



DC₂O THOUGHT LEADERSHIP PAPER

Minding the Gap

Response latency, the P–F interval and the real measure of an asset management system

Peter McCarthy, Managing Director, DC₂O Pty Ltd

October 2026

Every asset management system makes the same quiet promise: when something starts to go wrong, we will notice and act before it matters. This paper is about how to tell whether that promise is being kept.

Contents

- Executive summary.....1
- 1 Introduction: the gap nobody measures.....1
- 2 Latency as a measure of management system effectiveness.....2
- 3 The P–F interval: the clock that matters.....4
- 4 The price of delay.....5
- 5 Where the time goes: information is the bottleneck.....6
- 6 Implications for asset owners.....7
- 7 Designing for low latency: the Farsight approach.....7
- 8 Measuring what matters.....8
- 9 Conclusions and recommendations.....9
- References.....10

Executive summary

Every asset management system makes the same implicit promise: when an asset starts to deteriorate, the organisation will notice and act before the deterioration matters. Whether that promise is kept depends less on the quality of the strategy than on the speed of the loop connecting an observation in the field to a completed intervention. This paper argues that the time taken to close that loop, which we call **response latency**, is one of the most revealing and least measured indicators of asset management effectiveness.

Drawing on reliability-centred maintenance, control theory and quality management, we show that latency only has meaning relative to the **P–F interval**: the window between the point at which a failure becomes detectable and the point at which it occurs. Where latency exceeds that window, preventive effort is wasted and the organisation pays twice. Because the cost of intervening rises in steps rather than smoothly, the penalty for a slow loop is disproportionate.

We also show that in most property portfolios the bottleneck is not the contractor or the field team. It is information: how quickly a finding becomes a costed, prioritised decision, and how quickly the long-range plan learns from what happened. We set out the attributes of the latency relationship, a practical set of latency measures, and the design principles behind DC₂O's Farsight platform, whose central purpose is to close that gap.

1 Introduction: the gap nobody measures

Picture a familiar sequence. A condition assessor notices blistering on a lecture theatre roof membrane in March. It is noted on a form, typed into a spreadsheet in May, and the audit report is issued in July. The finding joins several thousand others in a capital works list considered in the following year's budget round. It is funded, if it is lucky, the year after that, procured over the summer and repaired. By then the water has found the insulation, the ceiling tiles and, eventually, the lecture theatre timetable.

Nobody did anything wrong. Every step followed procedure. And yet the organisation knew about the problem roughly three years before it acted. The roof, it should be said, was not consulted about the budget calendar.

This is the gap: the elapsed time between an event (a defect starting, a condition changing, a risk emerging) and the action that addresses it. Most organisations measure the two ends of the chain, condition scores at one end and completed work orders at the other. Very few measure the chain itself. We think that is a mistake, because the length and predictability of that chain say more about the effectiveness of an asset management system than almost any other single indicator.

Is there a recognised relationship between response time and management system effectiveness? Not as a single named law. Nobody has yet had the decency to name one after us. But the principle runs consistently through several well-established disciplines, and, more usefully, it can be expressed, measured and managed.

2 Latency as a measure of management system effectiveness

2.1 A principle with many parents

The idea that the speed of the decision loop determines the effectiveness of the system appears independently in at least four fields:

- **Military strategy.** Boyd's observe–orient–decide–act (OODA) loop holds that the party able to cycle through observation, interpretation, decision and action faster than the environment changes holds the initiative (Boyd 1987).
- **Control theory and system dynamics.** A feedback loop whose delay is long relative to the response time of the system it controls does not control it; it oscillates around it. Sterman (2000) shows how delays in management systems produce overshoot and boom–bust behaviour. In asset portfolios the familiar cycle of deferral, followed by a capital crisis, followed by renewed deferral is exactly this kind of oscillation.
- **Quality management.** The cost of correcting a defect rises the further it travels before it is caught, which is the basis of the 1-10-100 rule (Labovitz, Chang and Rosansky 1992).
- **Asset management standards.** ISO 55001:2014 requires an organisation to react to nonconformities and incidents and take action to control and correct them (cl. 10.1), and to establish processes that proactively identify potential failures in asset performance (cl. 10.2). The 2024 revision renames preventive action as *predictive* action and revises its content (ISO 2024; UKAS 2024). Neither edition sets a time limit, but a requirement to identify *potential* failures only has value if the organisation can act before they become functional ones.

