noisypeak

THE APEX OF VIDEO OVER IP SOLUTIONS

SERVER-SIDE AD INSERTION · PRODUCT BROCHURE

Prism X

Dynamic SSAI that stitches VAST ads into HLS/DASH streams per viewer, on live and VOD.



Prism X stitches VAST ads into HLS/DASH streams on the server — per viewer, on live and VOD, 700 ms to first segment. Primary and secondary VAST sources, fair CPM counting, ClickHouse and Grafana reporting. Cloud or on-premise, 2 to 50 session nodes.

700 ms
TO FIRST SEGMENT

2–50
SESSION PROCESSING NODES

2–100
CONTENT TRANSCODING NODES

500 TB
AD-CONTENT CACHE, MAX

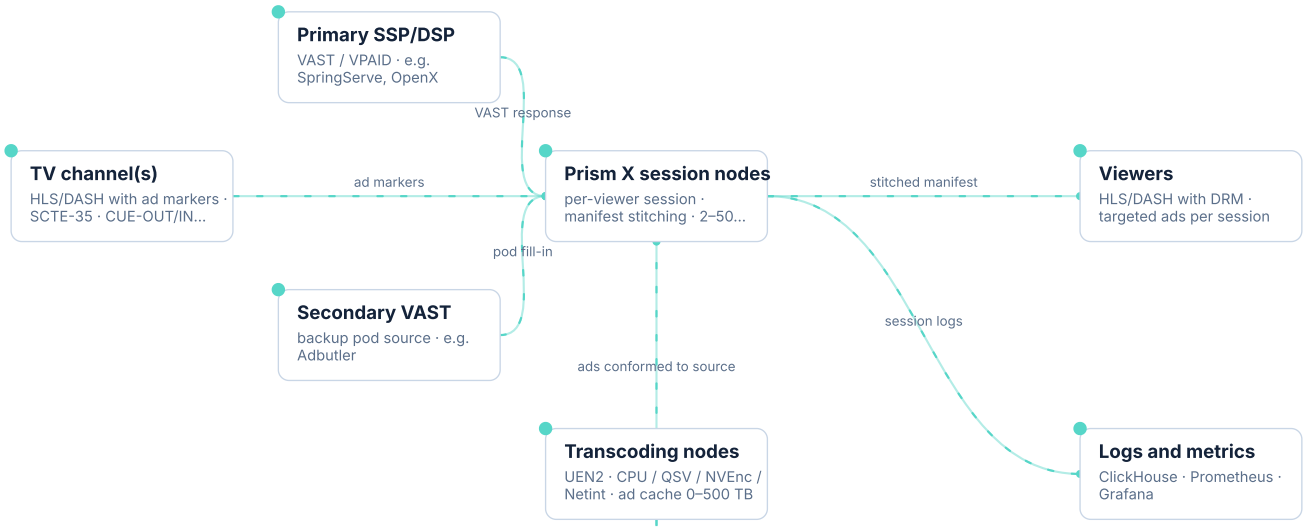
NOISYPEAK SÀRL · EPFL INNOVATION PARK · LAUSANNE · SWITZERLAND · 2026

PRISM X — SERVER-SIDE AD INSERTION

Prism X is Noisypeak's server-side ad insertion (SSAI) platform. It takes HLS/DASH channels carrying ad markers, requests ads from VAST/VPAID sources and stitches them into the stream on the server, with real-time targeting per session. It is built for broadcasters, FAST channel operators and streaming services, and runs as a service in the Noisypeak Cloud or on-premise on bare metal, VMware, Hyper-V, AWS, Google Cloud or Azure.

Because ads are inserted server-side, they are not affected by ad blockers and do not add buffering: the first segment is delivered in 700 ms. It scales from 2 to 50 session nodes and 2 to 100 transcoding nodes, and bills on a fair CPM model — only ads that reach the viewer are counted. Prism X is the monetisation stage behind the 130+ FAST channels of fastchannels.tv.

