LIFERAY WHITEPAPER

Operating Through the Mythos Era

Liferay's response to AI-accelerated vulnerability discovery, patch cadence, and continuity of service

Document	Version
Whitepaper / Trust Center	1.0
Audience	**Date**
Customers, Prospects, Risk & Procurement Teams	May 2026
Classification	**Owner**
Public	Liferay Security & Compliance

About this document

Following Anthropic's public announcement of Claude Mythos and the Project Glasswing initiative, Liferay has received inquiries from customers and prospects regarding the implications of AI-accelerated vulnerability discovery for the continuity, security, and patch cadence of Liferay's services. This whitepaper consolidates Liferay's position and serves as a reference for due-diligence questionnaires, supplier assurance reviews, and roadmap discussions. It is grounded in the controls, certifications, and programs published in the Liferay Trust Center.

1. Executive Summary

AI-assisted vulnerability discovery, exemplified by Anthropic's Claude Mythos and the Project Glasswing partnership, is expected to materially compress the time between vulnerability disclosure and exploitation, and to increase the volume of disclosures that vendors and customers must absorb. This is a meaningful change in operating conditions, but not a discontinuity in Liferay's security posture.

Liferay has operated a CVE-listed Vulnerability Disclosure Program since 2012, maintains an independently audited control framework, and certifies its AI Management System under ISO/IEC 42001. Our service-delivery and product-security programs are designed to operate continuously, to scale with disclosure volume, and to honor contractual and regulatory notification commitments — including those under DORA, the EU AI Act, and the EU Cyber Resilience Act (CRA).

This document summarizes (i) the controls and certifications that anchor Liferay's response, (ii) operational measures protecting service continuity, (iii) the expected impact on customer patching workflows, (iv) Liferay's medium-term roadmap for AI-assisted security in the product lifecycle, and (v) recommendations for customers building real-time patch operations of their own.

2. Background — The Mythos / Glasswing Context

Anthropic's Claude Mythos has been publicly described as a frontier AI model capable of autonomously discovering large numbers of zero-day vulnerabilities across operating systems, applications, and browsers, developing novel exploits, and automating significant portions of the attack chain. Project Glasswing is the partner initiative through which selected vendors and ecosystem participants receive coordinated, pre-disclosure access to findings.

The principal operational concerns for enterprise customers and their vendors are: (a) a higher volume of disclosures affecting upstream components (operating systems, browsers, libraries) that propagate into vendor products; (b) a compressed window between disclosure and weaponization; and (c) a likely shift away from calendar-based patching (e.g., monthly Patch Tuesday cycles) toward continuous, risk-based patching.

Liferay treats this as a meaningful evolution of the threat landscape rather than a categorical change. The controls described in this document already address the core requirements; the response is therefore one of scaling, tightening, and selective automation rather than wholesale redesign.

3. Liferay's Security and AI Governance Baseline

Liferay's response to AI-accelerated threats rests on a published, independently audited control framework. The following certifications and programs are maintained continuously and are documented in the Liferay Trust Center.

3.1 Independently audited certifications

Certificates and supporting documentation, including the ISO/IEC 42001 certificate, are available through the Liferay Trust Center: **https://www.liferay.com/trust-center**.

3.2 Standing security programs

Beyond certifications, Liferay operates a number of standing programs that are directly relevant to AI-accelerated vulnerability scenarios:

- Liferay DXP Vulnerability Disclosure Program, established in 2012, enabling coordinated reporting and remediation of internal and external findings.
- CVE Numbering Authority (CNA) participation, contributing to the public record of known security issues and giving customers consistent identifiers to track.
- Vulnerability and Patch Management Policy, defining the requirements for vulnerability identification, prioritization, testing, and patch deployment across Liferay systems.
- Secure Development Policy, governing secure acquisition, development, and maintenance of Liferay software, including SAST, DAST, and SCA tooling in the build pipeline.
- Incident Response Policy, defining the principles, escalation paths, and customer-communication workflow for security incidents.
- Business Continuity Plan and Disaster Recovery Policy, defining how operations continue and how IT services are restored after disruptions.
- Supplier Security Policy, governing third-party and upstream-vendor risk, including AI vendor due diligence under the AIMS.

3.3 Software Bill of Materials (SBOM)

Liferay provides Software Bills of Materials in ISO/IEC 5962 (SPDX) format, with CycloneDX available on request. SBOM coverage allows customers to map newly disclosed component vulnerabilities directly onto their Liferay deployment and prioritize accordingly. SBOM coverage will continue to be matured in line with the SBOM-related obligations under the EU Cyber Resilience Act, whose applicable deadline is December 2027.

This separation is important in the Mythos context: customers retain full control over which AI capabilities operate against their data, while Liferay's own AI use — including any AI applied to internal security operations — is independently governed and auditable.This separation is important in the Mythos context: customers retain full control over which AI capabilities operate against their data, while Liferay's own AI use — including any AI applied to internal security operations — is independently governed and auditable.

