



DISCUSSION PAPER | MAINTENANCE STRATEGY

Three Lines in the Sand

Using maintenance regimes to balance cost, risk and performance

Peter McCarthy

Managing Director, DC₂O Pty Ltd

October 2026

A single maintenance budget tells you what was spent. Three lines in the sand tell you what you are choosing.

Contents

Executive summary

1. One budget line, no choices
2. What the standards expect
3. Three regimes, every asset
4. What an example building shows
5. Below the line
6. From envelope to strategy
7. Putting the envelope to work
8. Conclusions and recommendations

References

Appendix A: Maintenance level descriptions

Executive summary

Most maintenance budgets are a single number, usually last year's figure plus indexation. That number says nothing about the level of service it buys or the risk it leaves behind.

DC₂O models every asset in a building or portfolio under three maintenance regimes, High, Medium and Low, to show the full range of operating (OPEX) and capital (CAPEX) cost the building could reasonably demand. For an example building, the 25-year cumulative cost runs from about **\$2.2 million** under Low to **\$3.9 million** under High: an envelope of around **\$1.7 million**.

Three points matter

- **No building should run on one regime.** The envelope frames the choice. The strategy comes from setting regimes by organisational criticality, asset by asset or space by space.
- **The Low line is a floor.** Every regime includes compliance and statutory maintenance, so a budget below the Low line suggests regulatory obligations are not being funded.
- **The question changes.** The budget conversation moves from "what did we spend last year?" to "what level of service and risk are we prepared to fund?"

1. One budget line, no choices

Ask a facilities manager how the maintenance budget was set and the honest answer is usually some version of "last year's figure, plus a bit". Sometimes it is a percentage of replacement value, borrowed from a benchmark that was never built for this portfolio. Either way, the result is a single number, and a single number hides a choice.

The problem is not new. A quarter of a century ago, *Strategic Asset Management* translated the typical maintenance bid as "just give me a bucket of money and I will do

my best", adding that management "knows this, and is unimpressed" (Burns 2000). The same issue made a sharper point that still stands: if you cannot say what will suffer when your budget is reduced, you are at risk.

Two things have changed since then. Standards and government policy now expect maintenance decisions to be risk-based and tied to service outcomes. And the data and tools to model maintenance at component level, across the whole life of a building, have become practical rather than heroic.

This paper sets out how DC₂O uses three maintenance regimes to give asset owners a range rather than a number, and how that range becomes a maintenance strategy they can defend.

2. What the standards expect

The guidance is remarkably consistent: maintenance strategy should follow criticality, not habit.

PAS 55-2 offers an example functional policy designed to stop a manager "simply employing the manufacturer's recommendations for maintenance activities":

"Maintenance strategies will be proactive and risk-based, taking due account of asset criticality and condition." It goes on to say that the control of maintenance should reflect the criticality and urgency of tasks, work-bundling opportunities and resource constraints (BSI 2008).

ISO 55002 asks organisations to define the asset information that connects service levels, future maintenance requirements and life cycle costing analysis (ISO 2014). The Institute of Asset Management frames the optimum as the point where the combined cost of preventive action and the risk of failure is lowest, and argues that the rigour of the decision method should be proportionate to the criticality of the decision. It also observes that cost and risk trade-off decisions "are very difficult to make without at least some mathematical rigour and 'what if?' decision-support tools" (IAM 2015).

For NSW public sector agencies, TPP19-07 requires asset management plans to include a prioritised program of lifecycle activities "consistent with the accepted level of risk", together with their capital and recurrent financial implications, and expects decisions that balance cost, risk and asset performance. Its illustrative asset management objectives even include a target ratio of programmed to corrective maintenance: 70:30 (NSW Treasury 2019).

The gap between policy and practice is also well documented. Reviewing New Zealand councils' long-term plans, auditors found that changes to capital expenditure budgets often did not result in changes to forecast levels of service, "suggesting weak relationships between financial and service performance forecasts" (Strategic Asset Management 2008a). When a budget can move without the service level moving, nobody is really making a choice.

3. Three regimes, every asset

3.1 How the regimes differ

In the Farsight model, each asset carries three maintenance regimes. They are not three price points for the same work. They are three different ways of looking after the asset, each with its own purpose, basis, task frequency, approach to condition monitoring and replacement policy. Table 1 summarises the differences that drive cost; Appendix A gives the full descriptions.

