OptiField

MaestroOS™

Deciding what to fix first

Integrity, maintenance and turnaround decisions under uncertainty

Two-day workshop · about 11 contact hours · up to 20 participants in teams of four

Classroom or virtual · no software and no network access required

This is a decision workshop, not an inspection or reliability course. It assumes the participants already run risk-based inspection, maintenance planning and turnaround processes, and works on the decisions those processes hand them: what to fund, what to defer, what to inspect first, and what to say when the money runs out before the scope does.

About this workshop

Every integrity and maintenance organisation produces more justified work than it can fund. The scope list is longer than the window, the window is shorter than the wish list, and each item arrives with a number attached that somebody derived, somebody else modelled, and nobody has questioned since. The decision is not whether the work is worth doing — most of it is — but which of it to do, in what order, and what it costs to wait.

Two days on exactly that. Participants work a real scope decision on a public-record asset: they frame it, class every number by where it came from, build and roll back a decision tree by hand, price what an inspection would be worth before commissioning it, put a cost per week on deferral, and sequence a programme against a single constrained window. The second afternoon is a capstone: a fixed budget, a fixed window, a deliberately ambiguous directive from management, and a late inspection finding that arrives after the plan is made.

The method is worked on paper throughout, because it has to survive back at the desk without a licence, an installation or an IT request. One short demonstration shows the same analysis run in software, so participants can see what is worth automating later.

Who should attend

- **Integrity and inspection engineers** who decide what gets looked at, and when.
- **Maintenance and reliability engineers** holding a backlog that will not fit the budget.
- **Turnaround and shutdown planners** who must sequence competing scopes through one window.
- **Asset and operations managers** who approve the programme and defend it upward.
- **Planners, cost engineers and economists** who fund the work and are asked whether it was worth it.

No mathematics beyond arithmetic, no coding, and no prior decision-analysis training is assumed. Mixed teams work better than uniform ones.

What participants will be able to do

1. State an integrity or maintenance decision at the level of the objective it serves, and identify genuinely distinct alternatives — including monitor, defer, and do nothing.
2. Classify every number behind a scope item by provenance — measured, interpreted, modelled, assumed or inherited — and separate uncertainty that data can reduce from uncertainty that must be managed.
3. Build and roll back a decision tree, and read an expected value without mistaking it for a forecast.

4. Say when an organisation's appetite for a bad outcome, rather than the numbers, is what decides the answer.
5. Put a value on an inspection or a monitoring campaign before commissioning it, and recognise data that cannot change the decision however interesting it is.
6. Price the cost of delay per week, and use it with duration to sequence approved work through a constrained window.
7. Challenge a recommendation — human or machine-generated — against its evidence, and state in advance what would change it.
8. Write a decision record their successor can audit two years later.

Programme · day one

Session	Focus
The decision, not the defect	Framing a scope decision at the level of what it is for. Alternatives that are really different, including deferral with mitigation. Why an alternative set of one is a request for approval rather than a decision
Where your numbers came from	Provenance classes applied to integrity data: inspection coverage, corrosion rates measured against modelled, assumptions inherited from a campaign nobody on the team was part of. Which of them data could still reduce
How well do you know what you know	A short calibration exercise, scored live. Most rooms discover their 90% confidence intervals are nothing of the kind
Consequence, and what it drags with it	One failure is rarely one cost. Working through the knock-on chain — production, safety margin, downstream commitments, the next window — and where the chain stops being knowable
Decision trees, by hand	Build it, roll it back, read it. Paper and argument, deliberately, so the mechanics are understood before anything computes them
When the average is the wrong answer	Expected value against certain equivalent; the failure you can only afford once. Where two defensible teams part company

Programme · day two

Session	Focus
What is an inspection worth	Value of information applied to inspection and monitoring: the most that resolving an unknown could be worth, what resolving one specific unknown is worth, and the test that says a campaign cannot change the decision and is therefore worth nothing
What it costs to wait	Cost of delay in money per week — deferred production, carrying cost, a window that closes, an obligation with a date on it. Step costs against linear ones, which in this industry is usually the step
What to do first	Sequencing competing scopes against one constrained resource by cost of delay divided by duration, and what that ordering saves against the two orderings people reach for instinctively
Reading a recommendation	A confident, well-written recommendation with a weak link in it. Teams find what is load-bearing, what the conclusion rests on, and what evidence would settle it
Capstone · the scope that does not fit	Fixed budget, fixed window, an ambiguous directive from management, a costed option list — and a late finding that lands after the plan is made. Teams allocate, re-plan, and defend the result out loud
Debrief and close	Where the teams read the directive differently, what every team funded regardless — the robust core — and each participant's own decision record to take home

THE CAPSTONE

Teams receive a budget that will not cover the list, a directive open to several honest readings, and an option list where every figure carries its provenance on its face — including one large number that is frankly assumed.

Partway through, a finding arrives that changes what is possible in the window. Teams revise under time pressure and present: how they read the directive, what they funded, what they changed, what they did not change and why, and what would change their mind.

The reasoning is the deliverable, not the portfolio. Nobody is scored.

What participants take away

- **A decision record for one live decision of their own.** Built up across the two days on their own paper. It is not collected, circulated or seen by the trainer.
- **Five worksheets they can reuse.** Framing, the tree, value of information, cost of delay, and the decision record.
- **A one-page challenge checklist.** The questions to ask of any recommendation before capital is committed.

Delivery and what the venue needs

- **Format.** Two consecutive days, classroom or virtual instructor-led. In-house delivery suits a single organisation; a mixed-organisation room produces a better argument.
- **Confidentiality.** All shared and assessed work runs on public-record data. No participant is asked to disclose their own organisation's information to the room or to the trainer; own-asset work stays on their own paper.
- **The room.** Tables that can be arranged as islands for teams of four, a screen, and printing. Theatre seating does not work.
- **Technology.** None required. No software is installed, no network access is needed, and no participant needs a laptop.

The trainer

Oddvar Spjeld is a mechanical engineer with 25 years across capital projects, industrial reliability and maintenance. He was Shut-down Manager for the Åsgard, Snorre, Visund and Gullfaks turnarounds, and owned the modification and maintenance portfolios for Troll B and Gullfaks C — annual programmes of roughly NOK 100–200 million, where the standing question was which work to fund, and in what order.

At Equinor he led predictive and condition-based maintenance implementation across offshore and onshore assets under the Technical Efficiency Programme, converting sensor data and maintenance history into operational decisions, and served as acting Department Leader for Reliability Analysis.

He holds a BSc in Mechanical Engineering from Louisiana State University and a Master of Management in business strategy and project management from BI Norwegian Business School, and completed NTNU's programme in Reliability and Maintenance Management. Earlier in his career he designed and installed active vibration-isolation systems at Caltech's LIGO gravitational-wave observatory. He is co-founder and CEO of OptiField.

A second facilitator supports the capstone, running the team work and playing the management role.