Burns (2005) put the information side of this bluntly in the title of an essay: if asset managers lose control of their information, they lose control of everything. Latency is the measurable symptom of that loss of control.

2.2 The four components of latency

Total latency is the sum of four distinct delays. Each has different causes and different remedies, so it pays to measure them separately.

Component	From	To	Typical drivers
Detection	Event or condition change	Recorded, usable fact	Inspection interval, data currency, transcription and re-keying
Decision	Recorded fact	Approved, funded action	Triage, costing, budget cycles, approval hierarchies
Execution	Approved action	Completed work	Procurement, contractor response, access and shutdowns
Learning	Completed work or new finding	Updated plan and model	Data entry, model refresh frequency, policy review cycles

Table 1. Components of response latency

2.3 Attributes of the relationship

It is relative, not absolute. A three-month response is excellent for a roof membrane and disastrous for a fire pump. What matters is latency relative to the rate at which consequence develops, which Section 3 formalises through the P–F interval.

It is non-linear in consequence. Cost and risk do not rise smoothly with delay; they jump when a threshold is crossed, such as corrosion initiating, a leak reaching the structure, or a component failing in service. Section 4 sets out the evidence.

Variance matters as much as the mean. A system with predictable latency is a capable system. The long tail, such as the small share of defects that sit unaddressed for two years, is where failures, compliance breaches and reputational damage come from. Average response time is a comfortable statistic; the 90th percentile is an honest one.

Backlog is latency in disguise. Little's Law (Little 1961) states that the average number of items in a system equals the arrival rate multiplied by the average time each item spends in it. If a portfolio generates 400 defects a year and each takes on average 18 months to resolve, the steady-state open backlog will be about 600 items. Halve the time and the backlog halves, with no change in workload. An ageing backlog is therefore a direct readout of loop latency, and arguably a more telling indicator than the dollar value of deferred maintenance.

Loops are nested, on different clocks. Operational loops (work orders) run in hours to weeks. Tactical loops (condition findings into the capital plan) run in months. Strategic loops (lessons into policy and the strategic asset management plan) run in years. Organisations are frequently quick in the first and glacial in the third, which is the difference between single-loop and double-loop learning (Argyris and Schön 1978).

A caveat worth stating plainly: speed is not the objective in itself. A deliberate, risk-assessed decision to defer is good asset management. An item nobody has looked at is unmanaged latency. An effective system can tell the two apart.

3 The P–F interval: the clock that matters

Reliability-centred maintenance describes the life of a failure mode with the P–F curve (Moubray 1997; SAE 2009). Point **P** is the earliest moment a developing failure can be detected; point **F** is functional failure. The time between them, the **P–F interval**, is the window in which an organisation can act on knowledge rather than react to events. Its length depends on the failure mode and on the detection technique: vibration analysis can reveal a failing bearing weeks before it becomes audible, and better detection moves P to the left.




Figure 1. The same failure mode, two management systems. Only the fast loop completes its response inside the P–F window. Source: DC₂O, after Moubray (1997).

Classical RCM guidance sets condition-monitoring intervals at no more than half the P–F interval, so that the defect is seen at least once with time left to respond. NASA's facilities RCM guide is more conservative, recommending that at least three monitoring tasks fall within the window and that intervals be cut to between a quarter and a third of their previous length once an alert limit is exceeded (NASA 2008). Moubray's *net* P–F interval is the window minus the inspection interval, which is the time genuinely available to act.

