.4 Availability Calculation. Monthly Uptime Percentage is calculated as:

Monthly Uptime % = (Total Minutes in Month - Downtime Minutes) / Total Minutes in Month × 100

Where "Downtime" means a period of five (5) consecutive minutes during which the affected Service is unavailable to Customer. Partial minutes of unavailability and intermittent issues lasting less than five

consecutive minutes are not considered Downtime. Availability is measured from Luna's monitoring systems using synthetic transactions from multiple geographic locations.

2.5 Multi-Cloud Architecture Considerations. Luna employs a multi-cloud architecture using AWS, Azure, and GCP to enhance resilience. However:

- Luna's ability to failover between cloud providers depends on the specific service and technical constraints

- Not all services are replicated across all cloud providers

- Failover, when available, may result in temporary degradation or feature limitations

- Customer data residency requirements may prevent cross-region or cross-provider failover

Luna's SLA commitments assume reasonable use of multi-cloud capabilities but do not guarantee immunity from cloud provider failures. Availability is measured based on Luna's application layer functioning on at least one cloud provider, not simultaneous availability across all providers.

3. SERVICE LEVEL EXCLUSIONS

3.1 General Exclusions. The following are expressly excluded from Downtime calculations and do not count against availability commitments:

- Scheduled maintenance performed during weekend windows with proper notice as specified in Sections 3.4-3.6

- Emergency maintenance required to address critical security vulnerabilities or prevent imminent system failure

- Minor updates and patches that do not impact availability

- **Any issues caused by third-party AI model providers (OpenAI, Anthropic, etc.) including outages, API failures, rate limiting, or performance degradation**

- Downtime caused by factors outside Luna's reasonable control, including Force Majeure events

- Problems arising from Customer's internet connectivity, networks, or devices

- Downtime resulting from Customer's violation of the Agreement or Acceptable Use Policy

- Service suspensions due to non-payment or credit exhaustion

- Customer-caused issues including misconfiguration, API abuse, or exceeding rate limits

- Beta features, preview releases, or experimental functionalities

- Degradation due to Customer exceeding purchased capacity or Credit allocations

3.2 Cloud Infrastructure Provider Exclusions. Luna's Services operate on third-party cloud infrastructure providers including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). Downtime or performance degradation caused by these underlying infrastructure providers is excluded from SLA calculations, including but not limited to:

- Regional or availability zone failures of cloud providers

- Cloud provider network, compute, storage, or database service outages

- DNS failures or routing issues at the cloud provider level

- Cloud provider API or control plane failures

- Degradation due to cloud provider capacity constraints or throttling

Luna commits only to the availability of its application layer and services running on top of cloud infrastructure. When cloud provider outages occur, Luna will:

- Use commercially reasonable efforts to failover to alternative regions or providers where technically feasible

- Provide timely communication about the infrastructure impact

- Work with cloud providers to restore service as quickly as possible

- Pass through a proportional share of any cloud provider SLA credits received by Luna to affected Enterprise customers, where such credits are directly attributable to Customer's impacted services

3.3 Third-Party AI Model Exclusions. Luna integrates with various third-party AI providers (including but not limited to OpenAI, Anthropic, Google, Meta, Mistral, Cohere, and other proprietary and open-source model providers) to deliver AI agent functionality. The following are completely excluded from SLA calculations and Service Credit eligibility:

- **Complete AI Provider Outages:** When third-party AI services are entirely unavailable or unreachable

- **API Failures:** Including rate limiting, throttling, or API errors from third-party AI providers

- **Model Performance Degradation:** Slow response times, timeout errors, or capacity constraints at the AI provider level

- **Model Quality Issues:** Degraded output quality, hallucinations, refusals, or unexpected model behavior changes

- **Model Deprecation or Changes:** When AI providers modify, update, replace, or discontinue models

- **Cost or Pricing Changes:** Service modifications due to third-party AI provider pricing changes

- **Geographic Restrictions:** AI services unavailable in certain regions due to provider policies