Table 1: Key differences between the three maintenance regimes

	High	Medium	Low
Purpose	Business continuity and safety	Acceptable service, controlled risk	Minimise cost on non-critical assets
PM coverage (vs OEM)	~90–110%	~60–80%	~20–40% (statutory and essential)
Corrective : preventive	20:80 or better	~50:50	70:30 or higher
Backlog tolerance	0–5%	10–20%	25%+
Life multiplier	0.9–1.1× nominal life	1.0–1.3×	1.3–1.8× or on failure
Budget profile	Higher steady OPEX, fewer CAPEX surprises	Moderate OPEX, occasional unplanned cost	Low steady OPEX, higher volatility

High is for assets whose failure would materially affect operations, safety or reputation. It follows OEM schedules supplemented by condition monitoring and risk assessment, keeps critical spares on hand and renews assets before performance degrades. **Medium** maintains acceptable service with controlled risk, stretching intervals and bundling tasks where the data supports it. **Low** minimises cost for non-critical assets, accepting more variability, and run-to-failure where the consequences are tolerable.

3.2 One asset type, three schedules

The regime changes the work, not just the budget. A single asset type can carry different task frequencies and intervals between activities under each regime, each set from best practice and basic preventive maintenance standards. An air handling unit serving a laboratory and an identical unit serving a storeroom need not share a schedule: under High it may carry condition checks between OEM services, while under Low it may receive only essential and statutory tasks at extended intervals.

Because each regime also carries a life multiplier, it moves capital as well as operating cost. Under High, an asset is replaced at around 0.9 to 1.1 times its nominal life; under Low, that stretches to 1.3 to 1.8 times, or replacement on failure. The regime therefore shifts both the size of the maintenance bill and the timing of renewals.

3.3 The compliance floor is built in

Every regime, including Low, contains the compliance and statutory maintenance the asset requires. Essential safety measures such as fire detection and suppression, emergency lighting, exit signage and lifts must be maintained to the relevant standards whatever regime is chosen; for fire protection systems in Australia, that is AS 1851 (Standards Australia 2012). Low is therefore not "do nothing". It is the minimum an organisation must do to stay safe and legal, plus whatever reactive work the assets demand.

4. What an example building shows

Figure 1 shows the cumulative maintenance cost for an example building when every asset is placed on the same regime, from 2026 to 2050.




Figure 1: Cumulative maintenance cost for an example building under each regime applied across all assets, 2026–2050. Source: *Farsight lifecycle cost model (illustrative)*.

Over the 25 years, the cumulative cost comes to about \$3.9 million under High, \$2.8 million under Medium and \$2.2 million under Low. The envelope between the High and Low lines is about \$1.7 million: on average, close to \$70,000 a year for this one building. Three features of the chart deserve a closer look.

The lines diverge early and keep diverging. In the early years, routine preventive work drives the difference. High carries more frequent tasks and condition monitoring, and the gap widens steadily.

The steps move to the right. Each line has sharp steps where major renewals fall. A step that appears under High in the mid-2030s arrives a few years later under Medium, and later again under Low. That is the life multiplier at work: the same renewal, deferred as the regime relaxes.

The Low line steepens at the end. Deferral is not avoidance. Towards the end of the horizon, the Low line climbs faster as deferred renewals start to fall due. A 25-year window flatters Low, because some of the renewals it defers fall beyond 2050. Over the 30- to 50-year horizons we recommend for whole-of-life analysis, the gap would be expected to narrow.

What the chart does not show matters just as much. These are direct maintenance and renewal costs. The cost of failure (downtime, lost teaching or research time, emergency call-outs, reputational damage) sits outside the lines, and it rises as the regime falls. The IAM optimum is the lowest combined cost of preventive action and failure risk, not the lowest maintenance spend (IAM 2015). The Low line looks cheapest partly because it leaves the risk off the page.

5. Below the line

If the Low line is the floor, what does it mean to sit under it?

Because Low includes all compliance and statutory work, an annual budget below the Low profile indicates that the organisation is not providing enough funding to meet its regulatory obligations. That is not a saving. It is a liability carried off the books, and it tends to surface at the worst possible moment: a failed inspection, an insurer's question, an incident.

