

F07 / AI OPPORTUNITY PRIORITISATION

Choose the next AI experiment

Start with the work and its alternatives. Resolve the conditions for a bounded test, then compare the evidence.

SESSION

120 minutes + preparation

TEAM

Operations + tech + risk

OUTPUT

One next intervention

01 COMPARE CREDIBLE ALTERNATIVES

Process change

Existing software

AI assistance

No change

02 RESOLVE ALL FIVE GATES

01

Purpose

Whose need?

02

Data

Permitted use?

03

Consequence

Contained harm?

04

Measurement

Credible test?

05

Ownership

Who can stop it?

FAIL

Stop this scope

UNRESOLVED

Prepare evidence

ALL PASS

Compare candidates

03 / SELECT THE NEXT INTERVENTION

Compare evidence. Keep uncertainty visible.

Improve process / Prepare / Experiment / Defer / Stop

Passing gates permits a scoped experiment, not production release.

CANDIDATE AND GATES

Define the opportunity

Describe a work episode and the decision it affects. Complete this page for each candidate.

PROBLEM / USER / CURRENT ALTERNATIVE

SYSTEM BOUNDARY / AFFECTED PEOPLE / CONSEQUENCE OF ERROR

01 RECORD EVIDENCE FOR EVERY GATE

GATE / QUESTION	STATE	EVIDENCE / GAP OWNER
Purpose Is there an observed operational need?		
Data Can this data be used as proposed?		
Consequence Can test failures be contained?		
Measurement Is a baseline and evaluation available?		
Ownership Who can operate and stop the test?		

Fail: stop the proposed scope. Unresolved: complete preparation before the dependent test. Pass: evidence supports this gate for the defined experiment only.

PORTFOLIO AND NEXT ACTION

Make the next choice explicit

Retain process change, conventional software and no change as credible alternatives. Do not add unlike ratings into one score.

Candidate / alternative	Gate status / evidence	Benefit / burden / unknowns	Owner / dependency / next action

Use low / medium / high / unknown only with written local anchors. Higher operating burden is a disadvantage. Unknown does not mean medium.

CHOSEN INTERVENTION / EXPERIMENT SCOPE / EVIDENCE THAT WOULD CHANGE IT

ACCOUNTABLE OWNER

COST CEILING / CAPACITY

REVIEW DATE

STOPPING EVENT / FALLBACK / NEXT DECISION

ILLUSTRATIVE EXAMPLE

Three ideas three next steps

A hypothetical manufacturer compares three opportunities. The decisions below depend on these assumed facts.

01 Policy search

PREPARE

Documents exist, but no one owns obsolete guidance.

Assign document ownership and remove outdated material.

02 Quote extraction

EXPERIMENT

A procurement owner, permitted samples and human approval exist.

Prepare a bounded offline comparison with simpler alternatives.

03 Machine adjustment

STOP THIS SCOPE

No tested containment or domain safety case is available.

Require specialist engineering analysis before reconsideration.

The quote project can proceed only once all five gates are cleared for the scoped experiment. If review effort is excessive, return to a simpler process intervention.

This is a contingent decision, not a universal ranking of AI use cases. A different evidence profile may produce a different sequence.

FACILITATION AND RESEARCH

Run an evidence-led session

Keep observation, interpretation and unknowns distinct. Use the linked research companion for full source notes and limitations.

10 MIN	Describe the episode Name the user, decision and consequence of error.
20 MIN	Generate alternatives Include process, software, AI and no change.
25 MIN	Resolve the gates Record scope-specific evidence and unknowns.
25 MIN	Compare evidence Retain separate benefit, feasibility and burden views.
15 MIN	Check dependencies Account for scarce reviewers and operating capacity.
15 MIN	Choose the intervention Explain experiment, preparation, defer or stop.
10 MIN	Write the test brief Define scope, cost ceiling, owner and stopping event.

SELECTED SOURCES / ACCESSED 16 SEP 2026

NIST / AI Risk Management Framework 1.0 / January 2023

NIST / Generative AI Profile / July 2024

UK GDS/CDDO / Technology Code of Practice

Original prioritisation protocol informed by risk-management practice. Gates and ratings are proposed facilitation rules, not validated risk scores. No certification, endorsement or legal clearance is implied. RMF 1.0 is the cited baseline.

Complete source ledger: [companion.html](#). Original dotSuper artwork; no third-party diagrams reproduced.