OTT Reference-Architecture Pack

The full data-plane and control-plane component list with the standard on each box, the one protection-boundary rule, and a readiness checklist to verify your own design. Engineering guidance — confirm vendor pricing live, it changes.

1 · DATA PLANE — the path the video bytes travel (own the marked boxes)

- ☐ **Ingest.** Accept the source: a mezzanine master file for VOD, or a live feed over RTMP / SRT / WHIP. Validate and normalise before anything downstream.
- ☐ **Encode → Package.** Build the ABR ladder, then package **once** as CMAF (ISO/IEC 23000-19) — one set of segments indexed by both HLS (RFC 8216) and DASH (ISO/IEC 23009-1).
- ☐ **Encrypt — plan day one.** Apply the **cbcs** scheme of Common Encryption (ISO/IEC 23001-7). cenc-only silently breaks every FairPlay device.
- ☐ **Origin → CDN → Player.** Origin is the source of truth; the CDN serves viewers and its egress sets margin; the player uses MSE + EME (W3C) and the device's native DRM.

2 · CONTROL PLANE — decides and records (first-class services, never per-player)

- ☐ **Identity + Entitlement.** One auth service and one entitlement service that answers "may this account watch this title, here, now?" on every play. Never copy the logic into each player.
- ☐ **Billing / monetization.** Subscriptions, transactions, or the ad relationship — SVOD, AVOD, or TVOD carry different requirements.
- ☐ **Ad decisioning (AVOD).** Server-side ad insertion stitches ads in; SCTE-35 (2023r1) marks the opportunity, VAST 4.3 requests the ad.
- ☐ **Metadata / recs + Analytics / QoE.** Discovery is the catalog's nervous system; analytics is how you see in. Define a "play" before you count it.

3 · THE PROTECTION BOUNDARY (the line a studio licence audits)

- ☐ Draw a box from the moment segments are encrypted to the moment a licensed player decrypts them. Inside it: cbcs-encrypted CMAF, at origin, across the CDN, in transit.
- ☐ **Encrypt once, license many.** One cbcs encryption → one license service issues Widevine, PlayReady, and FairPlay from the same files. Not three encodes.
- ☐ **Keys live separately.** A key store inside the boundary; the license server hands a key to a player only after entitlement says yes.
- ☐ **Know what DRM protects.** The decrypted bytes — not a camera at the screen (watermarking) and not the cable (HDCP output protection). Separate layers.

4 · READINESS CHECKLIST (tick before you build)

- ☐ Do we package once as CMAF and deliver both HLS and DASH from it?
- ☐ Do we encrypt with cbcs so FairPlay, Widevine, and PlayReady all work?
- ☐ Is entitlement a single service every player calls — not duplicated logic?
- ☐ Do all players emit the same QoE beacons to one analytics pipeline?
- ☐ Is delivery portable across more than one CDN, engineered for cache offload?

THE ONE RULE THAT PROTECTS YOUR CATALOG AND YOUR MARGIN

Two systems, one sentence: a data plane that carries the video and a control plane that decides and records. Package once as CMAF; encrypt once with cbcs; deliver as HLS and DASH; license Widevine, PlayReady, and FairPlay from the same segments — that single discipline keeps every device playing and your catalog protected. Then remember the cost truth: CDN egress is the recurring bill that sets your margin, so architect for a high cache-hit ratio and keep delivery portable. Make entitlement and analytics first-class services from day one; both are cheap in bytes and decisive in outcome.