Viewership Metrics: Definitions & Honest-Counting Checklist

Every viewership number is a counting rule you choose, not a fact you read off a sensor — so report the definition alongside the figure. This sheet gives each metric an honest definition, names the trap that makes the naive number lie, and shows how to count it so it survives an audit. The concurrency math at the bottom is the one that sizes your bandwidth bill.

1 · DEFINE THE EVENT — HONEST DEFINITIONS

- ☐ **Play / View** — one viewer starting a title. Count at the first frame (begins to render), not at buffer start, and past a qualifying threshold.
- ☐ **Watch time** — total elapsed playback time; may include rebuffering, seeking, startup. Not the same as content runtime consumed.
- ☐ **Playing time** — the stricter metric: only time content actually played, excluding rebuffering, seeking, and pauses.
- ☐ **Completion rate** — fraction of a title finished, read as a 25/50/75/100% quartile curve (IAB VAST events), never one average.

2 · HEADCOUNT — COUNT PEOPLE, NOT ACCOUNTS

- ☐ **Concurrency** — viewers watching at the same instant; track PEAK, not average. Sizes the delivery infrastructure.
- ☐ **Unique viewers** — distinct people/profiles, de-duplicated across devices by a stable Viewer ID, in a stated window.
- ☐ **Reach** — unique viewers an ad or title touched; for cross-media, MRC requires 100% pixels on screen for ≥ 2 s.
- ☐ **Accounts** — the base headcount only; one subscription can be several people on several devices — never quote as viewers.

3 · THE COUNTING TRAPS

- ☐ **Autoplay inflates plays** — preview tiles firing the plays counter can overstate genuine viewing many times over.
- ☐ **Watch time ≠ content consumed** — a 2-min video at 2× speed is one minute of watch time; buffering may be included.
- ☐ **Averages hide the drop-off** — '55% completion' fits both a healthy title and one that loses half the audience early.
- ☐ **Bots & invalid traffic** — filter at least GIVT (known bots, crawlers, data-center traffic) before quoting any figure.

4 · HOW TO COUNT IT HONESTLY

- ☐ **Publish the rule** — state the qualifying threshold and resume window alongside every play/view number.
- ☐ **Separate qualified plays from preview starts** — never blur autoplay previews into 'plays'.
- ☐ **Say if buffering is in** — quote watch time vs playing time explicitly; the gap is your QoE problem.
- ☐ **Read distributions, not means** — completion as a quartile curve; audience as unique viewers; size on PEAK concurrency.

THE CONCURRENCY MATH — WHAT SIZES THE BANDWIDTH BILL

Concurrency is the one viewership metric a daily-totals dashboard never shows you, and it is the number your delivery is provisioned against. Walk the arithmetic out loud. A live final peaks at 200,000 concurrent viewers and your adaptive encoding delivers an average of 5 Mbps per viewer at that moment: 200,000 × 5 Mbps = 1,000,000 Mbps = 1,000 Gbps ≈ 1 Tbps. That one-terabit-per-second figure — PEAK concurrency × per-viewer bitrate — is what you size content-delivery capacity for, not the comfortable daily average. A platform that planned around 'two million plays today' and ignored the 200,000-at-once spike fails at exactly the moment the most people are watching. To re-size for your own event, replace the two inputs: expected peak concurrent viewers, and the average delivered bitrate at peak (lower it with a smarter encoding ladder, and your peak bill falls in direct proportion). Provision for the peak, instrument a single consistent play/unique-viewer definition across web, mobile, and TV, and filter bots before any number reaches a board, an advertiser, or a content partner.