How Prism X stitches ads into a stream



Ads are requested per session, transcoded once to match the source and stitched into each viewer's manifest.

HOW STITCHING WORKS

Prism X sits between the origin and the viewer. The source channel arrives as HLS or DASH with ad markers — EXT-X-CUE-OUT/EXT-X-CUE-IN, EXT-X-SCTE35 or EXT-X-DATERANGE. Each viewer gets a session; the session node rewrites the manifest for that viewer, replacing the marked ad break with ad segments. Content segments and ad segments are served as one continuous stream, so the player sees no break and ad blockers have nothing to block.

Ad creatives are transcoded by the integrated UEN2 engine to strictly match the source signal: audio and video codecs, frame size, frame rate, bitrate, codec level and GOP structure. The conformed ad is kept in the ad-content cache (0 to 500 TB) so the same creative is not processed twice. The minimal first segment delivery latency is 700 ms.

- Markers: EXT-X-CUE-OUT/IN · EXT-X-SCTE35 · EXT-X-DATERANGE
- Per-session manifest with per-viewer ad pods
- Ads conformed to source codecs, frame size, fps, bitrate, GOP

- Transcoding on CPU, Intel QSV, NVEnc or Netint
- Output as HLS/DASH, with DRM

TARGETING AND PERSONALISATION

Targeting runs in real time per session. The ad request carries viewer demographics, location, device type and content preferences, so two viewers of the same channel receive different ad pods. Session data — user IP and country, user agent, origin URL and the rendition ladder of the master playlist — is visible in the session view for each stream.

- Demographics, location, device type, content preferences
- Each session identified and traceable in the dashboard
- Prefixes configure many channels over one SSAI config

LIVE AND VOD

On live channels the ad break is triggered by the marker in the stream and filled as it happens. On VOD, Prism X avoids the start-up delay usually needed to fill every ad pod: its ad-slot pre-filling engine handles pod requests over the duration of playback, so the asset starts immediately and later pods are requested while the viewer watches.

- Live: ad breaks driven by SCTE-35 / CUE markers
- VOD: pod requests spread over playback, no initial delay
- Same session nodes for live and on-demand

