

Implementation brief

CULTSYNC for Avid MediaCentral

Automation-first enrichment for Avid MediaCentral with governance and in-app review.

Highlights

- Searchable, governed archives inside Avid MediaCentral
- Automation-first: review only when confidence is low
- Pilot-backed sign-off package

Delivery overview

- Phase 0 (Weeks 1–3): discovery, sizing, acceptance criteria
- Phase 1 (Weeks 4–8): infrastructure staged; integration validated
- Phase 2 (Weeks 6–12): models calibrated on sample set
- Phase 3 (Weeks 13–16): pilot processed; governance validated
- Phase 4 (Weeks 17–28): rollout, backlog processing, handover

Scope at a glance

In scope

- Replay-safe metadata writeback
- Proxy acquisition + job visibility
- Local enrichment + governance
- Review queue + audit trail

Out of scope

- Avid licensing and platform changes
- Vantage licensing
- Storage/network expansions
- Integrations beyond Avid/Vantage

1. Executive Summary

What this is

- A technical implementation brief for **CULTSYNC** integrated with Avid MediaCentral.
- Designed to improve retrieval and reuse inside Avid MediaCentral workflows.
- Self-hosted by design: media and models stay in-house.
- Acceptance is evidence-based: pilot results and integration testing.

Key messages

- Governance-first enrichment: confidence scoring, review queue, audit trail.
- Practical integration: writeback into Avid MediaCentral custom metadata.
- Operations-light by default: automation-first; human effort on exceptions.
- Predictable delivery: phased rollout with clear dependencies and acceptance criteria.

1.1 Intended Outcomes

Outcome	Result
Keep Avid MediaCentral current	Enrichment runs during ingest and backlog processing
Make archives easier to use	Searchable transcripts, people, topics, and key moments
Reduce operational load	Automation-first; review effort focused on exceptions
Governance by default	Thresholds, review queue, policy gates, audit trail
Data stays local	Media and models remain in-house

1.2 Engagement Overview

Engagement checklist

- Prerequisites validated (platform access, proxy pipeline, storage paths)
- Hardware bill-of-materials and sizing delivered during Phase 0
- Integration approach validated (metadata writeback, retries, mapping)
- Models adapted and calibrated against representative sample
- Pilot processed with operator review and governance controls
- Production rollout with training, runbooks, monitoring, and handover

Workstream	Summary
Discovery & planning	Validate prerequisites; lock acceptance criteria
Integration & deployment	Deploy services; validate Avid MediaCentral + proxy pipeline
Model adaptation	Adapt and calibrate models on representative sample set
Pilot	Process pilot batch; validate quality and governance
Rollout & handover	Scale processing; handover with runbooks and training

2. Project Context

2.1 Current State

Attribute	Value
Archive volume	120TB (estimated 10,000–15,000 hours)
Annual growth	10–15% new ingest
Languages	Arabic, French, English
MaM	Avid MediaCentral