We extend that logic one step further. The inspection interval is only the detection component of latency; decision and execution time also consume the window. Defining total latency L as the sum of detection, decision and execution time gives a simple diagnostic ratio, $R = L \div P-F$:

Latency ratio R	What it means	What the organisation experiences
Below 0.5	In control, with margin	Potential failures are found and fixed while cheap
0.5 to 1.0	Racing the failure	Outcomes depend on luck and on the variance of latency
Above 1.0	Run-to-failure by default	Preventive effort is spent but reactive outcomes still occur

Table 2. Interpreting the latency ratio. The thresholds are DC₂O's proposed guide, extending RCM's half-interval rule.

The ratio offers two levers. An organisation can shorten L through faster detection, decision and execution, or lengthen the P–F interval through earlier detection technology. The second is often cheaper than reorganising the first, and the two combine.

3.1 The difficulty with buildings

P–F intervals in a building portfolio span several orders of magnitude, from weeks for rotating plant (rolling-element bearings are commonly quoted at two to ten weeks; Tractian n.d.) to years for roof membranes and sealants and decades for structure. A single organisation-wide response process will be over-engineered for some failure modes and dangerously slow for others.

Published P–F data for building fabric is also scarce. Reliability databases such as OREDA and Quanterion's NPRD give failure rates rather than warning windows, and most of the RCM literature concentrates on rotating and electrical plant. For fabric and finishes, the most practical source is an organisation's own record: the time components take to move between condition grades across successive audits, read alongside reactive work order history. That data usually exists, but rarely in a form that can be analysed.

4 The price of delay

Two widely cited heuristics capture the step-wise escalation of cost with delay. Neither is an empirical cost model, and both should be used as framing rather than evidence, but their shape is instructive.

De Sitter's "law of fives", developed for reinforced concrete durability, divides a structure's life into four phases and proposes that each dollar not spent in one phase becomes five dollars in the next (de Sitter 1984).




Figure 2. De Sitter's law of fives: relative cost of achieving the same durability outcome by phase of intervention. Source: DC₂O, after de Sitter (1984).

The 1-10-100 rule from quality management makes the same point about process stages: a dollar of prevention avoids ten dollars of internal correction and a hundred dollars of failure that reaches the customer (Labovitz, Chang and Rosansky 1992). Translated to asset information, an error caught at data capture is cheap, one caught at validation costs more, and one that has already shaped a twenty-year capital forecast and a funding submission is very expensive indeed.

The lesson is not the multiplier but the boundary. Costs jump when a threshold is crossed, and latency determines which side of the threshold the organisation acts on. In buildings, crossing F is rarely the end of the story: a roof leak becomes ceiling damage, then mould, then a closed teaching space. Each step is the next multiplier.

5 Where the time goes: information is the bottleneck

If latency were mainly a field problem, the remedy would be more inspectors and faster contractors. The evidence suggests otherwise. Wallsgrove (2004) reported a UK benchmarking study of six asset-intensive organisations across rail, electricity, water and nuclear. The study found that managing asset information cost 20% or more of total expenditure, whether or not it was managed well, and more than 25% where it was managed inefficiently. Spending on IT itself was modest, typically around 3% of expenditure. The waste lay in re-collecting data that had already been collected, in the time spent reaching information that was hard to access, and in too little effort to put usable information in front of the people who needed it.

Burns (2005) drew the corollary: organisations must distinguish good asset information, the business need, from IT, the enabler. A system that stores data without turning it into decisions adds cost without reducing latency.

Map those findings onto the four components of latency and the pattern is clear. Detection, decision and learning latency are almost entirely information latency. Only

execution is dominated by physical work. In our experience the property sector repeats the pattern: condition data collected at considerable expense, held in reports and spreadsheets, and re-collected at the next audit because the last one cannot be trusted, or found.