3.4 Responsible AI and the BYO-AI model

Liferay's use of AI follows a Model-as-a-Service (MaaS) approach: Liferay does not train, fine-tune, or host its own foundation models. Where AI is used — for internal productivity, software development, or operations — Liferay consumes models as a service from a vetted set of approved third-party providers offering enterprise-grade security and privacy assurances. Each provider, model, and use case is governed under Liferay's AI Management System (AIMS), certified to ISO/IEC 42001, which covers risk and impact assessment, vendor due diligence, data management, human oversight, incident response, and continuous monitoring.

Liferay does not introduce unvetted or high-risk models into its infrastructure: the set of models permitted to operate against Liferay or customer data is deliberately constrained to approved, vetted services. At the same time, Liferay applies these approved models constructively — including using them in a controlled environment to discover vulnerabilities

in its own products before adversaries do. All such use, together with any AI applied to internal security operations, is independently governed and auditable under the AIMS.

4. Operational Continuity in the Mythos-Era Threat Landscape

This section addresses the operational continuity questions most commonly raised by enterprise customers, particularly those operating under DORA or equivalent third-party risk regimes.

4.1 24/7 operational coverage and escalation

Liferay's Security Operations and Cloud Operations teams provide round-the-clock coverage with documented on-call rotations and defined escalation paths for emergency patching and incident handling. Major incidents are managed under our Incident Response Policy, with customer notification timelines aligned to contractual and regulatory commitments. The Customer Portal at support.liferay.com is the primary channel for customer-side coordination, supplemented by direct escalation paths for managed-service customers.

4.2 Emergency change procedures

Liferay operates an expedited change-management process that authorizes out-of-hours and weekend execution for security-critical changes. Where immediate patching of an upstream component is not feasible, compensating controls are deployed in the interim. These include web application firewall (WAF) rule updates, virtual patching, network segmentation adjustments, and configuration hardening, applied until the upstream fix can be ingested, tested, and rolled forward through normal release channels.

4.3 Monitoring and detection for rapid exploitation

Our detection stack ingests threat intelligence covering newly disclosed and actively exploited vulnerabilities, and our vulnerability management program prioritizes remediation based on exploitability and asset criticality rather than CVSS alone. In response to the Mythos / Glasswing context, Liferay is reinforcing monitoring for the compressed disclosure-to-exploitation window expected from AI-accelerated discovery, particularly for disclosures originating from Project Glasswing partners and upstream maintainers of components present in our SBOMs.

4.4 Service-delivery model and shared responsibility

Liferay is consumed under three principal models — SaaS, PaaS, and self-hosted — and the operational responsibility for patching is allocated accordingly:

Model	Patching responsibility	Compensating controls
Liferay SaaS	Liferay applies platform, runtime, and product patches centrally under managed change windows.	WAF, virtual patching, and segmentation managed by Liferay.
Liferay PaaS	Liferay patches the platform	Shared — Liferay platform controls plus customer-side configuration.

	; customer manages tenant configuration, runtime and custom code.	
Self-hosted	Customer applies hotfixes and updates; Liferay provides advisories and releases.	Customer-owned, informed by Liferay advisories and configuration guidance.

5. Patch Management Impact on Customers

Liferay anticipates a measurable increase in upstream-driven patch frequency as AI-accelerated discovery enters routine use. The intent is to shield customers from disclosure-volume shock without compromising the rigor of testing and release.

5.1 Release channels

- Product releases and Updates remain the primary channel for bundled, tested security and quality fixes. Bundling is preferred over per-CVE releases to keep change management tractable for customers.
- Out-of-band releases continue to be available for actively exploited or critical vulnerabilities, with customer advisories distributed through the Customer Portal.
- For Liferay-managed environments, fixes are applied transparently under agreed maintenance windows, with emergency procedures invoked where warranted.

5.2 Advisory quality

Each Liferay security advisory provides severity, exploitability context, affected versions, and required customer actions. Where applicable, advisories include CVE identifiers issued under Liferay's CNA participation, mapping cleanly to customer vulnerability-management systems.

5.3 Anticipating elevated patch volume

Liferay's scaling response to elevated upstream disclosure volumes includes: tightened triage SLAs prioritized by exploitability and exposure rather than CVSS alone; expanded threat-intelligence ingestion covering newly disclosed and actively exploited components; to shorten lead times for upstream patches. Customers receive clearer prioritization guidance and mitigation steps per advisory in order to avoid unscheduled releases.

6. AI in Liferay's Supply Chain and Product Lifecycle

Liferay's medium-term roadmap embeds AI-assisted security throughout the software development lifecycle and supply chain. All such use is governed by Liferay's ISO/IEC 42001-certified AI Management System.

6.1 AI-assisted secure development and testing

Liferay is expanding AI-assisted static analysis, dependency review, and code review in its CI/CD pipelines, complementing existing SAST, DAST, and SCA tooling. The objective is shorter time-to-detection for code-level weaknesses and tighter coupling between upstream 3rd party open source component disclosures and the components usage actually present in Liferay releases.