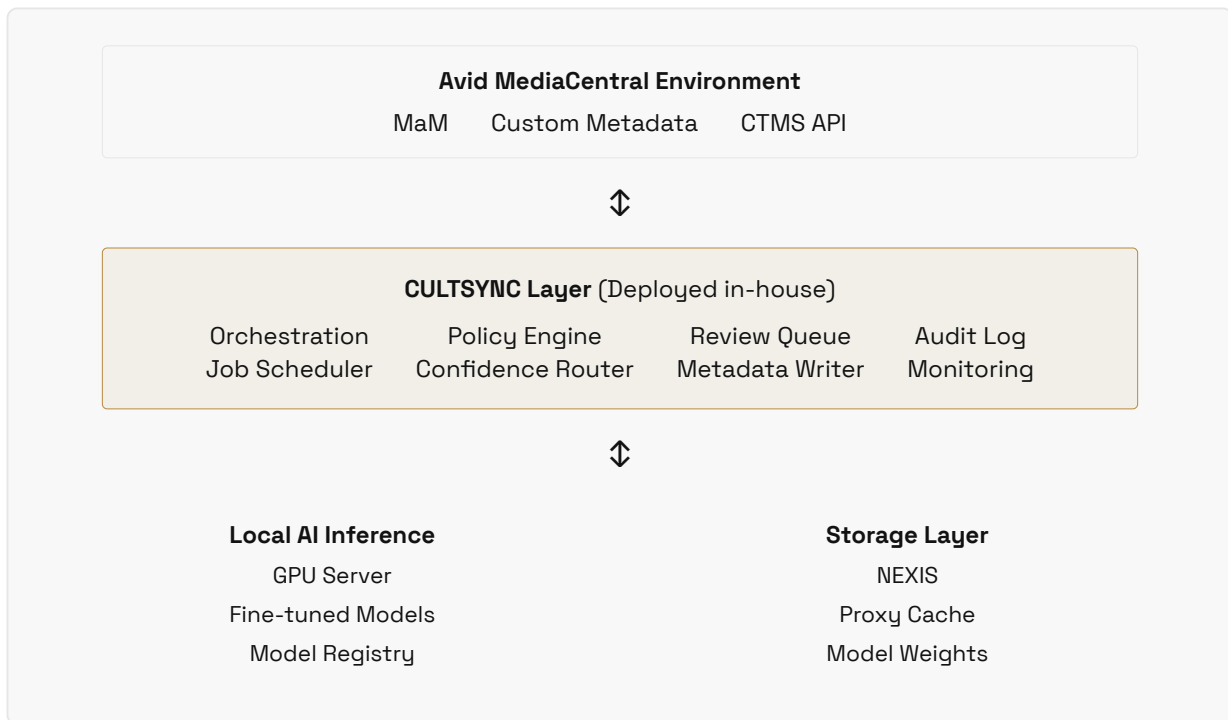
Current scope: improve find, review, and reuse of archive material without adding operator workload.

2.2 Objectives

- Align enrichment with operational requirements and governance
- Keep media and models in-house
- Enable practical search and reuse of archive content
- Keep review effort focused on exceptions (low-confidence fields)

3. Technical Architecture

3.1 System Overview



3.2 Pipeline: Decode → Inspect → Enrich → Writeback

The pipeline is designed to support predictable sizing of compute, storage, and review capacity.

Stages (per asset):

Stage	Purpose
Acquire	Select proxy for analysis; high-res when needed
Decode & inspect	Extract audio + sample frames/keyframes (optional GPU decode)
AI enrichment	ASR, OCR, and visual analysis; summaries + keywords (local)
Aggregate & score	Merge results, compute confidence, apply per-field policies
Review (if needed)	Route low/medium confidence fields for review
Writeback & interaction	Write to custom metadata; provide in-app review/curation
Audit & metrics	Record model versions, thresholds, actions, and metrics

3.3 Deployment Topology

The MaM Intelligence Layer runs as isolated containers on dedicated, self-managed infrastructure for predictable performance and controlled change.

Resources are reserved per service (CPU and accelerator assignment) to prevent contention.

Service domain	Responsibility	Compute profile	GPU
Platform	Container runtime + monitoring	CPU	N/A
CULTSYNC core	Orchestration, policy gates, audit log	CPU	N/A
Integrations	CTMS mapping/writeback; Vantage job I/O	CPU + API I/O	N/A
Media processing	Decode, audio, frame sampling, QC	CPU + I/O	Optional
Inference services	ASR, vision/OCR, summarization	GPU + CPU	GPU
Index + search	Metadata store + search index	CPU + storage	N/A
Operator review	Review queue + approvals	CPU	N/A
Observability	Metrics, logs, health checks	CPU	N/A

Resource allocation is estimated in Phase 0 and confirmed during the Phase 3 pilot to meet throughput and facility constraints.

Stage	Primary resource	Primary container	Notes
Decode & inspect	CPU	decode-inspect	Audio, frame sampling, integrity checks; optional GPU decode
Inference	GPU	asr, vision-ocr, llm	ASR, vision/OCR, and summarization (local)
Aggregate & writeback	CPU + API I/O	cultsync-core / integrations	CTMS mapping, policy, audit, retries

3.4 Local AI Stack

AI workloads run locally on dedicated GPU infrastructure. Models are versioned and pinned per job.