6 Implications for asset owners

The following illustration is indicative rather than measured, but it reflects patterns we see regularly in university and public-sector portfolios. Consider a roof membrane failure mode with an assumed P–F interval of about three years.

Component	Conventional loop	Low-latency loop
Detection	Four-yearly condition audit: average lag about 2.0 years	Continuous capture by FM staff plus annual walk-over: about 0.5 years
Decision	Report, then the next budget round and approval: about 1.5 years	Finding costed and prioritised on entry, reviewed quarterly: about 0.25 years
Execution	Procurement and works: about 0.5 years	Procurement and works: about 0.5 years
Total latency	about 4.0 years	about 1.25 years
Latency ratio R	about 1.3: run-to-failure by default	about 0.4: in control

Table 3. Illustrative latency profiles for a roof membrane failure mode (assumed P–F interval about 3 years). Figures are indicative only.

Two things stand out. First, the improvement comes almost entirely from detection and decision latency; nobody worked faster on the roof. Second, the conventional loop is not negligent. It is simply built on calendars, audit cycles and budget rounds, that are longer than the failure mode they are meant to govern.

The implications differ by audience:

- **Finance teams** are forecasting renewals from condition data that may be years old. A long learning latency means the capital plan describes the portfolio as it was, not as it is.
- **Facility managers** inherit the consequences as an ageing backlog. By Little's Law, the backlog will not shrink until latency does, however hard the team works.
- **Executives and boards** need defensible evidence that the organisation identifies and acts on potential failures, as ISO 55001 expects. Latency measures provide it; condition scores alone do not.

7 Designing for low latency: the Farsight approach

Farsight is DC₂O's lifecycle cost and asset information platform. Its central purpose follows directly from the argument above:

Farsight closes the gap between what your assets are telling you and what you do about it, so that the right action lands inside the window where it is still cheap.

Four design principles map onto the four components of latency:

Principle	Latency it reduces	How
Capture once, at the asset	Detection	Recording on mobile or desktop at the point of observation, structured by a catalogue of more than 13,000 asset types, so data is captured correctly the first time and not re-collected
Make every record mean something immediately	Decision	Each record is costed on entry using component-level lifecycle costs under High, Medium and Low maintenance regimes, all including statutory work, so a budget below Low signals unfunded compliance obligations
Put the right list in front of the right person	Decision to execution handover	Filters, saved queries and attention and overdue views deliver prioritised work without anyone compiling a report first
Keep the plan alive	Learning	The long-range lifecycle forecast recalculates as data changes, rather than waiting for the next audit cycle

Table 4. Farsight design principles mapped to latency components

It is worth being equally clear about what Farsight does not claim. It does not make contractors arrive sooner or procurement run faster. What it does is remove information delay from the loop and make the remaining delay visible, which is the precondition for managing it.

8 Measuring what matters

If every observation, finding, approval, work order and plan line carries a timestamp, latency becomes measurable with tools most organisations already own. We recommend the following indicators, reported by asset criticality:

Indicator	Definition	Reading it
Data currency	Median age of condition data	Compare against the P–F interval of each failure mode
Time to record	Observation to usable record	Should be days, not months
Time to decision	Record to approved, funded action	Report by priority band, including the 90th percentile
Time to complete	Approval to completed work	Report the distribution, not only the mean
Plan refresh latency	Completion or audit to updated forecast	Measures the learning loop directly
Latency ratio R	Total latency divided by P–F interval	Above 0.5 for a critical failure mode flags a system weakness
Aged backlog profile	Open items by age band	A Little's Law cross-check on the indicators above

Table 5. Recommended latency indicators

9 Conclusions and recommendations

The time between an asset telling us something and our acting on it is a direct measure of whether an asset management system works. It only has meaning relative to the P–F interval of the failure mode concerned, it penalises delay in steps rather than smoothly, and in property portfolios it is dominated by information rather than physical work. That makes it both a diagnostic and an opportunity: the cheapest lever in asset management is a faster loop, not a bigger budget.