- **Content Filtering:** Rejections or blocks imposed by third-party AI safety systems

- **Token Limitations:** Constraints imposed by AI providers on request/response sizes

- **Scheduled or Unscheduled AI Provider Maintenance:** Any maintenance windows of third-party AI services

Luna commits only to maintaining the availability of its routing and orchestration layer that connects to these AI providers. Luna is not responsible for the availability, performance, quality, or behavior of third-party AI models themselves. When third-party AI services are unavailable, Luna will use commercially reasonable efforts to route requests to alternative providers where technically feasible, but makes no guarantee of equivalent functionality or performance. Customer acknowledges that AI model availability is entirely outside Luna's control and that Luna cannot be held liable for third-party AI service disruptions.

3.4 Planned Maintenance Windows. Luna reserves the right to perform planned maintenance during non-peak hours on Saturdays and Sundays as determined by Luna in its sole discretion based on platform usage patterns, geographic distribution of users, and operational requirements. Luna will use commercially reasonable efforts to schedule maintenance during periods of lowest weekend activity to minimize disruption, though Customer acknowledges that global usage patterns may make it impossible to avoid all user impact.

- **Standard Tier (Default) / Premium Tier:** Maintenance scheduled during Luna-determined non-peak hours on weekends (Saturday-Sunday)

- **Elite Tier:** Maintenance timing coordinated with Customer's technical contact where feasible, though Luna retains final scheduling authority within weekend windows

Unless Customer's Enterprise Agreement explicitly designates Elite tier status, Luna has complete discretion over maintenance scheduling within weekend non-peak hours. Planned maintenance performed with proper notice is excluded from Downtime calculations regardless of actual impact, timing, or duration.

3.5 Maintenance Notice Requirements.

- **Planned Maintenance:** Luna will provide at least seventy-two (72) hours advance notice via the status page (status.lunabase.ai) for all planned maintenance, including the expected duration and affected services. Maintenance will be scheduled during non-peak hours on Saturdays and Sundays as determined by Luna. Elite tier customers with explicit designation also receive email notification.

- **Emergency Maintenance:** Luna may perform emergency maintenance with no advance notice when required to address critical security vulnerabilities, prevent data loss, restore service availability, or prevent imminent system failure. Luna will communicate about emergency maintenance as soon as practicable.

- **Minor Updates:** Luna may deploy minor updates, patches, and configuration changes that do not impact availability without notice at any time.

3.6 Extended Maintenance Rights. Luna reserves the right to extend maintenance windows or schedule additional maintenance outside typical weekend windows when necessary for:

- Major platform upgrades or migrations (with five (5) days notice)

- Critical infrastructure changes (with five (5) days notice)

- Compliance or regulatory requirements (with reasonable notice as circumstances permit)

Extended maintenance may be scheduled outside weekend hours if technical or operational requirements make weekend maintenance impractical or insufficient. Such maintenance may exceed any informal targets mentioned in Customer's Enterprise Agreement if technically necessary, though Luna will use commercially reasonable efforts to minimize impact. Customers may not claim Service Credits for maintenance performed with proper notice, regardless of duration, timing, or impact.

3.7 Maintenance Frequency and Limits.

- **Standard/Premium Tiers:** Luna may perform unlimited maintenance during weekend non-peak hours with proper notice

- **Elite Tier:** Luna will use reasonable efforts to limit planned maintenance to four (4) hours per month, though this is a target, not a binding commitment

Luna determines non-peak hours within the Saturday-Sunday window based on real-time platform analytics, historical usage patterns, and regional considerations. The specific timing within weekends may vary based on usage patterns and may change over time. Luna may aggregate maintenance activities to minimize frequency but is not obligated to do so. The timing and grouping of maintenance activities within weekend windows are at Luna's sole discretion. Customers acknowledge that modern SaaS platforms require regular maintenance for security, performance, and reliability, and that such maintenance is essential for service delivery.

