



TAMOSS

+



TAMS


+



lwy.fyi/phrame-gcp

Cloud-native live broadcast production on Google Cloud

TAMOSS is a Kubernetes-native implementation of the BBC TAMS v8.2 API: an operator-managed, time-addressable media store. Phrame is the live production layer: SRT/RTMP ingest, queue-driven GPU rendering, SRT programme egress and a producer web UI.



- Prometheus metrics from both stacks are collected into Managed Service for Prometheus.
- Clock sync via Google's internal NTP service — Phrame requires NTP for frame-accurate timing.
- PostgreSQL 15 and Redis run in-cluster; managed equivalents exist on Google Cloud.
- Cloud Storage interoperates with S3-style tooling through its XML API (HMAC keys).

Live production and a standards-based media store. One Kubernetes stack, run by operators.

TAMOSS and **Phrame** run side by side in one managed Kubernetes cluster: Phrame takes contribution feeds from the wire to rendered programme output, and TAMOSS keeps every frame addressable in a BBC TAMS v8.2 media store. Two operators — `tamoss-operator` and `phrame-operator` — carry the entire lifecycle.

One operator-driven lifecycle

Deploy, manage, control and extend the whole stack through Kubernetes custom resources. The operators reconcile desired state, correct drift and report readiness — day-2 operations become declarative.

Elastic GPU reasoning

Phrame's render farm is queue-driven: GPU workers scale out on dedicated GPU nodes when demand spikes and scale back when the show is over.

One storage foundation

Ingest, mezzanine and rendered media live in cloud object storage, and TAMOSS manages its store the same way — the media library and the production pipeline share one S3-compatible storage foundation.

Open standards, portable by design

TAMOSS implements the BBC TAMS v8.2 API — a published, open specification — and both products are Kubernetes-native, so the same stack runs wherever certified managed Kubernetes runs.

Producer-friendly web workflow

Producers drive the show from Phrame's web app, while operations teams watch state and browse the store through TAMOSS's operational web UI.

Built on Google Cloud

One storage tier

Ingest, mezzanine, renders and the TAMS store all live in Cloud Storage. Media never crosses a boundary between production and archive, and the XML API keeps S3 tooling working.

Frame-accurate timing

Google's in-region NTP holds every node well inside a frame period, and C3 and C4 node pools can sync to the host clock for accuracy within 1 ms.

Standard machine types

CPU-only node pools on standard machine types with no GPU quota to secure. Spot VMs suit render workers because the queue tolerates interruption.



Talk to a specialist

LiveWyer builds the platform and Phrame delivers the show: cloud native infrastructure, an open BBC TAMS media store and live production on one Kubernetes stack. Tell us about your production goals and we will show you what it looks like on your Google Cloud project.

ibc@livewyer.com · livewyer.io/contact · lwy.fyi/phrame-gcp