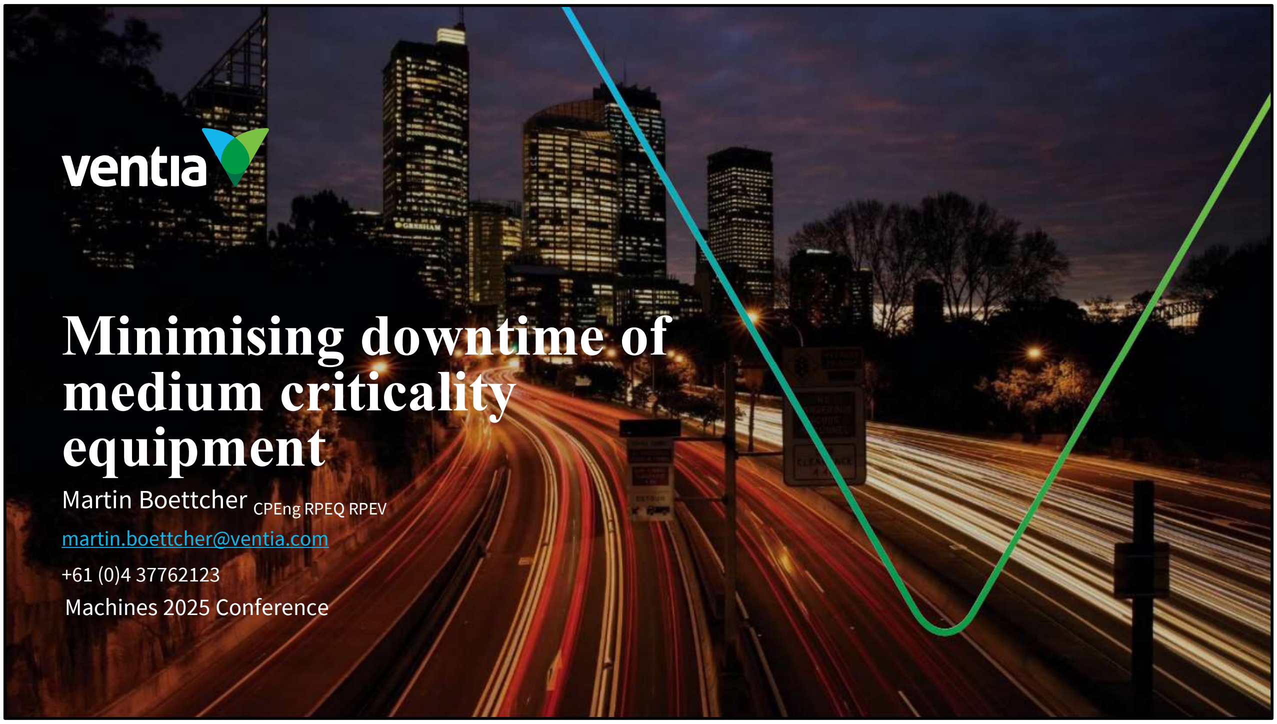

Today I'll be talking about how we can better manage and maintain medium criticality equipment—those assets that often sit just below the radar but can quickly become critical when things go wrong.
We'll explore how maintenance thinking has evolved across the industrial ages, look at examples of rapid change and what they mean for our mindset, and dive into practical strategies for dynamic maintenance—covering reliability categories, shifting criticality, and when to intervene.
I'll also touch on how emerging technologies like AI, drones, and digital twins are reshaping what's possible, and how we can 'nurse' equipment safely until the right opportunity for repair arises.
By the end, I hope to leave you with a clear framework for making smarter, more adaptive decisions—especially for those medium criticality assets that are often overlooked but increasingly vital.

Let's take a moment to look at how maintenance has evolved across the industrial ages—because understanding this journey helps us see where we're heading, especially with medium criticality equipment.

In Industry 1.0, the Steam Age, maintenance was entirely reactive. Machines were mechanical, and when something broke, you fixed it—often with a hammer and a lot of guesswork.