3.8 Notice Methods and Customer Responsibilities.

- Primary notice via status.lunabase.ai is deemed sufficient for all customers

- Email notifications are a courtesy for Elite tier customers only

- Customers are responsible for monitoring the status page and planning accordingly

- Failure to receive or review maintenance notices does not entitle Customer to Service Credits

- Luna is not responsible for Customer's failure to adapt to announced maintenance schedules

- Customer may not dispute Luna's determination of what constitutes "non-peak hours" within the Saturday-Sunday window or request alternative maintenance timing unless explicitly granted such rights in an Enterprise Agreement

3.9 Maintenance Summary Table.

Maintenance Type	Notice Required	SLA Credit Impact	Frequency/Timing
Planned Maintenance	72 hours	None (fully excluded)	Weekends (Sat-Sun) during Luna-determined non-peak hours
Emergency Maintenance	None	None (fully excluded)	Any time, as needed

Maintenance Type	Notice Required	SLA Credit Impact	Frequency/Timing
Minor Updates/Patches	None	None if no impact	Any time, unlimited
Major Upgrades	5 days	None (fully excluded)	Any time Luna deems necessary
Critical Infrastructure	5 days	None (fully excluded)	Any time Luna deems necessary
Compliance-Required	Reasonable	None (fully excluded)	As required by law

All maintenance performed with proper notice per this table is completely excluded from Downtime calculations and Service Credit eligibility, regardless of actual duration or impact on Customer operations. Luna's determination of specific non-peak hours within weekend windows is final, at Luna's sole and absolute discretion, and not subject to dispute.

3.10 Weekend Non-Peak Hours Determination. Luna's determination of non-peak hours within the Saturday-Sunday maintenance window is based on proprietary analytics and operational considerations including but not limited to:

- Global platform usage patterns during weekends

- Geographic distribution of weekend users across time zones

- Weekend system resource utilization metrics

- Third-party service weekend availability

- Luna personnel weekend scheduling preferences

- Historical weekend incident patterns

- Cost optimization for weekend operations

Luna's designation of specific non-peak hours within weekends is final, conclusive, and not subject to challenge or dispute by Customer. The specific hours chosen within Saturday-Sunday may vary by week, season, or region and may change without notice. Luna may perform maintenance at any time during weekends it designates as non-peak hours, even if such timing impacts Customer's specific weekend operations. Customers acknowledge that as a global platform operating across all time zones, Luna must balance competing weekend usage patterns and that maintenance timing must be optimized for the platform as a whole, not for individual customers.

3.11 AI Provider Dependency Acknowledgment. Customer explicitly acknowledges and agrees that:

- Luna's AI features are entirely dependent on third-party AI providers whose availability Luna cannot control or guarantee

- Luna has no contractual ability to compel third-party AI providers to maintain service availability

- Luna receives no advance notice of most third-party AI provider outages or maintenance

- Different AI models may have varying availability, and fallback options may not provide identical functionality

- During AI provider outages, certain Luna features may be completely unavailable with no workaround

- Luna has no obligation to maintain redundant subscriptions with all possible AI providers

- Customer's inability to use AI features due to third-party provider issues does not entitle Customer to any refunds, credits, or compensation

- Third-party AI providers may experience simultaneous outages (e.g., multiple providers down concurrently), rendering AI features entirely unavailable

- Luna's status page will indicate when issues are caused by third-party AI providers, but such transparency does not create any obligation for Luna to provide remedies

Customer assumes all risk associated with dependency on third-party AI services and waives any claims against Luna arising from third-party AI provider unavailability, performance issues, or service changes.

3.12 Third-Party AI Provider Outage Scenarios - No Credits Due:

To ensure complete clarity, the following scenarios involving third-party AI providers result in ZERO Service Credits:

Scenario	Luna Platform Status	AI Providers Status	Credits Due
OpenAI API down	✅ Operational	✖ OpenAI down	**None**
All AI providers down simultaneously	✅ Operational	✖ All providers down	**None**
Anthropic rate limiting	✅ Operational	⚠ Anthropic throttled	**None**
GPT-4 model deprecated	✅ Operational	⚠ Model unavailable	**None**
AI provider maintenance	✅ Operational	🔧 Provider maintenance	**None**
Mixed cloud + AI issues	✖ Partial outage	✖ AI also affected	**None**
AI quality degradation	✅ Operational	⚠ Poor AI outputs	**None**

Service Credits are ONLY applicable when Luna's own infrastructure (as defined in Section 2.1) fails to meet uptime commitments. The unavailability of AI features due to third-party providers never triggers SLA remedies, regardless of business impact.