We recommend that asset owners:

- 1. Timestamp the loop.** Record when each finding was observed, recorded, approved, completed and reflected in the plan. What is not timed cannot be managed.
- 2. Estimate P–F intervals for critical failure modes.** Start with RCM references for plant and your own successive condition data for fabric and finishes.
- 3. Separate managed deferral from unmanaged latency.** Deferral decisions should be explicit, risk-assessed and visible; everything else is delay.
- 4. Attack information latency first.** Capture at source, cost on entry and remove re-keying before investing in faster field response.
- 5. Shorten the strategic loop.** Replace periodic plan rebuilds with a living lifecycle model that updates as data arrives.
- 6. Report latency to the executive** as a leading indicator, alongside condition and expenditure.

None of this requires heroics from the field team, only a system that respects the clock. The roof, after all, has never once waited for the budget committee.

References

- Argyris, C. and Schön, D.A. (1978) *Organizational Learning: A Theory of Action Perspective*. Reading, MA: Addison-Wesley.
- Boyd, J.R. (1987) *A Discourse on Winning and Losing*. Unpublished briefing. Maxwell AFB, AL: Air University Library. [Details to be verified]
- Burns, P. (2005) 'If asset managers lose control of their information, they lose control of everything', *Strategic Asset Management*, Issue 174, 2 September. Salisbury, SA: AMQ International.
- de Sitter, W.R. (1984) 'Costs for service life optimization: the "law of fives"', *CEB-RILEM Workshop on Durability of Concrete Structures*, Copenhagen, 1983. CEB Bulletin d'Information No. 152. Lausanne: Comité Euro-International du Béton. [Bulletin number and date to be verified]
- ISO (2014) *ISO 55001:2014 Asset management: Management systems: Requirements*. Geneva: International Organization for Standardization.
- ISO (2024) *ISO 55001:2024 Asset management: Asset management system: Requirements*. Geneva: International Organization for Standardization.
- Labovitz, G., Chang, Y.S. and Rosansky, V. (1992) *Making Quality Work: A Leadership Guide for the Results-Driven Manager*. New York: HarperBusiness. [Publisher and year to be verified]
- Little, J.D.C. (1961) 'A proof for the queuing formula: $L = \lambda W$ ', *Operations Research*, 9(3), pp. 383–387.
- Moubray, J. (1997) *Reliability-centred Maintenance*. 2nd edn. Oxford: Butterworth-Heinemann.
- NASA (2008) *Reliability-Centered Maintenance Guide for Facilities and Collateral Equipment*. Washington, DC: National Aeronautics and Space Administration.
- SAE International (2009) *JA1011: Evaluation Criteria for Reliability-Centered Maintenance (RCM) Processes*. Warrendale, PA: SAE International.
- Sterman, J.D. (2000) *Business Dynamics: Systems Thinking and Modeling for a Complex World*. Boston: Irwin/McGraw-Hill.
- Tractian (n.d.) 'P-F curve (potential failure curve)'. Available at: <https://traction.com/en/glossary/p-f-curve-potential-failure-curve> (Accessed: 15 October 2026).
- UKAS (2024) *Technical Bulletin: Asset Management Systems (AMS) Transition to ISO 55001:2024*, 24 October. Staines-upon-Thames: United Kingdom Accreditation Service.
- Wallsgrove, R. (2004) 'How much does asset information cost?', *Strategic Asset Management*, Issue 143, 25 June. Salisbury, SA: AMQ International. Based on 'The real cost of asset information: how better costs less', ICOMS 2004, Sydney.

Items marked [to be verified] could not be confirmed against the original publication at the time of writing and should be checked before external release.

Peter McCarthy

Managing Director

DC₂O Pty Ltd

M: +61 (0) 400 XXX XXX

E: peter.mccarthy@dc2o.com.au

W: www.dc2o.com.au
