Learning Technologies Awards 2026

Category: Best Use of AI to Transform L&D in an Organization

OnPoint Digital · MetaLark.ai for T-Mobile USA

Category Entry

1. The Rationale: L&D was being asked questions it could not answer

T-Mobile's learning content ecosystem had grown the way large enterprise libraries grow — by accretion. Content arrived from different business units, agencies, acquired telcos and authoring standards over two decades. By 2026 T-Mobile's production environment held 7,003 in-scope learning objects: 5,101 courses and 1,902 content nuggets, authored across eight different tools. Their catalog listed titles and audiences but could not describe what was inside any of them, and the applied metadata was widely perceived as incomplete and inaccurate.

That was tolerable while L&D was measured on delivery. It ended when three demands arrived that the L&D team experienced challenges addressing.

Demand 1 — An enterprise skills agenda blocked by L&D data. T-Mobile had committed to a skills-based talent model and engaged a skills intelligence vendor to map earned skills to learning content. That mapping depends on content metadata being present and accurate. At this size and age, T-Mobile's library was neither: metadata for more than a third of all content considered was absent or inaccurate, delaying the skills strategy on what was fundamentally a large-scale data problem.

Demand 2 — A compliance obligation with no manual answer. T-Mobile regularly faced legal mandates to identify, isolate and document every instance of specific language deemed inappropriate in a highly-charged regulatory and cultural environment — locating every occurrence of an offending word or phrase in every file, documenting it for scheduled remediation, and reporting progress to authorities.

Demand 3 — A content investment conversation L&D could not win. Every new content request was judged on its own merits, because nobody could say what already existed covering the same ground. L&D was funding duplication it had no way to see.

The thesis of this entry follows: none of these were learning design problems. All three were consequences of L&D operating without an evidence base regarding its own assets. What follows is the acquisition of that evidence base, and what it changed about the function.

2. Why this was not achievable without AI

A content audit looks like a matter of effort rather than technology. The composition of this library is what makes that wrong: none of the standards-based courseware objects were machine-readable by any commercial or frontier LLM.

In total, 4,959 of the learning objects were packaged courses — 2,664 in SCORM v1.2, 1,926 in xAPI, 369 in SCORM v2004 — in which instructional text sits inside zipped archives of fragmented HTML, XML manifests and

JavaScript, addressable by no search tool. 291 were videos whose content existed only as spoken audio. 1,495 were PDFs, including many image-only files with no text layer at all. Eight authoring tools, led by Articulate Rise (3,109 objects) and Storyline (1,642), each packaged their output differently.

Prior to MetaLark, manual review was the only alternative, and it fails on three counts. At a conservative fifteen minutes per object — enough to launch it, skim it and record findings — 7,003 objects represent roughly 1,750 person-hours, or eleven person-months of designer and SME time diverted from their actual work. It leaves no durable artifact: the day the review ends, the library starts drifting again. And sampling cannot deliver the assurance a general counsel requires — consistency across 7,003 files is precisely what manual processes struggle with.

AI changed the unit of work from reading individual assets to querying a structured corpus. That is a difference in kind, not in degree.

3. What was Implemented

MetaLark's content extraction pipeline connected directly to T-Mobile's production LMS environment, automatically unpacking every course package, transcribing all media, and extracting text from each layer of every digital resource. For each object it generated a full-text transcript, AI-authored title and description, a summary with overview and main topics, inferred skills, and diagnostic quality flags, plus JSON and markdown outputs optimized for analysis via a secure private LLM.

The bulk run for all legacy content executed on 29 May 2026, with a delta catch-up two weeks later; MetaLark's AI content services were formalized that day.

Extraction Outcome — Content Items	Count	Rate
Learning objects evaluated (in scope)	7,003	—
Successfully extracted	6,834	97.59%
Courseware packages extracted	4,970 of 5,101	97.43%
Digital files (PDFs, videos, images) extracted	1,864 of 1,902	98.00%
Clean content extracts	6,323	92.52%
Content items flagged for human review	511	7.48%
Failed content extraction	169	2.41%

Scope note: The analysis covered active and inactive, published and unpublished courseware packages and digital files, excluding by agreement objects with fewer than six assigned users, HTML/text and URL nuggets, shared replica objects, and objects with missing assets.

Extraction Outcome — Metadata & Skills	Count	Based on
Unique metatags surfaced/discovered	31,309	6,323 objects
Unique skills mapped to content	1,901	6,323 objects
Content objects with incomplete metadata	2,421 of 7,003	34.58%
Content objects with updated descriptions	7,003 of 7,003	100.00%

That is the whole of the technical description, deliberately. What follows is what the L&D team did with it.

4. Three Shifts in How L&D Operates

Shift One — From sampling to systematic assurance, and a new standing with Legal

With 6,834 objects indexed to their full text, any compliance research question became a query rather than a project. Flagged terminology can now be located with pinpoint accuracy — to the slide, transcript segment, video frame, displayed graphic or embedded document.

All new content is now scanned and indexed automatically on upload to the T-Mobile LMS or content management platform, keeping trends, skills alignment and tagging current for discoverability and future agentic access.

The operational change is the marquee outcome. Remediation became a triaged worklist with named owners and an audit trail instead of an open-ended review, and L&D delivered Legal something it had never previously been able to offer: a complete, defensible, re-runnable analysis. A working relationship exists between the two functions because L&D turned out to be the team that could answer the question.

Shift Two — From metadata custodian to data supplier in the enterprise talent architecture

MetaLark assessed metadata across every extracted object, then filled the gaps at scale, applying AI-generated metatags and updating descriptions with generated summary overviews and main topics. Crucially, tags were applied only where no manually assigned tag already existed — human-authored metadata was preserved by design, at T-Mobile's direction.