4. SERVICE CREDITS

4.1 Credit Structure. If Luna fails to meet the applicable Monthly Uptime Percentage for any Service in a calendar month, Customer may be eligible for Service Credits calculated as a percentage of the monthly fees paid for the affected Service:

Monthly Uptime Percentage	Service Credit
Less than committed target but ≥ 99.0%	5%
Less than 99.0% but ≥ 98.0%	10%
Less than 98.0% but ≥ 97.0%	15%
Less than 97.0% but ≥ 95.0%	20%
Less than 95.0%	25%

4.2 Credit Conversion. Service Credits are converted to platform Credits at the following rates:

- **Standard Tier (Default):** 1:1 ratio (1 USD of Service Credit = 1 platform Credit)

- **Premium Tier:** 1:1.1 ratio (1 USD of Service Credit = 1.1 platform Credits)

- **Elite Tier:** 1:1.25 ratio (1 USD of Service Credit = 1.25 platform Credits)

Unless Customer's Enterprise Agreement explicitly specifies Premium or Elite tier status, Standard Tier conversion rates apply. Luna reserves the right to modify conversion rates with thirty (30) days' notice, except that rates applicable to already-earned Service Credits remain fixed. Platform Credits received through SLA remedies are subject to the same usage terms and restrictions as purchased Credits under Customer's subscription.

4.3 Credit Limitations. Service Credits are Customer's sole and exclusive remedy for Luna's failure to meet SLA commitments. The maximum total Service Credits for any calendar month shall not exceed twenty-five percent (25%) of Customer's monthly fees for the affected Services, regardless of the number or duration of SLA breaches. Service Credits are converted to platform Credits (the consumable currency for AI services) at Luna's standard conversion rate applicable to Customer's subscription tier and are not refundable or convertible to monetary compensation. Converted platform Credits are added to Customer's account balance and expire twelve (12) months after issuance if unused.

4.4 Credit Request Process. To receive Service Credits, Customer must submit a claim through Luna's support portal within thirty (30) days after the end of the month in which the SLA breach occurred. Claims must include:

- Specific dates and times of unavailability

- Affected Service components

- Customer's logs or monitoring data supporting the claim

- Description of impact on Customer's operations

Luna will evaluate claims within fifteen (15) business days. Approved Service Credits will be converted to platform Credits at the conversion rate specified in Customer's subscription tier and deposited into Customer's Credit balance within five (5) business days of approval. The Credit conversion rate is determined by Luna based on Customer's subscription tier and any negotiated Enterprise Agreement terms. Luna's monitoring data shall be definitive for determining Downtime unless Customer provides compelling contradictory evidence.

4.5 Shared Responsibility Credit Adjustments. In cases where Downtime is partially attributable to both Luna-controlled services and excluded third-party infrastructure:

- If Luna can demonstrate that more than 50% of the Downtime was caused by cloud provider failures, no Service Credits are due

- If any portion of the Downtime was caused by third-party AI model provider failures, no Service Credits are due regardless of percentage

- If the cause cannot be definitively determined, Service Credits may be reduced by 50%

- Luna's monitoring and root cause analysis shall be the authoritative source for attribution

- Customer may request detailed incident reports for transparency, subject to confidentiality constraints

Special provision for AI functionality: If AI features are unavailable due to third-party AI provider issues but Luna's platform remains otherwise operational, this is not considered Downtime and no Service Credits apply. The ability to route requests to AI providers is distinct from the availability of the AI providers themselves.