By Industry 2.0, the Electric Age, we saw mass production and more standardized equipment. Preventive maintenance started to emerge, but it was still mostly manual and time-based.

In Industry 3.0, the Digital Age, automation and control systems changed everything. We began using sensors and early SCADA systems, and maintenance became more data-informed. This is where Condition-Based Maintenance really started to take hold.

Then came Industry 4.0, the Smart Age. Connectivity, IoT, and AI allowed us to monitor assets in real time. Predictive maintenance became possible, and we started using digital twins and performance models to anticipate failures—especially useful for medium criticality assets that don't always get top priority.

Now we're entering Industry 5.0, the Human-Centric Revolution. It's about collaboration between humans and machines, and tailoring maintenance strategies to operational context and business impact.

Looking ahead to Industry 6.0, we're talking about biological-digital integration—systems that are self-aware, adaptive, and possibly even self-maintaining. Imagine AI agents embedded in equipment that negotiate their own maintenance schedules.

The key takeaway is this: maintenance has gone from reactive to predictive, and now it's becoming adaptive. And for medium criticality equipment, that means we need strategies that are flexible, intelligent, and context-aware—because these assets often sit right on the edge of risk and opportunity.

Examples of rapid change

Easter morning 1900: 5th Ave, New York City. Spot the automobile.

Source: US National Archives.

Easter morning 1913: 5th Ave, New York City. Spot the horse.

Source: George Grantham Bain Collection.

#	Example	Timeframe	Description
1	Disappearance of Horses in Cities	1900–1910	Cars replaced over 100,000 horses in NYC within a decade.
2	Smartphone Adoption	2007–2015	iPhone launch triggered global smartphone adoption.
3	Remote Work Shift	2020–2021	COVID-19 forced a global transition to remote work.
4	Rise of E-Commerce	1995–2005	Amazon and eBay popularized online shopping.
5	Film to Digital Photography	2000–2010	Digital cameras and smartphones replaced film.
6	Music Industry Shift	1999–2009	Napster and iTunes changed how music was consumed.
7	AI and Chatbots	2022–2025	Tools like ChatGPT and Copilot became widely used.

This slide shows how quickly entire systems can transform. In just a decade, cars replaced over 100,000 horses in New York City. The iPhone reshaped global communication in under eight years. And more recently, AI tools like ChatGPT and Copilot have gone from novelty to everyday tools in just a few years.

These aren't just interesting historical facts—they're a reminder that change is accelerating, and it's not limited to consumer tech. It's happening in our industry too. Equipment design, materials, control systems, and even the way we define reliability are evolving faster than ever.

So what does that mean for us? It means we can't rely on static plans or legacy assumptions. We need to expect change—in asset performance, in failure modes, and in what's considered critical. And we need to build maintenance strategies that are flexible, responsive, and data-driven.

This shift in mindset—from planning for stability to preparing for disruption—is especially important for medium criticality equipment. These assets often sit just below the radar, but in a rapidly changing environment, they can quickly become the weak link.

The takeaway here is: if we've seen this much change in the past 10 years, imagine what the next 10 will bring. Our maintenance mindset needs to be ready—not just for what we know, but for what's coming.

How equipment is changing – Solar panel example

Powermag, April 2025

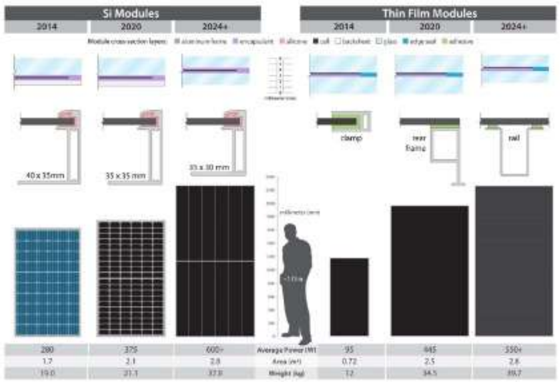

- Impact of US Tarriffs on the design of equipment.
- Considering use of steel vs aluminium in design:
  - Corrosion failure mechanisms introduced.