Capability	Implementation
Speech-to-text	Local ASR model, fine-tuned for Arabic and domain vocabulary
Speaker diarization	Local diarization for speaker segmentation and identification
Face detection	Local face detection model
Face recognition	Fine-tuned embedding model trained on provided known-persons dataset
Scene classification	Vision model fine-tuned for agreed content taxonomy
OCR	Multi-script OCR for Arabic and Latin in-frame text and graphics
Summarization	Local language model for asset descriptions and key topics

4. Dependencies

4.1 Infrastructure Prerequisites

Prerequisites checklist

- Avid MediaCentral available with CTMS licensed and reachable
- Custom metadata enabled and schema approvals available
- Vantage proxy workflow operational; proxy profile agreed
- Storage paths validated; connectivity to CTMS and Vantage endpoints
- Facility readiness (power, cooling, monitoring)

These components are required prior to project start.

Component	Requirement	Owner
Avid environment	Avid MediaCentral (2024.10 LTM or 2025.10; final version agreed with Avid) Custom metadata enabled CTMS licensed and accessible Credentials with read/write permissions	Organization
Telestream Vantage	REST API / Notify actions enabled Proxy workflow operational Proxy approach validated Access to job status endpoints	Organization
Storage & network	NEXIS (2025.10) proxy access path Connectivity to CTMS + Vantage endpoints Access to proxy storage Internet required for deploy + maintenance Minimum 10GbE to storage	Organization
Facility	Power: 1.5–2.5 kW typical / 3.0 kW peak Cooling target <85°C GPU thermals 10GbE NVMe health + lifecycle monitoring	Organization

Organization-Procured Infrastructure

The organization purchases and operates the infrastructure below. CULTSCALE provides specifications and performs deployment + integration.

Component	Specification	Owned by
GPU server (minimum)	See Appendix A (Hardware Specifications). Final sizing confirmed during Phase 0 and the Phase 3 pilot.	Organization
Base software stack	Pinned organization standard server OS Docker Engine + Docker Compose v2 (pinned) GPU driver + container GPU runtime (pinned)	Organization

4.2 Deliverables

This section lists delivered components and the documentation pack. Acceptance is via the criteria and evidence pack.

Deliverables summary checklist

- Components delivered (orchestration, enrichment, indexing, integrations)
- Avid MediaCentral operator module delivered and documented
- Runbooks delivered (deployment + operations)
- Acceptance criteria met with evidence artifacts

Deliverables (components)

Component	Notes
CULTSYNC orchestration	Containerized orchestration, job scheduling, and policy gating
AI enrichment pipeline	Speech/language/vision enrichment with model pinning + validation report
Review + governance	Confidence routing, review queue, audit trail, and operator workflows
Metadata index + provenance	Searchable index plus per-job provenance for traceability
Avid MediaCentral integration	CTMS writeback + Avid MediaCentral operator module for review/curation
Vantage integration	Proxy acquisition + job status integration for operational visibility

Documentation pack

Document	Includes
Architecture + dataflow diagrams	System overview, pipeline stages, and integration boundaries
Deployment runbook	Install, configuration, upgrades, rollback, and environment pinning
Operations runbook	Monitoring, alerts, incident playbooks, and routine maintenance
Integration guide (CTMS)	Endpoints, field mapping, retries/idempotency, and error semantics
Avid MediaCentral modules manual	Install guide + operator guide for review/curation workflows
Model registry + evaluation notes	Model versions, thresholds, calibration notes, and validation metrics

Acceptance Criteria (detailed)		
Area	Acceptance criteria	Evidence
Avid MediaCentral (CTMS)	CTMS read/write validated on pilot set Idempotent writeback + retry/backoff tested Metadata mapping agreed and implemented	Integration test report Pilot writeback logs
Avid MediaCentral modules	Custom Avid MediaCentral modules deployed Operators can manage interaction + curation workflows Changes write back cleanly with audit trail	Operator walkthrough UAT sign-off
Vantage	Proxy acquisition workflow operational end-to-end Representative sample validates codec/container compatibility Job status integration visible to operators	Vantage integration report Sample validation set
Governance + review	Review queue operational with roles/permissions Approval/correction flows validated Audit log captures model version, thresholds, reviewer action, timestamps, asset IDs	Governance checklist Audit log sample export
Indexing + search	Index populated for pilot set End users validate retrieval speed + relevance Time-aligned outputs available where applicable	Search demo End-user validation notes
Throughput + reliability	Sustained throughput meets backlog schedule Safe restarts (queues + provenance persist) Replays are safe and idempotent	Throughput benchmark Restart/recovery test log
Operational readiness	Monitoring + alerting in place Runbooks reviewed Handover session completed	Monitoring dashboard Runbooks Handover notes

