

Discovery-and-Retention Audit — One-Page Checklist

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 paralyzes 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.