



PLATFORM RELEASE · ENGINEERING · SAAS · AI

Cho Oyu — From observation to enforcement

One pane of glass for engineering license denials, decentralized SaaS allocations, and generative AI consumption. Move from passive observation to proactive, automated endpoint and cloud enforcement.

LAC Enforcement

MCP Connector

AI Usage Dashboard

License File Management

License Parser

Mass Agent Upgrades

Software Discovery

Zendesk · Google Chat · monday.com

20+

NEW FEATURES

10

UI AREAS REFRESHED

3

NEW INTEGRATIONS

20+

BUG FIXES

servicenow.

Select
Partner

Build

www.openlm.com



— THE PROBLEM

Licensing **multiverse** zero shared visibility

⚠ THE BLIND SPOT YOU PAY FOR

Engineering concurrent licenses, SaaS subscriptions, and AI tools are managed in silos — each with its own blind spots. High-value licenses consumed from workstations without the Workstation Agent are invisible. Your consumption data is wrong exactly where licenses are most expensive. Policies observed and reported. They couldn't act.

— THE CHO OYU ANSWER



Enforce Endpoint Policy

Stop passive observation. Turn on the global toggle to automatically detect unagented machines and block non-compliant software checkouts on your next deployment cycle.



MCP Connector

Ask your license data anything in plain language via Claude, ChatGPT, Cursor, or any MCP-aware AI client. No CSV exports, no BI setup.



AI Usage Dashboard

AI subscriptions are licenses too. Track token consumption, cost by model family, dormant seats, and top spenders — across Anthropic, GitHub Copilot, and Google Gemini.

— EXPANDED COVERAGE

SaaS Agent now monitors Canva, Zoom, Claude AI allocations (with LAC support), Cursor AI, Apollo.io, LinkedIn company pages, Altair, and JetBrains Cloud. **Cloud Broker** adds GitHub Copilot usage (credits and requests), initial Google Gemini support, and monday.com API monitoring. SaaS Agents now self-update remotely. Infrastructure foundations are expanded with automated Database Configuration Tools via Docker, standalone self-signed certificate parsing, and on-demand Broker standby pausing.

200+

Groups onboarded
in one action

13

ServiceNow tables
in daily sync

10

UI areas rebuilt
with Angular Material

4+

AI providers
monitored



SPOTLIGHT · LICENSE ACCESS CONTROL

Agent enforcement: no Agent, no license.

- ✓ **Active verification:** LAC automatically cross-references allocations against the Agent Activity Manager to stop workstations from drawing licenses without an active local agent running.
- ✓ **Allocation automation:** Highlight multiple features and user groups simultaneously. The platform builds the configuration matrix instantly, dropping a manual all-day option file chore down to a single action.
- ✓ **Hybrid control:** Flip one toggle to block unauthorized checkouts globally. Policies push natively to both on-premises and SaaS license servers, preserving legitimate users during brief network disconnects (Offline ≠ Missing).

DEEP DIVE · MCP CONNECTOR

Ask your license data anything.

OpenLM now speaks the Model Context Protocol — the open standard that connects AI assistants to live business data. Sign in once with OAuth; your plain-language prompt becomes a GraphQL query against your reporting database.

Claude

ChatGPT

Gemini CLI

Windsurf

Cursor

US: cloud-us.openlm.com/mcp
EU: cloud-eu.openlm.com/mcp

Returns data your way

Tables and summaries in any MCP client. Fully interactive dashboards on higher-tier AI plans. No new BI tool. No exported CSVs. No hand-built filters.

Example prompts

"Which features were denied most often last month?"

"Show me underused AutoCAD seats by office."

"Who are my top 5 AI spenders this quarter?"

☀ Good afternoon, Victoria

Find software that hasn't been run in the last 30 days and calculate the guaranteed savings if you remove those restrictions.



Sonnet 4.6 ▾






AI FINOPS

AI usage reporting with AI Proxy




Executive KPIs

Total cost, total tokens, top spender. Boardroom-ready at a glance.




Licensed vs. Active Users

Dormant seats quantified. Reclaim AI budget from users who never logged in.




Cost by Model Family

Opus vs. Sonnet vs. Haiku — see exactly where your AI budget is going, by version.




Peak Burst

5-hour token burst identifies tier needs. Exportable day-by-day usage and cost logs.



New Homepage dashboard



The redesigned Homepage surfaces license server health, denial trends, top features, and license pool utilization in a single post-login view.



Newly introduced BI dashboards

License Utilization



Set a Quality of Service target with a slider and get a statistically grounded recommendation for the license count that meets it — capped at your actual observed peak, so it never recommends more than reality required.

Multi License Consumption

A new report flags users simultaneously holding more than one license key for the same feature and product — redundant consumption you can reclaim — with KPIs for the servers, licenses, users, and features involved.



Explore detailed BI reports at [OpenLM Analytics](https://openlm.com).

— ADDITIONAL ENHANCEMENTS & UPDATES



Secure Single Sign-On (SSO)

Log in once to access the Usage, Allocation, Denial, and Identity apps seamlessly. One global logout ends sessions across all endpoints, isolating tokens from malware.



Software Discovery for SAM

AI-assisted discovery of vendors and products, enriched with feature codes and deployment types. Tabs for vendors, products, and job history.



License File Management

Edit, validate, compare, and deploy license files in one workspace, synchronized with SLM. No more emailing license files around.



Platform Broker Standby Console

Pause or resume independent Broker activity on demand via the console. Avoid configuration dismantling during routine system maintenance windows.



Workstation Agent for macOS

The OpenLM Workstation Agent now runs natively on macOS, on both Intel and Apple Silicon. Install it interactively with the .pkg or shell installer, or roll it out across a managed fleet with your MDM. macOS joins Windows and Linux as a first-class platform for license and process monitoring. See the [Workstation Agent changelog](#) for the full per-version history.



Zendesk

Alerts auto-create tickets



Google Chat

Alerts + slash query commands



monday.com SAM

SAM inside your project boards



ServiceNow

Daily sync · 13 tables · live status



Ready to **upgrade** your fleet?

- 01 Log into your dashboard** — navigate to Platform Administration → Products → Downloads. Platform and Legacy tabs are separate.
- 02 Mass-upgrade your agents** — select agents in Agent Activity Manager, pick a target version, upgrade the entire fleet in one action.
- 03 Enable Endpoint Enforcement** — flip the global toggle in the administration console to automatically block license allocations from unagented endpoints.
- 04 Connect the MCP Connector** — point your AI client at cloud-us.openlm.com/mcp or the EU endpoint. Sign in with OAuth. Start asking questions.
- 05 Open the AI usage dashboard** — Power BI, four report pages. See your total AI cost, dormant seats, model mix, and top spenders in one view.

SUPPORTED SOLUTIONS



CERTIFICATIONS & STANDARDS



Supported deployments: Cloud · On-Premises · Docker · Azure Cosmos DB · Redis/Valkey

Request a Cho Oyu demo or read the technical release notes.

Full changelog, per-component release notes, and client setup guides at openlm.com/documentation

support@openlm.com

info@openlm.com

sales@openlm.com