6.2 AI-assisted vulnerability triage

Within security operations, AI assists in triaging the higher volume of upstream disclosures expected from AI-accelerated discovery — clustering related issues, reducing false positives, and shortening time-to-fix. Human oversight is preserved for prioritization decisions, advisory authoring, and customer-facing communications, consistent with the AIMS principles of human oversight and accountability.

6.3 AI vendor and supply-chain governance

AI vendors used in Liferay's development and operations are assessed under the AIMS, with documented risk and impact assessments, vendor due diligence, and a centralized AI use-case registry. Liferay only leverages AI offerings from providers with enterprise-grade privacy and security assurances.

6.4 SBOM, provenance, and signed releases

Investment in SBOM coverage, signed releases, and software provenance continues, both to satisfy customer due-diligence requirements and to meet the obligations of the EU Cyber Resilience Act. Liferay's CRA readiness program tracks the December 2027 SBOM deadline and the broader regulation timeline.

7. Regulatory Alignment

Liferay's response to the Mythos / Glasswing landscape is aligned with the regulatory expectations our customers face — particularly in financial services, public sector, healthcare, and other regulated industries.

7.1 EU Cyber Resilience Act (CRA)

Liferay is aligning its product security and SBOM practices with the EU Cyber Resilience Act. The CRA imposes obligations on manufacturers of products with digital elements covering vulnerability handling, coordinated disclosure, security updates throughout the product support period, and the provision of SBOMs. Liferay's CRA readiness program is structured around the regulation's phased timeline, with SBOM-related obligations targeted for the December 2027 deadline.

7.2 EU AI Act

Liferay's AIMS, certified to ISO/IEC 42001, has been designed in alignment with the EU AI Act. Internal AI use cases are assessed for risk classification, transparency, and human oversight; AI vendors are subject to due diligence; and AI incidents are managed under a defined response policy.

7.3 DORA (Digital Operational Resilience Act)

For customers in scope of the Digital Operational Resilience Act using Liferay for critical or important functions, Liferay's incident response, business continuity, disaster recovery, and supplier-management programs are designed to support DORA's third-party ICT risk-management requirements. Specific contractual provisions are addressed in the applicable agreement.

7.4 NIS2 and national schemes

Liferay's operations have been independently audited against NIS2 requirements in Hungary, supporting customers whose use of Liferay falls within the scope of the EU NIS2 Directive on cybersecurity for essential and important entities. Additional certifications and attestations are maintained for specific markets, including HIPAA for US healthcare workloads and the Spanish Esquema Nacional de Seguridad (ENS) for Spanish public-sector use.

8. Recommendations for Customers

Mythos-class capabilities will shift the operating point of customer patch operations toward continuous, risk-based execution. Based on patterns Liferay observes across its customer base, the following are worth planning for:

- Move from calendar-based to risk-based continuous patching, with automation in the test and deploy stages and human oversight focused on exceptions.
- Strengthen exposure management — asset inventory, SBOM ingestion from vendors, internet-exposure mapping — so triage is driven by real exploitability rather than raw CVE volume.

- Invest in compensating-control capability — WAF, virtual patching, zero-trust network segmentation, EDR response automation — to cover the disclosure-to-patch window that Mythos-class capabilities are likely to compress.
- Tighten vendor coordination through machine-readable advisories (CSAF, VEX), aligned notification SLAs, and pre-agreed emergency change paths. CRA-driven SBOM and advisory practices across the vendor ecosystem should make this materially easier by 2027.
- Rehearse high-volume disclosure scenarios via tabletop exercises so the operational flow is exercised before it is needed.

9. How to Engage Liferay

Customers and prospects can engage Liferay's security, compliance, and account teams through the following channels:

- Liferay Trust Center — primary repository for certifications, policies, and documentation: https://www.liferay.com/trust-center.
- Customer Portal — incident reporting and managed-service coordination: https://support.liferay.com.
- Vulnerability disclosure — Liferay DXP Vulnerability Disclosure Program, in operation since 2012. Reporting guidance is published at https://www.liferay.com/trust-center/security-compliance.
- Security Vulnerabilities and Reports can be found on Customer Portal: https://support.liferay.com/security-vulnerabilities

Appendix — References

- Liferay Trust Center: https://www.liferay.com/trust-center
- Security Compliance: https://www.liferay.com/trust-center/security-compliance
- Responsible AI and ISO/IEC 42001 certificate: https://www.liferay.com/trust-center/responsible-ai
- Security Controls: https://www.liferay.com/trust-center/security-controls
- Data Protection: https://www.liferay.com/trust-center/data-protection
- FOSS / IP Compliance: https://www.liferay.com/trust-center/foss-ip-compliance
- Trust Center Documents: https://www.liferay.com/trust-center/documents