That output feeds T-Mobile's separately selected skills intelligence vendor, which uses the mapped skills to align earned skills to completed learning. Two complementary AI systems from different vendors now compose into one pipeline: MetaLark supplying the content truth layer, the skills platform supplying the approved ontology, talent model and recommendations for skills-centric growth.

The shift for L&D was in standing. It moved from providing content lists with incomplete metadata to being the authoritative upstream source of vetted content and associated skills for T-Mobile's enterprise talent architecture.

Shift Three — From producing content to managing a portfolio

Once every object carried an improved description, accurate metatags and inferred skills, overlap became visible and actionable for the first time. L&D strategists can see clusters of duplicate and near-duplicate material, topics covered repeatedly by different business units, coverage concentrated in some domains while others go unserved, and content no longer matching any current business priority.

MetaLark reviews objects for content-level similarity using inferred skills and thematic clusters, focusing the team on duplicative or overlapping material to support consolidation and retirement decisions, and enabling evidenced "build versus reuse" choices that avoid extraneous cost.

The extraction also produced a complete picture of authoring tool sprawl, now driving a standardization decision to consolidate on Articulate Rise, Storyline and PDF. Content requests are assessed against known coverage, so L&D enters budgeting conversations with evidence rather than instinct — the difference between an order-taking function and a strategic one.

5. Evidence of Impact

For T-Mobile's L&D Team

- Audit capability: 7,003 objects assessed and 6,834 indexed in an end-to-end content review completed in 36 hours, against an estimated ~1,750 person-hours (≈11 person-months) for equivalent manual review.
- Metadata coverage: 6,834 objects updated with more than 31,000 AI-generated tags, descriptions and summaries, improving concept-level discoverability.
- Time to answer: content questions previously taking weeks now resolve in minutes or hours. L&D leadership recently requested every occurrence and variation of a phrase from a legacy HR program deemed inappropriate and culturally misaligned with new program values; the MetaLark query returned exact and semantic results in under an hour, enabling near-immediate action.
- Cost recovery: T-Mobile has substantially reduced the direct expense of content audits while improving library health, freshness and consistency.
- Capacity returned: designers and SMEs recovered measurable time from review exercises, redirected to new content design and delivery.

For T-Mobile's Learners

- Discoverability: a streamlined, concise library returns more accurate, directed results to learner queries.
- Currency: retiring outdated and duplicate content also means enterprise LLM and agentic search draw on more accurate, concise sources — learners are no longer trained on superseded or non-compliant material.
- Skills visibility: vetted skills are now tied to content completions, measured by the skills intelligence platform and surfaced through T-Mobile's LMS and in-the-flow learning agent.

Statement of Maturity

MetaLark services went live at T-Mobile in early June 2026. The extraction layer is fully enabled and in continuous use across all content uploaded to T-Mobile's LMS and document management solution. Review and analysis practices are still being documented, with scheduled reviews of flagged and failed objects and periodic audits of metadata and inferred skills. With ingestion proven, the L&D team is now investigating learner-facing concept-level retrieval through a MetaLark content player and enterprise agentic interface, using its vector database and content chunking for pinpoint delivery rather than object-level launch.

6. Responsible Practice and Human Impact

Four commitments governed the deployment, each evidenced in the delivered outcomes.

The AI defers to human judgment. AI-generated metatags were applied only where no manually assigned tag existed, by explicit agreement with T-Mobile. Nothing authored by an L&D resource was overwritten via AI inference except where specifically dictated. No object was automatically retired, altered or republished; disposition recommendations are flags on a review queue with named human owners.

Quality was measured and disclosed, not assumed. 511 objects (7.48%) were flagged for human review and root-caused rather than passed through silently — 453 as minimal extractable content, 58 as empty extracts.

295 were examined individually and the patterns documented: externally linked content, videos with no spoken audio, image-only infographics. A remediation plan with named owners and status was issued to T-Mobile. Treating AI output as provisional and evidencing its limits is the practice, not an admission.

Data stayed governed. All processing ran in T-Mobile's dedicated LMS environment under SOC 2 controls with enterprise SSO. No T-Mobile content was used to train external AI models, per strict data residency and retention terms.

The human consequence was the subtraction of drudgery, not of roles. The alternative was never "no work" — it was roughly 1,750 person-hours of manual file review absorbed by instructional designers and SMEs. No T-Mobile L&D role was reduced or eliminated as a result of using MetaLark; reclaimed capacity was reapplied to content design and development, SME and stakeholder interactions and strategic planning.

7. Why This Transition is Durable

A content audit is an event. What T-Mobile has now is a flexible instrument, and that is the distinction we ask the panel to weigh. The corpus is defined, understood, and updated continuously as content changes. Compliance questions are answered in hours rather than weeks. New proposals are assessed against known coverage. The skills intelligence platform is continuously fed. Flagged and failed objects are reprocessed as the pipeline improves. And the L&D team is positioned to exploit cleaner, more concise content across all planned AI and agentic services.

L&D at T-Mobile did not become better at making learning content. It became the function that knows what the organization knows — and that changed which conversations it is in.

Supporting Materials

All supporting materials are posted to the web address below and can be accessed with any standard web browser.

https://metalark.ai/learning-technologies-submission/

Available Materials on submission mini=site include:

1. Executive Summary document (as PDF)
2. Category Submission document (as PDF)
3. T-Mobile AI Content Extraction Results Report (11 June 2026) — extraction rates, content type and authoring tool distribution, quality root-cause analysis.
4. Screenshots of extracted content managed within MetaLark for T-Mobile
5. Link to a Recorded Demonstration (via YouTube private link)