Sign-off evidence pack	
Evidence artifact	Used for
Pilot results summary	Accuracy + throughput KPIs on representative dataset
Integration test report	CTMS idempotent writeback + Vantage proxy acquisition verification
Field mapping + taxonomy reference	As-built schema ownership + validation criteria
Runbooks + monitoring baseline	Operational readiness and handover verification

5. Delivery Approach & Scope of Work

Phase 0 Discovery & Planning

Weeks 1-3

Objective: Confirm prerequisites, lock scope, and finalize the implementation plan.

Activities (checklist):

- Validate CTMS connectivity (read/write) and metadata schema approach
- Confirm proxy availability and decodability on a small representative sample
- Review the existing proxy-generation pipeline (Vantage) and confirm integration points
- Confirm storage paths, throughput targets, and review workflow owners
- Produce final architecture and bill-of-materials specifications

Deliverables (checklist):

- Discovery report + final architecture specification
- Metadata field mapping for CTMS writeback
- Proxy profile validation results
- Final implementation plan (timeline, milestones, acceptance criteria)

Inputs (lightweight):

- Access to technical endpoints/credentials for validation
- A small sample set of representative assets
- A short working session to confirm review workflow and sign-off roles

Phase 1 Infrastructure & Base Integration

Weeks 4-8

Objective: Deploy infrastructure, establish connectivity, and validate base integration.

Activities (checklist):

- Hardware procurement and staging
- Hardware readiness support (remote) and go-live coordination
- Base software deployment (container runtime + GPU drivers; versions pinned)
- CTMS API integration development and testing
- Vantage integration development and testing
- Proxy retrieval pipeline implementation
- Metadata write-back pipeline implementation
- Custom field schema deployment in Avid MediaCentral
- End-to-end connectivity validation

Deliverables (checklist):

- Hardware and base environment validated for AI workloads (thermal baselines; NVMe health established)
- CTMS read/write integration validated (idempotent writeback; retry logic tested)
- Vantage integration validated (idempotent retries; API backoff tested)
- Proxy retrieval pipeline operational (queue persistence verified)
- Container resource isolation validated (CPU/GPU limits enforced per table)
- Integration test report

Dependencies (checklist):

- Hardware procured and staged per bill-of-materials
- API credentials provisioned
- Custom metadata fields approved

Phase 2 Model Fine-Tuning & AI Pipeline

Weeks 6-12

Objective: Adapt models and deploy the inference pipeline.

Notes:

- Overlaps with Phase 1; fine-tuning begins once sample data is collected.

Activities (checklist):

- ASR fine-tuning (Arabic optimization; domain vocabulary injection; accuracy validation)
- Known-persons dataset processing and threshold calibration
- Taxonomy definition and scene classification calibration
- OCR + summarization deployment
- Inference pipeline integration
- Confidence scoring and routing logic
- Operator review queue implementation
- Policy engine configuration

Deliverables (checklist):

- Fine-tuned ASR model with accuracy report
- Trained face recognition model with known-persons index
- Adapted scene classifier
- Deployed inference pipeline
- Operator review operational in Avid MediaCentral (custom UI SDK module)
- Model validation report with metrics

Dependencies (checklist):

- Annotated sample content for validation
- Complete known-persons dataset
- Category taxonomy sign-off
- Review workflow approval

Phase 3 Pilot Processing

Weeks 13-16

Objective: Validate quality, governance, and throughput on a representative pilot batch.

Activities (checklist):

- Pilot batch selection (500-1,000 hours)
- Full pipeline processing
- Operator quality review in Avid MediaCentral
- Confidence threshold tuning
- False positive/negative analysis
- Model refinement iterations (if needed)
- Performance optimization
- Throughput benchmarking
- User acceptance testing

Deliverables (checklist):

- Pilot batch processed and validated
- Quality metrics report
- Tuned confidence thresholds
- Performance benchmark report
- UAT sign-off

Dependencies (checklist):

- Review resources for pilot validation

- UAT participants identified
- Sign-off authority designated

Phase 4 Production Rollout

Weeks 17-28

Objective: Process full archive and transition to production operations.