- Reliability of “NEW” equipment almost guaranteed to be worse than those they are replacing – consider from AM perspective.

- Breakage of new panels increasing issue:
  - “50% of modules using 2mm glass suffered breakage within nine months of deployment”.
  - Larger size. Reduced glass thickness. Different glass composition.
  - Reduced structural support designs.

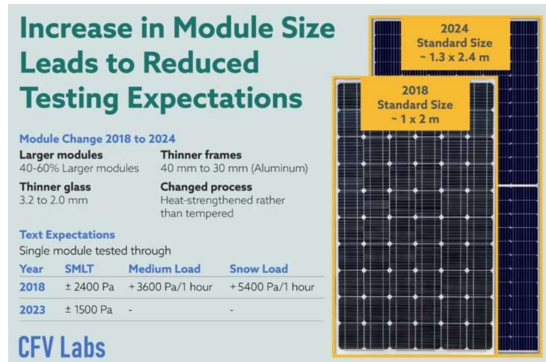

This slide highlights how even well-intentioned design improvements can introduce new reliability challenges. Take solar panels, for example. In the push for higher efficiency and lower costs, manufacturers have moved to thinner glass and larger formats. But studies show that 50% of modules using 2mm glass suffered breakage within nine months—a huge reliability issue.
These changes aren’t just about materials. They’re also driven by external factors like tariffs, which influence whether we use steel or aluminium—each with its own corrosion risks.
The key takeaway here is that newer doesn’t always mean better from a maintenance perspective. When we replace older, proven equipment with newer designs, we often inherit unknown failure modes. That’s why asset managers need to be proactive—anticipating reliability risks before they show up in the field.
This is a great example of how equipment evolution is outpacing traditional maintenance planning, and why dynamic, data-driven strategies are becoming essential.

Trends in maintenance + Power Industry example

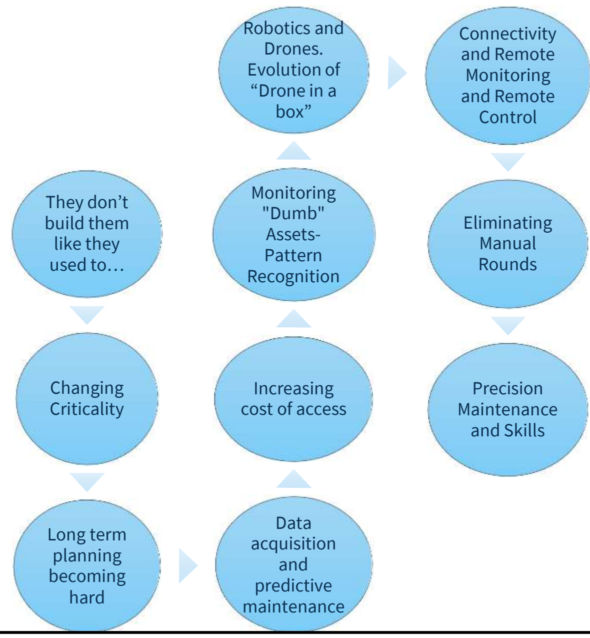

Period	Power Industry changes
Pre-1970	Operations experienced based. Mechanical/ pneumatic controls
Seventies	Analogue control (0-10mA, 4-20mA). Maintenance paper-based. Early days of SCADA systems. RCM gaining ground.
Eighties	Integrated SCADA systems. Preventive Maintenance (PM) becomes more widespread, some Condition-Based Maintenance (CBM).
Nineties	Advancements in SCADA, Distributed Control Systems (DCS). Increased use of CBM, Computerized Maintenance Management Systems (CMMS). Mobile phone connectivity (1G)
2000's	Integration of DCS and SCADA systems with improved real-time data collection and remote monitoring capabilities. Growth of Predictive Maintenance (PdM). Data interfacing between SCADA and IT such as Pi.
2010's	Early development of IoT and smart sensors. Initial adoption of digital twins. Proactive Maintenance. Penetration of wind and solar starts impacting on the operational profile of large units.
2020's	Rise of digital twins. Linear systems monitoring using mass data capture and AI or model-based assessment. Cyber. Reduced & cycling loading requirements. Criticality becomes increasingly important due to cost cutting. Lots of equipment moved to a corrective maintenance strategy. Closure planning.