"As litigation becomes ever more prevalent, maintenance managers will be expected by the courts to identify maintenance needs and argue competently for the resources needed to meet these needs."

Strategic Asset Management, Issue 40 (Burns 2000)

Twenty-five years on, the Low line gives maintenance managers the evidence to do exactly that. For boards and accountable officers, a budget sitting below it is a governance issue, not just a facilities one.

It also reframes the backlog conversation. TPP19-07 requires NSW agencies to report the size of their maintenance backlog each year and identify measures to address it (NSW Treasury 2019). Funding below the floor is one of the most reliable ways to grow one.

6. From envelope to strategy

Figure 1 applies a single regime to every asset to show the envelope. No organisation should run a building that way. Treating a decorative light fitting like a main switchboard wastes money; treating a main switchboard like a decorative light fitting invites trouble.

The real value comes when organisations nominate regimes based on the criticality they have determined themselves, **asset by asset or space by space**. The space matters because the space is where the required functionality is delivered. A university's PC2 laboratory, data hall or main lecture theatre carries a very different consequence of failure from a storeroom or staff kitchenette, even when the assets inside are the same type.

Assigning regimes by space lets the criticality of the activity flow through to every asset that supports it.

Figure 2 sets out the process.




Figure 2: From criticality to funding: setting maintenance regimes asset by asset or space by space.

- 1. Map assets to spaces.** Build or confirm an asset register that records where each asset sits and which spaces it serves.
- 2. Rate criticality.** Apply the organisation's own criteria (life safety, compliance, business continuity, teaching or research impact, reputation) to spaces and, where needed, to individual assets.
- 3. Assign regimes.** Translate criticality ratings into High, Medium or Low, asset by asset or space by space.
- 4. Model the envelope.** Run OPEX and CAPEX under all three regimes, then the blended profile that results from the assignments.
- 5. Fund and review.** Compare the funded budget with the Low floor and the blended profile, agree the level of service, and revisit as use, condition and priorities change.

The blended profile will sit somewhere inside the envelope. Where exactly depends on how much of the portfolio is genuinely critical, which is a question the organisation should answer deliberately rather than leave to the budget to answer by default.

7. Putting the envelope to work

Once the envelope and the blended profile exist, they change the budget conversation in several practical ways.

- Budget defence.** A bid framed as "the blended profile for the service levels we agreed" is far easier to defend than "last year plus CPI". Any proposed cut can be expressed as a change of regime for specific assets or spaces, and so as a change in risk that someone has to accept.

- **Service-level trade-offs.** Executives can see what it would cost to move a building or group of spaces from Medium to High, what would be saved by moving non-critical areas to Low, and where the floor sits.
- **Capital planning.** Because the regime drives renewal timing, the same model informs the capital forecast, so renewal peaks are planned rather than discovered.
- **Policy and assurance.** The approach maps directly onto policy expectations: a prioritised program consistent with the accepted level of risk, with its capital and recurrent implications set out (NSW Treasury 2019), and maintenance that is risk-based and criticality-driven (BSI 2008). The regime targets also provide measurable objectives, such as the corrective-to-preventive ratio.

As *Strategic Asset Management* put it in 2008: "think twice about cutting the maintenance budget. Maintainers of tomorrow will be in a position to advise the cost to the business and its stakeholders of cutting the maintenance budget" (Strategic Asset Management 2008b). The three-regime model is one practical way to give that advice.

8. Conclusions and recommendations

A single maintenance number tells you what was spent. Three lines in the sand tell you what you are choosing. The High, Medium and Low regimes frame the full range of OPEX and CAPEX a building could reasonably demand. The Low line marks the floor below which statutory obligations are at risk. And criticality, set by the organisation asset by asset or space by space, determines where within the envelope the right answer lies.

We recommend that asset owners:

1. **Model the envelope.** Run every asset under all three regimes to establish the range of maintenance and renewal cost for each building and for the portfolio.
2. **Find the floor.** Compare current funding with the Low profile. Where funding sits below it, treat and report it as a compliance and governance risk.
3. **Set criticality deliberately.** Agree organisational criteria and apply them to spaces and assets, using the space as the unit where function is delivered.
4. **Fund the blended profile.** Build maintenance budgets from the blended profile and the service levels it represents, not from last year's figure.
5. **Use a long enough horizon.** Test results over 30 to 50 years so deferred renewals are not hidden beyond the planning window, and review regime assignments as use, condition and priorities change.

