Discovery-and-Retention Audit — One-Page Checklist

Fora Soft Learn

On a large catalog, viewers watch what the system surfaces — recommendations drive ~80% of hours (Netflix, 2015). Discovery is the retention engine, not decoration. Audit the three surfaces, the sixty-second window, and the one metric that keeps it honest. Benchmarks move in 2026 — confirm live.

1 · THE THREE DISCOVERY SURFACES (build all three)

- ☐ **Recommendations** — personalized rows ("because you watched"); the ~80%-of-hours engine.
- ☐ **Search** — find a known title; rescue typos/near-misses and hand them back to recommendations.
- ☐ **Merchandising** — which rows, in what order, and per-viewer artwork. The home screen is a recommendation.
- ☐ **Shared foundation** — all three run on one metadata pipeline + one viewing-event stream.

2 · THE SIXTY-SECOND WINDOW (find it or lose them)

- ☐ **The clock** — a viewer reviews 10-20 titles in 60-90s, then leaves; they do NOT browse harder.
- ☐ **It compounds** — repeated failed sessions teach "nothing's ever on" — the belief that cancels.
- ☐ **The cost** — 19% abandon a session over discovery failure; 49% would cancel (Gracenote/Nielsen 2025).
- ☐ **QoE counts too** — a perfect rec behind a spinner still loses the viewer; measure discovery + delivery together.

3 · THE CHOICE-OVERLOAD TRAP (show less, better)

- ☐ **More can convert worse** — 24 jams → 3% bought; 6 jams → 30% (Iyengar & Lepper, 2000).
- ☐ **The wall fails** — an undifferentiated grid of thousands of posters paralyses the chooser.
- ☐ **Narrow for them** — good discovery cuts thousands of titles to a hand-picked few per viewer.
- ☐ **Personalize the rows** — row order and artwork are personalization decisions, not fixed design.

4 · MEASURE WATCH TIME, NOT CLICKS (the honest target)

- ☐ **Clicks lie** — tune for clicks and discovery learns clickbait artwork and tappable-but-mediocre titles.
- ☐ **The honest metric** — watch time, completed sessions, and return rate predict retention; clicks don't.
- ☐ **Prove it** — every discovery change ships behind an experiment measured against retention, not taps.
- ☐ **Mind the privacy line** — the personalization pipeline runs on viewing data; respect the consent boundary.

THE ORDER OF OPERATIONS — TREAT DISCOVERY AS THE RETENTION ENGINE FROM DAY ONE

A streaming service spends its budget on the catalog and the subscribers; both are wasted if viewers open the app, scroll for a minute, find nothing, and leave. Do it in order. First, accept the thesis: on a large catalog, what a viewer watches is mostly what the system surfaces — Netflix reports recommendations drive about 80% of hours streamed and search the other 20%, and that personalization reduced its churn by several percentage points. Second, build all three discovery surfaces — recommendations, search, and merchandising — on a shared metadata pipeline and viewing-event stream, not as a late-stage UI bolt-on. Third, respect the sixty-second window: a viewer who finds nothing leaves rather than browses harder, and repeated failures become churn. Fourth, avoid the choice-overload trap by showing less, better — narrow thousands of titles to a hand-picked few. Fifth, and most important, measure discovery by watch time and return, never by clicks, and prove every change with an experiment tied to retention. Benchmarks (search-time and abandonment figures, 2025) move — re-verify before you plan against them.



GO DEEPER · THE 7-VOLUME HANDBOOK

Fora Soft Video Engineering Handbook

Seven volumes, 9,000+ pages: encoding, streaming, audio and AI for video products — from the pixel to production.

By Nikolay Sapunov, CEO at Fora Soft · Kindle & paperback on Amazon

Browse the handbook on Amazon →

This audit checks the three discovery surfaces, the sixty-second window and the metric that keeps them honest. Our seven-volume handbook covers the stack behind discovery; start with Volume 6 for search and retrieval over a video catalog, then Volume 3 for players and QoE, and Volume 1 for encoding.

WHAT YOU WILL FIND INSIDE

- Search and RAG over a video archive (Vol. 6)
- Players and startup time (Vol. 3)
- QoE monitoring of playback (Vol. 3)
- Encoding the catalog (Vol. 1)

THE FULL HANDBOOK · 7 VOLUMES · 9,000+ PAGES · ALL ON AMAZON

Vol. 1 · Video Encoding — pixels and color to H.264, HEVC, AV1, quality metrics and AI

Vol. 2 · Video Streaming I — network transport, contribution and ingest, HLS, MPEG-DASH, CMAF

Vol. 3 · Video Streaming II — ABR, CDNs, players, WebRTC, DRM and the economics of delivery

Vol. 4 · Audio for Video I — PCM, codecs, loudness, accessibility, real-time voice processing

Vol. 5 · Audio for Video II — jitter buffers, lip sync, Atmos, quality metrics, neural codecs

Vol. 6 · AI for Video I — models and inference, computer vision, speech, search, generative video

Vol. 7 · AI for Video II — real-time pipelines, agents, serving, cost, compliance, verticals

Viewers browse for a minute and leave without playing anything?

We will audit your rows, search and artwork, set up watch-time experiments and fix the slow starts behind the spinner. 250+ video products shipped since 2005.

Book a 30-min call → **WhatsApp →** **Email us →**