Let's talk about how maintenance is evolving—because it's not just about fixing things anymore. It's about adapting to a world where equipment, data, and expectations are all changing faster than ever.

First, we're seeing a shift in criticality. What used to be considered low or medium criticality can suddenly become high-risk depending on market conditions, weather, or operational context. That means our maintenance strategies need to be dynamic, not static.

Second, data acquisition and predictive maintenance are becoming the norm. We're no longer relying solely on scheduled inspections. Instead, we're using sensors, AI, and pattern recognition to predict failures before they happen—even on so-called "dumb" assets.

Third, access costs are rising. Whether it's remote locations, safety restrictions, or just the cost of sending people out, we need to eliminate manual rounds wherever possible. That's where drones, robotics, and remote monitoring come in. The idea of a "drone in a box" that can autonomously inspect assets is no longer science fiction—it's happening now.

And finally, there's a growing emphasis on precision maintenance and skills development. As systems become more complex, we need technicians who can interpret data, work with digital twins, and understand the interplay between

mechanical, electrical, and software systems.
So the big takeaway here is this: maintenance is no longer just about keeping things running—it's about staying ahead of change. And that means embracing technology, rethinking criticality, and building systems that can adapt in real time.

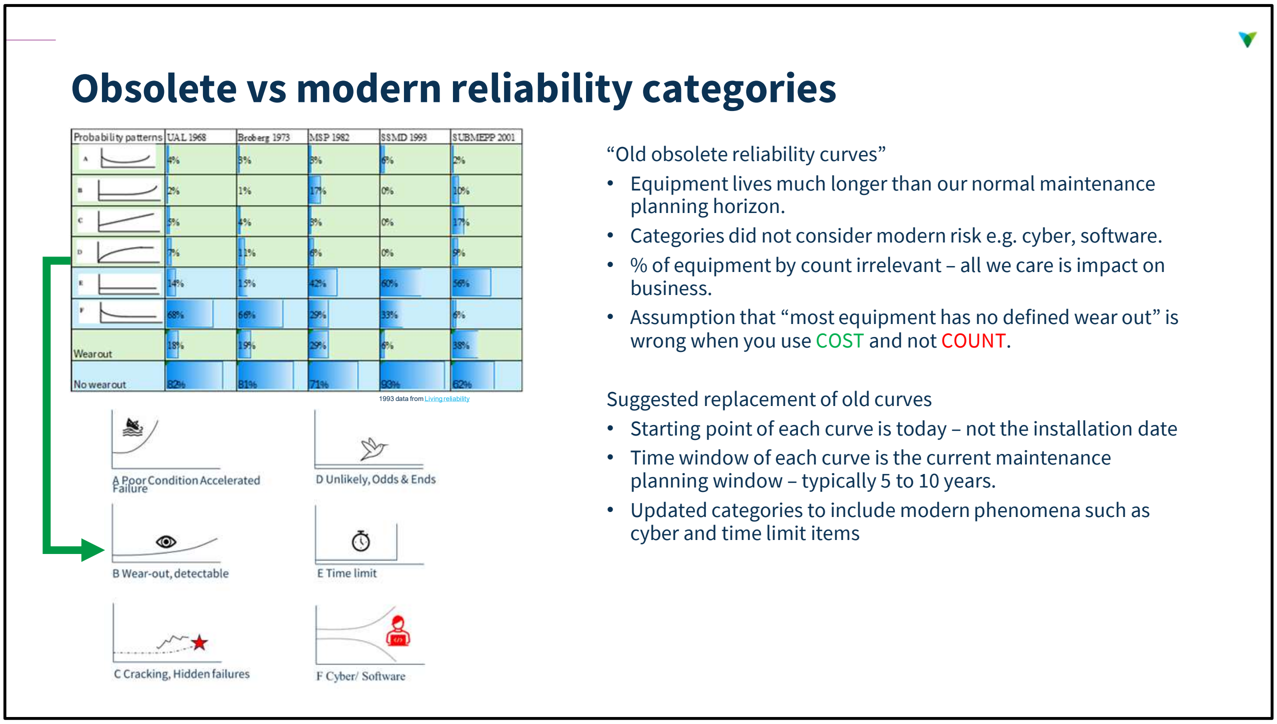

This slide is about how our traditional understanding of equipment reliability is becoming outdated—and why we need to rethink how we categorize and plan for failure.

Historically, we used what I call the 'old reliability curves'—these assumed that most equipment didn't have a defined wear-out phase, and that failure rates were relatively stable over time. But those models were built in a world where equipment was simpler, and risks like cybersecurity, software bugs, and planned obsolescence didn't exist.

Today, that's no longer the case. Equipment is more complex, more interconnected, and more vulnerable to non-physical failure modes. So we need a new approach—one that starts from today, not from the installation date, and that looks at the next 5 to 10 years, which is the real window we plan maintenance in.

The modern approach also shifts focus from counting equipment to cost and impact. It doesn't matter if 90% of your assets are low-risk if the 10% that fail can shut down your plant.

We also now include categories for things like cyber risk, software failure, and time-limited components—like sensors or fire systems that must be replaced on a fixed schedule regardless of condition.

