Metadata Readiness Checklist — One Page

Software cannot watch your videos — it finds, recommends, and searches titles through metadata. Discovery is only as good as the metadata under it. Capture the four kinds, adopt the identity and interchange standards, blend your tagging sources, and build the pipeline before launch. Standards versions move in 2026 — confirm live.

1 · THE FOUR KINDS YOU MUST CAPTURE

- ☐ **Descriptive** — title, synopsis, genre, cast, mood, theme, tags. The fuel for recs and search; go past title+genre.
- ☐ **Structural** — series → season → episode, chapters, trailer ↔ feature. Powers grouping, resume, "next episode."
- ☐ **Technical** — resolution, codec, bitrate, audio/subtitle tracks, DRM. Mostly filter badges; the encoder fills it in.
- ☐ **Administrative** — owner, territory, window dates, model, rating. Gates whether a title may appear at all.

2 · IDENTITY & INTERCHANGE (standards, not custom code)

- ☐ **EIDR** — adopt the universal content ID (DOI-based; series→season→episode); insist licensed content arrives with it.
- ☐ **Canonical ID** — assign internal IDs for everything without EIDR, so "is this the same title?" is never a guess.
- ☐ **MovieLabs MEC** — accept catalogs in the Common Metadata / MEC schema so fields land in fields, not custom parsers.
- ☐ **Schema.org** — publish VideoObject / Movie / TVSeries JSON-LD so Google and AI assistants can read your titles.

3 · TAGGING & ENRICHMENT (blend three, always QC)

- ☐ **Human tagging** — nuanced tags (mood, theme); the richest fuel, but slow, costly, inconsistent without a vocabulary.
- ☐ **Providers** — buy ready metadata for the mainstream catalog; weak on the niche/original content you differentiate on.
- ☐ **AI enrichment** — auto-tag scenes, objects, speech, mood at catalog scale; QC it (one mislabel pollutes every rec).
- ☐ **Fingerprinting (ACR)** — identify unlabeled content and catch duplicates; link every asset to its canonical record.

4 · THE PIPELINE (build it before launch)

- ☐ **Ingest** — accept many supplier dialects and formats without a custom importer for every deal.
- ☐ **Normalize** — map every field to one canonical vocabulary and one ID; deduplicate.
- ☐ **Enrich** — fill the gaps with human + AI + provider data; add the descriptive depth discovery needs.
- ☐ **Govern** — own the truth: accuracy, freshness, and who fixes a wrong tag.

THE ORDER OF OPERATIONS — IDENTIFY, INTERCHANGE, ENRICH, GOVERN; DESCRIPTIVE DEPTH BEFORE LAUNCH

Metadata is the fuel discovery burns, and the discovery layer can never be better than the metadata under it — garbage in, garbage out. Do it in order. First, capture all four kinds, and refuse a catalog that arrives with only title and genre, because descriptive and structural metadata are what recommendations, search, the home-screen rows, and the cold-start path for new titles all read. Second, solve identity and interchange with standards, not custom code: give every title one universal ID (EIDR) so all your systems agree which title is which, and accept catalogs in a shared schema (MovieLabs Common Metadata / MEC) so a supplier's fields land in your fields without a per-deal parser; publish Schema.org structured data so Google and AI assistants can read your catalog too. Third, enrich by blending human tagging, third-party providers, and AI plus content fingerprinting — and quality-control the machine output, because one mislabel pollutes every downstream recommendation. Fourth, build the ingest-normalize-enrich-govern pipeline before launch, not as a backlog item — thin metadata makes a good catalog invisible and turns licence fees into write-offs. Standards versions (MEC, Google's required structured-data fields) move — re-verify before you plan against them.