5. SUPPORT RESPONSE COMMITMENTS

5.1 Support Response Targets. Based on Customer's support tier as specified in their Enterprise Agreement (defaulting to Standard Tier if not explicitly stated), Luna commits to the following initial response times for support requests:

Priority Level	Standard Tier (Default)	Premium Tier	Elite Tier
Critical (P1)	Best Effort	4 hours	1 hour
High (P2)	Best Effort	8 hours	2 hours
Medium (P3)	2 business days	1 business day	4 hours
Low (P4)	5 business days	2 business days	1 business day

5.2 Priority Definitions.

- **Critical (P1):** Complete platform unavailability or data corruption affecting production use

- **High (P2):** Significant functionality degradation with no acceptable workaround

- **Medium (P3):** Moderate impact with acceptable workaround available

- **Low (P4):** Minor issues, questions, or enhancement requests

5.3 Support Limitations. Response time represents initial acknowledgment only, not resolution. Luna reserves the right to reclassify priority levels based on actual impact and available workarounds. Support coverage hours are:

- **Standard Tier (Default):** Business hours only (9 AM - 5 PM Pacific Time, Monday-Friday, excluding holidays)

- **Premium Tier:** Extended hours (6 AM - 8 PM Pacific Time, Monday-Friday, plus limited weekend coverage)

- **Elite Tier:** 24/7/365 coverage

Unless Customer's Enterprise Agreement explicitly specifies Premium or Elite support, Standard Tier hours apply. Support is provided in English only. Support scope excludes custom development, training, and third-party integration assistance unless covered by separate Professional Services agreements.

6. PERFORMANCE BENCHMARKS

6.1 Response Time Targets. While not subject to Service Credits, Luna targets the following performance benchmarks under normal operating conditions:

- API response time (p50): < 200ms

- API response time (p95): < 1000ms

- AI agent response time (simple requests): < 5 seconds

- AI agent response time (complex generation): < 30 seconds

- Project save operations: < 2 seconds

- IDE synchronization: < 3 seconds

6.2 Performance Disclaimers. Performance targets are aspirational and may vary based on request complexity, system load, third-party AI model performance, network conditions, and geographic location. Luna makes no binding commitments regarding performance metrics beyond the availability targets specified in Section 2.

7. CUSTOMER RESPONSIBILITIES

7.1 Infrastructure Requirements. Customer is responsible for maintaining adequate infrastructure to effectively utilize the Services, including:

- Stable internet connectivity with sufficient bandwidth

- Compatible browsers and development environments

- Properly configured firewalls and security settings

- Adequate local computing resources for IDE operation

7.2 Cooperation Obligations. Customer must:

- Promptly report service issues through designated channels

- Provide reasonable assistance in diagnosing problems

- Implement recommended configurations and updates

- Maintain current contact information for incident communication

- Follow Luna's Acceptable Use Policy and rate limiting guidelines

7.3 Failure to Comply. Customer's failure to meet these responsibilities may void SLA commitments and disqualify Customer from receiving Service Credits for related incidents.

8. MONITORING AND REPORTING

8.1 Service Monitoring. Luna monitors Service availability continuously using automated systems that perform synthetic transactions from geographically distributed monitoring points. Monitoring data is retained for ninety (90) days and is available to Customer upon request for SLA verification purposes.

8.2 Incident Communication. Luna will notify Customer of Critical (P1) incidents affecting service availability through:

- Email to designated technical contacts (Elite tier only, if explicitly designated in Enterprise Agreement)

- Status page updates at status.lunabase.ai (all tiers, including default Standard tier)

- In-product notifications where feasible (all tiers)

Luna's status page will distinguish between:

- **Luna application layer issues** (SLA-applicable, may result in Service Credits)

- **Cloud provider infrastructure issues** (SLA-excluded, no Service Credits)

- **Third-party AI model provider issues** (SLA-excluded, no Service Credits)

- **Multiple provider failures** (SLA-excluded even if all AI features unavailable)

This distinction helps Customer understand whether Service Credits may apply to any given incident. For third-party AI provider issues, Luna will indicate which provider(s) are affected but has no obligation to provide estimated resolution times, as these are entirely outside Luna's control. Customer acknowledges that when the status page shows third-party AI provider issues, no SLA remedies apply regardless of impact on Customer's operations.