The key message here is: if we're still using reliability models from the 1990s, we're flying blind. We need to update our thinking to reflect the real risks and realities of modern equipment—and that means embracing data, context, and a more dynamic view of reliability.

Practical reliability categories – Industry 4.0

Category	Concept sketch	Overview/ Examples	Management Approach
A Poor Condition Accelerated Failure		Accelerated deterioration due to poor installation/ condition. E.g. Pump not aligned, corroded structure	Urgent intervention to avoid catastrophic/ consequential damage
B Detectable Wear-out	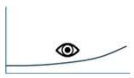	Typical equipment subject to wear. Deterioration visible to operations. E.g. Pump pumping slurry, conveyor systems	Monitor over time. CBM.
C Cracking, Hidden failures		Fatigue and Creep failures – detectable during major inspections only. E.g. Cracks in turbine blades, piping creep damage	Planned routine inspections. Compliance to regulations.
D Unlikely, Odds & Ends		Unlikely that failure will occur in current planning horizon. Covers many small items, civil structures etc.	Occasional inspections. Watch for multiple small issues accumulating.
E Time limit	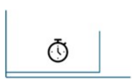	Fixed contract duration replacements, planned obsolescence	Consider alternate commercial agreement or just replace on the day
F Cyber/ Software	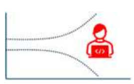	Cyber impacts or software (e.g. Crowdstrike) can take down whole plant/ entire organization.	Comprehensive cyber strategy

This slide introduces a practical framework for categorizing equipment reliability in the context of Industry 4.0. It's not just about whether something will fail—it's about how, when, and what we can do about it.

We start with Category A – Poor Condition Accelerated Failure. These are assets that are deteriorating rapidly due to poor installation or environmental exposure. Think of a misaligned pump or a corroded structure. These need urgent intervention to avoid catastrophic damage.

Then we have Category B – Detectable Wearout. These are your typical workhorses—pumps, conveyors—where wear is visible and predictable. Condition-based monitoring works well here, and maintenance can be planned.

Category C – Cracking and Hidden Failures includes fatigue and creep damage that's only detectable during major inspections. Turbine blades and high-pressure piping fall into this group. These require scheduled inspections and compliance checks.

Category D – Unlikely, Odds & Ends covers low-risk items like civil structures or small components. They're unlikely to fail in the current planning horizon, but we still need to keep an eye on them.