AD DECISIONING — VAST, PRIMARY AND SECONDARY

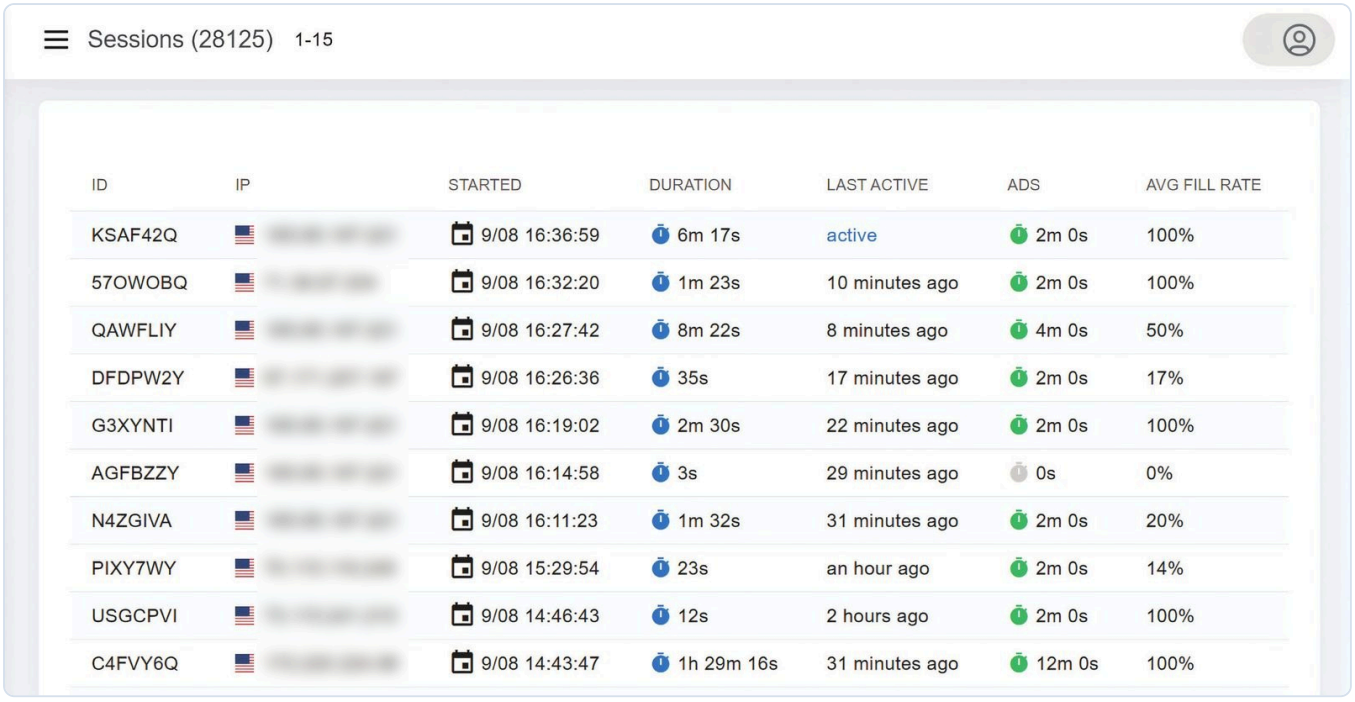
Prism X speaks VAST and VPAID upstream to the SSP/DSP of your choice; in the Noisypeak FAST pipeline this is typically SpringServe or OpenX. A secondary VAST/VPAID source fills pods the primary leaves empty — the datasheet gives Adbutler as the example backup source, and Adbutler integration is available for evaluation out of the box.

- Upstream protocols: VAST, VPAID
- Primary SSP/DSP plus secondary fill-in source
- Adbutler integration as reference SSP/DSP

MEASUREMENT AND REPORTING

Billing follows a fair CPM counting model: an ad is counted only when it is factually delivered to the viewer. Every session is logged visually and collected into ClickHouse; run-time metrics are exposed to Prometheus with Grafana dashboards. The SSAI dashboard lists all sessions with start time, duration, last activity, ad time and average fill rate, and opens a per-session timeline of content and ad pods.

- Fair CPM — delivered ads only
- Logs to ClickHouse, metrics to Prometheus/Grafana
- Per-session fill rate, pod count, viewed share