Activities (checklist):

- Batch processing scaled to full archive
- Progress monitoring and reporting
- Exception handling and re-processing
- Real-time ingest workflow activation
- Operator training
- Documentation finalization
- Handover to operations
- Hypercare support period

Deliverables (checklist):

- Full archive processed (120TB / 10,000–15,000 hours)
- Real-time ingest pipeline operational
- Trained operators
- Operations runbook
- System documentation
- Handover complete

Dependencies (checklist):

- Ongoing review capacity for exceptions
- Operations staff identified for training
- Production sign-off authority

6. Timeline Summary

Total duration: 28 weeks (7 months)

Phase	Activities	Duration
Phase 0	Discovery & Planning	Weeks 1-3
Phase 1	Infrastructure & Base Integration	Weeks 4-8
Phase 2	Model Fine-Tuning & AI Pipeline	Weeks 6-12
Phase 3	Pilot Processing	Weeks 13-16
Phase 4	Production Rollout	Weeks 17-28

Critical path:

- Hardware procurement begins in Phase 0; staging completes before Phase 1
- Sample data collection gates fine-tuning start
- Integration validation gates pilot processing
- Pilot UAT gates production rollout

Parallel tracks:

- Phase 2 (fine-tuning) overlaps with Phase 1 (infrastructure)
- Archive batch processing runs continuously during Phase 4

7. Capabilities & Controls

7.1 AI Capabilities (core)

Capability	Description	Languages	Output
Transcription	Full audio-to-text with timestamps	Arabic, French, English	Timestamped transcript, word-level timing
Speaker Diarization	Speaker segmentation and labeling	Language-agnostic	Speaker segments with IDs
Face Detection	Face localization in video frames	N/A	Bounding boxes, timestamps
Face Recognition	Known-person identification	N/A	Person ID, confidence score
Scene Classification	Shot-level content categorization	N/A	Category tags per segment
Object Detection	In-frame object identification	N/A	Object labels, positions
OCR	In-frame text extraction	Arabic, French, English	Text content, position, timestamp
Summarization	Asset-level content description	Source language	Summary text, key topics

7.2 Confidence & Governance

- Confidence thresholds configurable per capability
- Low/medium confidence routes to operator review in Avid MediaCentral (UI SDK module)
- Approval workflow defined for sensitive fields
- Audit trail captures model version, thresholds, operator actions, timestamps, and asset IDs
- Policy engine controls auto-approval vs. manual review

8. Assumptions & Exclusions

8.1 Assumptions

1. **Avid environment operational:** Avid MediaCentral + CTMS installed, licensed, and reachable.
2. **Vantage operational:** Proxy workflows exist and produce usable proxy files.
3. **Proxy decodability:** Proxies decode cleanly and include usable audio for ASR.
4. **Network connectivity:** Sufficient bandwidth to storage and API endpoints.
5. **Sample content available:** 100–200 hours available for Phase 0 fine-tuning.
6. **Known-persons dataset:** Photos + identifiers provided for face recognition training.
7. **Review resources:** Operators available for pilot validation and exceptions.
8. **Facility ready:** Power and cooling available for installation.
9. **Single site:** Single-location deployment.

8.2 Exclusions

1. **Avid licensing:** Avid MediaCentral/CTMS licensing is out of scope.
2. **Vantage licensing:** Telestream Vantage licensing is out of scope.
3. **Storage expansion:** No additional NEXIS capacity included.
4. **Network infrastructure:** No switches, firewalls, or WAN changes included.
5. **Content migration:** No migration between storage systems.
6. **Electricity costs:** Ongoing power consumption is out of scope.
7. **Third-party integrations:** No systems beyond Avid MediaCentral + Vantage.
8. **24/7 on-call support:** Remote support only; no 24/7 on-site coverage.

9. Governance & Quality

9.1 Confidence-Based Routing

Routing band	Behavior
High confidence	Auto-approved; metadata written directly to Avid MediaCentral
Medium confidence	Queued for review; visible in Avid MediaCentral (operator module)
Low confidence	Held pending review; not written until approved

Thresholds are configurable per capability and tuned during pilot phase.

