

AI Content Release Gate

Eight synthetic slides turned from fluent-looking risk into an auditable release decision.

8	8	7	0
SLIDES REVIEWED	ISSUES LOGGED	CLAIMS REGISTERED	CLIENT RECORDS USED

Executive result

The observed carousel is held before publication. An unsupported absolute claim, a tenfold cost error, inaccurate privacy copy, a missing workflow fallback and an unsafe CTA are made visible and repaired.

What this sample proves

AI-assisted content can be reviewed as an operational object: every important claim, calculation, design-density problem and release risk can be tied to evidence, a severity, a repair and an acceptance condition.

Use boundary

This is self-directed technical work using a fictional product and synthetic source pack. It is not client work, a reproduced marketplace brief, legal advice or a claim that a named product was tested.

State transformation

01 Chaos

An eight-slide AI-education carousel looks polished, but mixes fact, hype, arithmetic, privacy statements, diagrams and visual density without a release gate.

02 Operation

Slide decomposition → claim register → source request → technical and arithmetic QA → logic-path review → information-density review → risk-ranked repairs.

03 Usable state

A reviewer receives release-safe copy, a visible HOLD / REVIEW / READY state, a change log and an editable workbook that preserves every unresolved product fact.

Uncompressed residual

Product-owner intent, current product labels, legal/privacy interpretation and final aesthetic approval remain human decisions.

Observable proof

- Editable six-sheet QA workbook
- Observed and repaired text side by side
- Claim, source and change registers
- Explicit release gate and residual uncertainty

Portfolio integrity

Independently constructed self-directed technical work using public or synthetic inputs. Not a reproduced marketplace brief, confidential client engagement, paid result or claim of deployment.