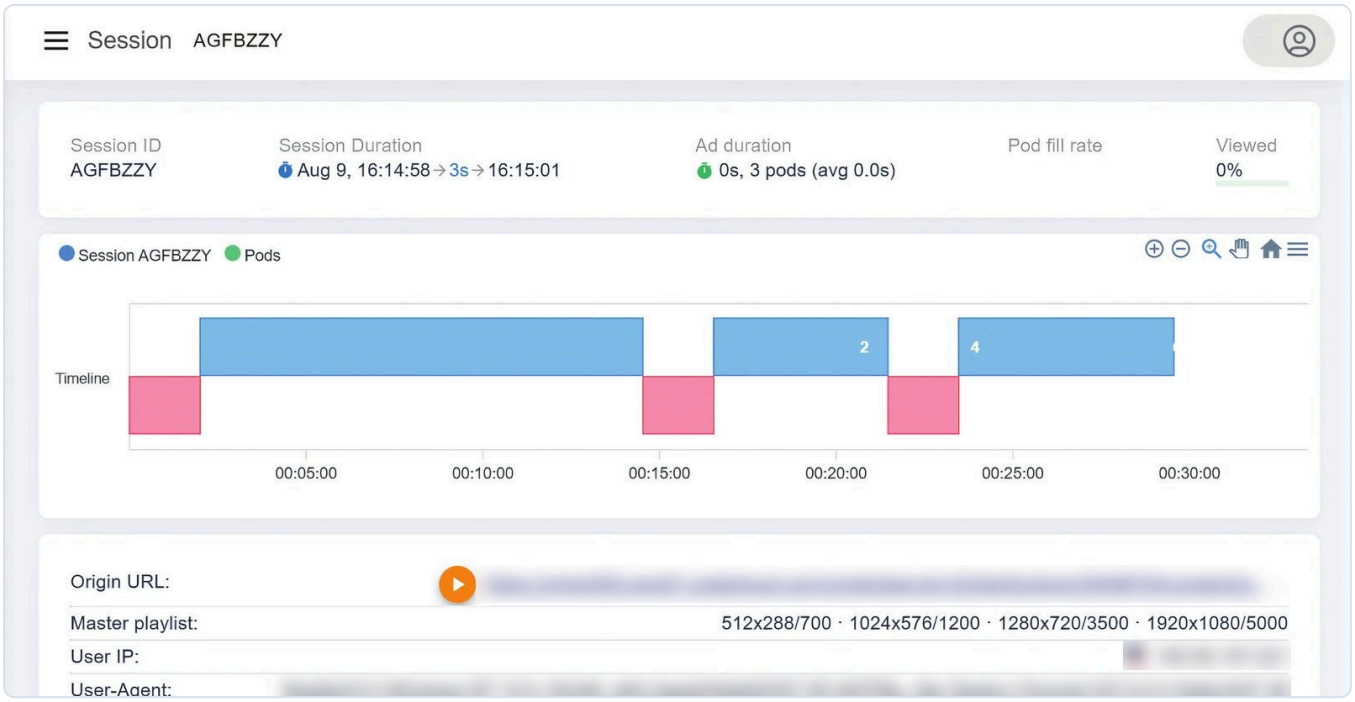
ID	IP	STARTED	DURATION	LAST ACTIVE	ADS	AVG FILL RATE
KSAF42Q		9/08 16:36:59	6m 17s	active	2m 0s	100%
57OWOBQ		9/08 16:32:20	1m 23s	10 minutes ago	2m 0s	100%
QAWFLIY		9/08 16:27:42	8m 22s	8 minutes ago	4m 0s	50%
DFDPW2Y		9/08 16:26:36	35s	17 minutes ago	2m 0s	17%
G3XYNTI		9/08 16:19:02	2m 30s	22 minutes ago	2m 0s	100%
AGFBZZY		9/08 16:14:58	3s	29 minutes ago	0s	0%
N4ZGIVA		9/08 16:11:23	1m 32s	31 minutes ago	2m 0s	20%
PIXY7WY		9/08 15:29:54	23s	an hour ago	2m 0s	14%
USGCPVI		9/08 14:46:43	12s	2 hours ago	2m 0s	100%
C4FVY6Q		9/08 14:43:47	1h 29m 16s	31 minutes ago	12m 0s	100%

Sessions list — ID, viewer IP and country, start, duration, last activity, ad time and average fill rate for each session.

SCALING AND DEPLOYMENT

Prism X is delivered as dockerised software with Ansible playbooks for setup and updates. A deployment runs 2 to 50 session processing nodes and 2 to 100 content transcoding nodes on Ubuntu or Amazon Linux, in the Noisypeak Cloud, on bare metal, VMware, Hyper-V, AWS, Google Cloud or Microsoft Azure. The cloud option gives managed scaling; on-premise gives full cost control for operators who want the lowest CPM overhead.

- Docker delivery, Ansible setup and update
- Ubuntu, Amazon Linux
- Noisypeak Cloud, bare metal, VMware, Hyper-V, AWS, GCP, Azure
- Trial from the Noisypeak Cloud with fastchannels.tv playout channels



Session view — timeline of content (blue) and ad pods (red), origin URL, rendition ladder, viewer IP and user agent.

