

The AVS AI Production Index

A quarterly assessment of what AI can and cannot do in audio/video production. No hype. Just what works, what is close, and what is still a demo.

The quarter moved from generators to production workspaces.

Video

Editing and revision became more important than one-shot generation.

Audio

Realtime voice and open-weight music both became practical categories.

Agents

AI moved into Microsoft 365, codebases, local runtimes, and governed workflows.

THE ONE-LINE FINDING

70-90%

AI can now compress the middle of many AV workflows. The final quality, taste, trust, and approval layer still belongs to people.

Q2 made AI video more useful by making it more editable.

READY

Controlled revisions

Keyframe and source-image workflows are now the best path to useful video output.

CLOSE

Longer continuity


Multi-shot work is improving, but consistency still fades as sequences grow.

NOT YET

Prompt-only final video

Long-form polished outputs still require human direction and finishing.

The strongest current ROI is still inside the edit and finishing stack.



Realtime voice and open-weight audio moved from experiments to useful tools.

Realtime voice

Streaming transcription, translation, and voice agents are now practical with human boundaries.

Open-weight music

Stable Audio 3 makes local or self-hosted generation credible for drafts and controlled uses.

Agents became a governance problem, not just a capability story.

Microsoft Scout, Work IQ, Windows Agent Framework, and Execution Containers show the enterprise pattern: context, tools, identity, logs, and policy controls.

Internal tools are now cheaper to build.

Cursor / Windsurf

Best for IDE-native development and multi-file edits.

Claude Code / Codex

Best for interactive codebase work and well-scoped implementation.

Antigravity / Jules

Best for parallel or asynchronous task workflows.

The next tools manage the revision loop.

Memory

Assets and references persist across projects.

Timelines

Generated clips move into edit structures sooner.

Shot revision

Change one shot or keyframe without starting over.

Clear ROI categories

Transcription & captions

Mature, fast, and reviewable.

Audio cleanup

High confidence for common production problems.

Controlled video revision

Useful when the input frame or footage is clear.

Internal tools

Coding agents make small software cheaper.

Production docs

Reports, summaries, and structured handoffs benefit from context.

Draft voice/music

Useful with review and clear rights records.

Still human-owned

Creative direction

AI can propose options; people decide.

Final QC

Captions, outputs, sync, rights, and trust need review.

Autonomous production

Agents need scoped permissions and approval points.

Use AI to make the first version faster. Use people to make the final version worthy.