The example building in this paper is drawn from a demonstration model and is for illustration only. Validated benchmark values for specific building types are available from the Farsight catalogue, which covers more than 13,000 asset types.

Peter McCarthy

Managing Director

DC₂O Pty LtdE: peter.mccarthy@dc2o.com.auW: www.dc2o.com.au

References

- BSI (2008) *PAS 55-2:2008 Asset management. Part 2: Guidelines for the application of PAS 55-1*. London: British Standards Institution.
- Burns, P. (2000) 'Maintenance budgets', *Strategic Asset Management*, Issue 40, 14 July. Salisbury, SA: AMQ International.
- IAM (2015) *Asset Management Decision Making. Subject Specific Guidelines: Life Cycle Value Realisation*. London: Institute of Asset Management.
- ISO (2014) *ISO 55002:2014 Asset management. Management systems: Guidelines for the application of ISO 55001*. Geneva: International Organization for Standardization.
- NSW Treasury (2019) *TPP19-07 Asset Management Policy for the NSW Public Sector*. Sydney: NSW Treasury.
- Standards Australia (2012) *AS 1851-2012 Routine service of fire protection systems and equipment* (incorporating amendments). Sydney: Standards Australia.
- Strategic Asset Management (2008a) Issue 237, 24 March. Salisbury, SA: AMQ International.
- Strategic Asset Management (2008b) 'Business maintenance', Issue 243, 16 June. Salisbury, SA: AMQ International.

Appendix A: Maintenance level descriptions

All three regimes include compliance and statutory maintenance.

Dimension	High	Medium	Low
Purpose	Assure business continuity and safety for assets whose failure would materially impact operations, safety or reputation	Maintain acceptable service with controlled risk where some downtime can be tolerated	Minimise cost for non-critical assets with inconspicuous impact; accept higher variability and outages
Basis	OEM schedules + reliability / condition monitoring (RCM/CBM) + risk assessment	OEM schedules adjusted by utilisation/duty and funding envelope	Minimal PM to keep safe and legal; otherwise reactive / run-to-failure
PM coverage (vs OEM)	~90–110% (no deferral of critical tasks)	~60–80% (selective deferral / combination of tasks)	~20–40% (statutory and essential care only)
Typical task frequency	Equal to or shorter than OEM (may add CBM checks)	Slightly stretched vs OEM; tasks bundled by access windows	Extended intervals; only essential checks
Condition monitoring	Routine CM (e.g. vibration, oil, thermal, BMS analytics) with thresholds and work triggers	Light CM or periodic sampling on high-risk subcomponents	None, unless required by law
Corrective : preventive (target)	20:80 or better	~50:50	70:30 or higher
Backlog tolerance	0–5%	10–20%	25%+
Downtime tolerance	Minutes to hours; rapid restore and redundancy expected	Hours to days acceptable with workarounds	Days or more acceptable; fix when failed and resources available
Spares strategy	Stock critical spares; vendor framework for fast replacements	Limited spares; rely on supplier lead times	No dedicated spares
Replacement policy	Proactive: age/condition-based before end of life; plan renewals	Mixed: extend life where data supports; replace at failure-risk inflection	Reactive: replace on failure unless statutory or insurance requirements drive earlier
Life multiplier vs nominal life	0.9–1.1× (replace before performance or risk degrades)	1.0–1.3× (stretch where evidence supports)	1.3–1.8× or on failure (accept efficiency/comfort penalties)
Documentation and QA	Full digital records, SOPs, risk registers, asset health scores	Work orders documented; key assets tracked	Minimal records beyond compliance
Budget profile	Higher steady OPEX; lower outage risk and fewer CAPEX surprises	Moderate OPEX with occasional unplanned costs	Low steady OPEX; higher volatility and outage costs
Typical asset examples	Fire/life-safety systems, main switchboards, critical IT cooling, lifts in essential buildings, medical gases, main chillers/boilers serving mission-critical spaces	General AHUs/FCUs, lighting controls, non-critical pumps/fans, domestic hot water, secondary distribution boards	Decorative lighting, irrigation, minor signage, sundry fit-out items, low-impact site furniture