FEATURES

Dynamic ad insertion

Real-time targeting and personalisation per session, based on viewer demographics, location, device type and content preferences.

Server-side stitching

Ads are stitched into the HLS/DASH stream on the server — no ad blocking, no extra latency or buffering on the client.

700 ms first segment

Minimal first segment delivery latency of 700 ms, on live and VOD.

Source-strict transcoding

Integrated UEN2 GPU transcoding conforms every ad to the source codecs, frame size, fps, bitrate, codec level and GOP structure.

VOD pod pre-filling

Pod requests are handled over the duration of VOD playback, removing the initial delay needed to fill all ad pods.

Primary and secondary VAST

VAST/VPAID upstream to the primary SSP/DSP; a secondary source fills the pods the primary leaves empty.

Fair CPM

Charging only for ads that are factually delivered to the end user.

Observability

Visual logging with collection to ClickHouse; Prometheus run-time metrics with Grafana dashboards.

Mass deployment

Prefixes configure many channels over one SSAI configuration at once; Docker and Ansible delivery.

SPECIFICATIONS

FORMATS

Stream formats	HLS · MPEG-DASH, in and out
Output	HLS/DASH with DRM, targeted ads per session
Source compliance	Strict — audio/video codecs, frame size, fps, bitrate, codec level, GOP structure

AD MARKERS

Supported markers	EXT-X-CUE-OUT / EXT-X-CUE-IN · EXT-X-SCTE35 · EXT-X-DATERANGE

AD STANDARDS

Upstream protocols	VAST · VPAID
Secondary source	Yes — pod fill-in from an alternative/backup VAST/VPAID source (e.g. Adbutler)
VOD	Optimised — pod requests handled over the duration of playback
Counting	Fair CPM — ads factually delivered to end users

SCALING

Session processing nodes	2 to 50
Content transcoding nodes	2 to 100
Content transcoding	CPU · Intel QSV · NVEnc · Netint
Ad-content cache	0 to 500 TB
First segment latency	700 ms minimal
Mass deployment	Prefixes to configure many channels over one SSAI config

INTEGRATIONS

SSP/DSP	Any VAST/VPAID source · SpringServe, OpenX in the FAST pipeline · Adbutler reference integration
Logging	Visual logging · ClickHouse
Metrics	Prometheus · Grafana dashboards
Transcoding engine	Noisypeak UEN2 (Noisypeak Cell)

DEPLOYMENT

Host OS	Ubuntu · Amazon Linux
Environments	Noisypeak Cloud · bare metal · VMware · Hyper-V · AWS · Google Cloud · Microsoft Azure
Delivery	Dockerised software · Ansible playbook setup and update
Evaluation	From the Noisypeak Cloud, with fastchannels.tv playout channels for trial

WHERE IT IS USED

FAST channel monetisation

Ad breaks from Noisypeak Playout filled per viewer through SpringServe/OpenX; the model behind fastchannels.tv's 130+ channels on 60+ platforms.

Live TV with targeted ads

Broadcast channels with SCTE-35 markers, re-monetised per device, location and audience on OTT.

VOD ad insertion without start-up delay

Pre-roll and mid-roll pods on on-demand libraries, filled during playback instead of before the first frame.

On-premise SSAI at lowest CPM overhead

Operators who want full cost control run the same stack on their own hardware or cloud account.

Noisypeak Sàrl

EPFL Innovation Park, Bâtiment D
CH-1015 Lausanne, Switzerland
CHE-270.718.196

Sales and evaluation

sales@noisypeak.com
+41 22 575 37 57
noisypeak.com

Works with

Noisypeak Playout · Noisypeak Cell · Noisypeak DRM

Noisypeak, Noisypeak Connected and Noisypeak DigitalEye are service marks and trademarks of Noisypeak. All other trademarks, service marks, trade names, product names and logos are the property of their respective owners. Specifications are subject to change without notice. © 2026 Noisypeak Sàrl. All rights reserved.