8.3 Monthly Reports. Elite tier Customers (as explicitly designated in their Enterprise Agreement) receive monthly service level reports detailing:

- Actual availability percentages by Service

- Incident summaries with root cause analysis

- Service Credit calculations if applicable

- Performance metrics and trends

Standard (default) and Premium tier Customers may access availability data through the customer portal but do not receive proactive monthly reports unless separately negotiated.

9. REMEDIES AND LIABILITY

9.1 Exclusive Remedy. Service Credits (converted to platform Credits) constitute Customer's sole and exclusive remedy, and Luna's entire liability, for any failure to meet the service level commitments in this SLA. Customer waives any right to terminate the Agreement, claim monetary damages, or seek refunds based solely on SLA breaches, except as explicitly provided in Section 1.3. The conversion of Service Credits to platform Credits represents full satisfaction of Luna's obligations under this SLA.

9.2 Liability Cap. Luna's total liability for all SLA breaches in any twelve-month period shall not exceed the Service Credits cap of twenty-five percent (25%) of annual fees for affected Services (calculated before conversion to platform Credits). This limitation is in addition to, and does not replace, the liability limitations in the Agreement.

9.3 No Consequential Damages. In no event shall Luna be liable for any indirect, incidental, special, consequential, or punitive damages arising from SLA breaches, including but not limited to lost profits, business interruption, or data loss, even if advised of the possibility of such damages.

10. DEFINITIONS

10.1 Key Terms.

- **"Downtime"**: Period of five or more consecutive minutes during which a Service is unavailable

- **"Monthly Uptime Percentage"**: Percentage of total monthly minutes the Service was available

- **"Service Credits"**: Monetary value credits earned for SLA breaches that are converted to platform Credits (consumable AI service units) and deposited into Customer's account balance

- **"Platform Credits"**: The consumable currency units used to access AI services and computational resources on the Luna platform

- **"Luna-controlled Infrastructure"**: The application layer, services, and logic directly developed and operated by Luna, excluding underlying cloud provider infrastructure and third-party services

- **"Standard Tier"**: The default service level tier that applies to all Customers unless their Enterprise Agreement explicitly specifies Premium or Elite tier status

- **"Planned Maintenance"**: Scheduled service work performed during non-peak hours on Saturdays and Sundays as determined by Luna with at least 72 hours advance notice via status.lunabase.ai

- **"Non-Peak Hours"**: Time periods during weekends (Saturday-Sunday) determined by Luna in its sole discretion based on platform usage analytics, typically when weekend system utilization is lowest

- **"Third-Party AI Provider Issues"**: Any unavailability, performance degradation, API failures, rate limiting, quality issues, or service disruptions of external AI model providers (OpenAI, Anthropic, etc.) that are completely outside Luna's control and for which Luna bears no SLA responsibility

- **"Emergency Maintenance"**: Unscheduled maintenance required to prevent or address critical issues, performed without advance notice

- **"Force Majeure"**: Events beyond Luna's reasonable control as defined in the Agreement

11. GENERAL PROVISIONS

11.1 Entire SLA. This SLA, together with the Agreement, represents the complete service level understanding between the parties. No other service level commitments, whether oral or written, shall be binding unless incorporated into this SLA through written amendment.

11.2 Governing Law. This SLA is governed by the same law specified in the Agreement, without regard to conflict of law principles.

11.3 Survival. The limitations of liability and exclusive remedies provisions survive any termination or expiration of this SLA.

END OF SERVICE LEVEL AGREEMENT

This SLA is effective only when explicitly incorporated into a Customer's Enterprise Agreement or subscription plan. All service level commitments, response times, and credit conversion rates default to Standard Tier unless Premium or Elite tier status is explicitly documented in the Customer's Enterprise Agreement. For questions regarding SLA coverage or tier designation, contact legal@lunabase.ai.