Video content plan

One row per task. The plan is the sheet.

any2.video

$ cat columns.tsv

task	One job with a nameable finish line. Not a feature, not a screen.
product_area	Your own grouping. Used to spot areas with no coverage at all.
search_demand	Monthly searches for how people phrase this outside your product.
ticket_count	Tickets on this task last quarter. Export it from your helpdesk.
priority	1 = high on both numbers. 2 = high tickets. 3 = high search. 4 = wait.
mode	live | rendered | designed. Decide per task, not per programme.
languages	Decide before recording. One run serves every language.
platform_exceptions	Where the path genuinely differs. Usually 'none'.
version_tag	The build it was recorded against. This is what triggers a re-run.
owner	The team, not a person. Support owns support answers.
status	planned | recorded | published | stale
video_url	Fill on publish.
last_verified	Date the result was last checked against the product's evidence.

$ cat method.md

1	**Inventory every task**	Docs table of contents + support tag cloud + product analytics click paths.
2	**Score two columns**	search_demand and ticket_count. Both come from data you already own.
3	**Set priority**	High on both first. There are fewer of these than anyone expects.
4	**Choose mode per task**	Reachable on a screen is live. Hardware is rendered. Concepts are designed.
5	**Fix languages now**	Before recording. One verified run becomes every language you sell in.
6	**Bind to a version**	On release, re-run only the tasks the changelog touched.

The one rule

The first ten videos should be the ten answers your team is tired of giving.

Everything else on the sheet is ranked, visible and waiting. That is the difference from a backlog.