9.2 Queue Persistence & Resilience

Control	Implementation
Writeback queue	Persistent in-database queue; survives container restarts
API resilience	Exponential backoff (1s → 60s) with idempotent retries to prevent duplicate writes
Thermal monitoring	Throttle detection triggers alert and safe workload pause; resume after recovery
Container isolation	CPU/memory/I/O limits enforced; GPU assignment isolated per workload

9.3 Operator Review Experience (Avid MediaCentral)

Operator review runs inside Avid MediaCentral via a custom UI SDK module. Review stays in the system-of-record with auditable, controlled writeback.

Capability	Operator workflow
In-context review	View AI outputs alongside source media within Avid MediaCentral
Decision actions	Accept, reject, or modify suggestions; changes write back with audit trail
Queue + assignment	Queue-based review with assignment, batching, and prioritization
Interaction management	Operators can curate outputs and manage the intelligence layer interaction experience
Audit trail	Reviewer actions recorded with timestamps, thresholds, and model/provenance metadata

9.4 Quality Metrics

Metric	What it validates
Processing throughput (hours/day)	Backlog burn-down rate and capacity planning
Confidence score distributions	Threshold calibration and routing quality
Review rates by capability	Operational workload and exception volume
Correction rates (AI vs. human)	Model quality + training opportunities
Approval latency	Queue health and operator efficiency
Error rates + exception types	Stability, integration issues, and retry effectiveness

10. Success Criteria

10.1 Technical KPIs

KPI	Target	Measured via
ASR quality	WER ≤ 10%	Pilot validation set + spot checks (Arabic, French, English)
Face recognition quality	Precision ≥ 95%	Known-persons validation set at calibrated threshold
Processing throughput	Meets backlog schedule	Throughput benchmark + run history
Metadata write success	≥ 99.5%	CTMS writeback logs + retry metrics
System availability	≥ 99%	Monitoring uptime window during production operations

10.2 Business Outcomes

Outcome	Target	Measured via
Archive searchability	100% indexed	Sampling audits + search validation via Avid MediaCentral tools
Search efficiency	Material reduction	Operator time-on-task comparison (baseline vs. post-rollout)
Content reuse	Measurable uplift	Access pattern reporting + stakeholder review
Review efficiency	≤ 15% manual	Review queue rates + exception sampling

11. Project Organization

11.1 CULTSCALE Team

Project Lead:	Single point of contact; overall delivery responsibility
Architecture Lead:	Technical design; integration architecture
ML Engineer:	Model fine-tuning; inference optimization

Integration Engineer: CTMS/Vantage integration; deployment
 Support Engineer: Post-deployment support; monitoring

11.2 Required Roles

Project Sponsor: Executive ownership; sign-off authority
 Technical Lead: Avid/Vantage expertise; API access coordination
 IT/Infrastructure: Network, facility, and server access
 Content SME: Domain expertise; training data curation
 Review Operators: Pilot validation; ongoing exception review

12. Next Steps

1. **Review:** Confirm technical scope, access, and success criteria
2. **Phase 0 kickoff:** Schedule discovery sessions and sample collection
3. **Infrastructure readiness:** Procure hardware and confirm facility/network prerequisites
4. **Implementation start:** Begin Phase 1 deployment and integration

CULTSCALE

Maroun Najm

mm@cultscale.com

cultscale.com

This document is a technical implementation brief.
 Scope and sequencing are validated during Phase 0
 discovery.

Appendix A: Hardware Specifications

GPU Server (Minimum Specification)

CPU	64-core class server CPU (AVX2 + AES-NI)
RAM	256GB ECC
GPU	NVIDIA RTX 4090 (24GB) class (initial quantity estimated in Phase 0 and confirmed during the Phase 3 pilot)
Storage	2× 4TB NVMe
Network	10GbE

Note: Hardware specifications are estimated in Phase 0 against facility constraints and target throughput, and confirmed during the Phase 3 pilot.

Appendix B: Reference Materials

System	Links
CULTSYNC deck	cultsync-deck.pdf
Avid MediaCentral (CTMS)	CTMS overview CTMS API reference
Telestream Vantage	Documentation hub REST API reference (PDF)
Avid NEXIS	Learn & support Administration guide (PDF)