Category E – Time Limit is for assets with planned obsolescence or fixed replacement schedules—like fire systems, sensors, or cranes. These are replaced

based on contract duration or regulation, not condition.
Finally, Category F – Cyber/Software. This is a newer but critical category. A cyberattack or software failure can take down an entire plant. These risks require a comprehensive strategy, not just technical fixes.
The key takeaway is that reliability isn't one-size-fits-all. By categorizing assets this way, we can tailor our maintenance strategies to the actual risk and behavior of each type of equipment—making our planning more effective and our interventions more timely.

Reliability groupings over time

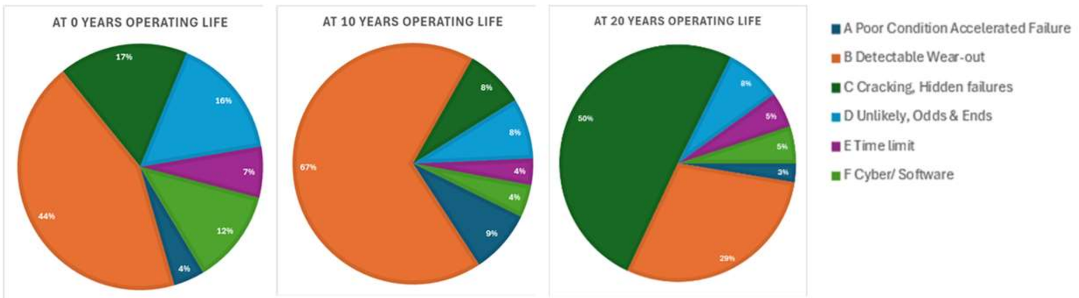

Significant investments for first period in B Detectable Wear-out. As the facility ages, C Cracking, Hidden Failures become dominant

Paper contains sanity check example from the power industry, which is consistent with these results

- Data based on water facility using maintenance capital replacement and overhaul plan.
- Each graph represents a 10 year window starting from the operating life as shown.
- % based on total maintenance investment, NOT number of equipment
- E grouping assigned to planned replacements e.g. chemical sensors, fire equipment, cranes. Pumps and motors assigned to B. High wearing solids applications assigned to A. DCS system and equipment assigned to F.

This slide shows how the focus of maintenance investment shifts over the life of a facility—and why our reliability planning needs to evolve with it.
The data here comes from a water facility’s capital replacement and overhaul plan. What’s important is that the percentages shown are based on total maintenance investment, not the number of assets. That’s a key distinction—because it’s not about how many items you have, it’s about where your money and risk are concentrated.
In the early years of operation, most investment goes into Category B – Detectable Wear-out. These are your pumps, motors, and other assets that wear down in predictable ways. You can monitor them, plan for them, and manage them efficiently.
But as the facility ages, the focus shifts. Category C – Cracking and Hidden Failures becomes dominant. These are issues like fatigue in turbine blades or creep in piping—problems that are harder to detect and more expensive to fix. They often require major inspections or shutdowns.
You’ll also see consistent investment in Category E – Time Limit items, like sensors, fire systems, and cranes. These are replaced on a schedule, regardless of condition, due to compliance or contract terms.
And increasingly, we’re seeing investment in Category F – Cyber/Software. As

digital systems become more integrated, the risk of a cyber event or software failure grows—and the impact can be massive.

The key takeaway is this: reliability isn't static. What you focus on in year one won't be the same in year ten. By understanding how reliability groupings shift over time, we can better allocate resources, anticipate risks, and ensure long-term performance.

Changing criticality in dynamic maintenance context

- Organisational expectations constantly changing depending on market conditions, demand, weather etc.
- Condition of equipment becoming "live"
  - Better monitoring available
  - Much better ability to understand the condition given improved data availability.
  - Better understanding of "dumb" assets
- Criticality of equipment changing depending on what the overall plant status is
  - Criticality MUST include likelihood to work under dynamic conditions
  - Traditional definition of "worst case" criticality is obsolete when considering Dynamic Maintenance.

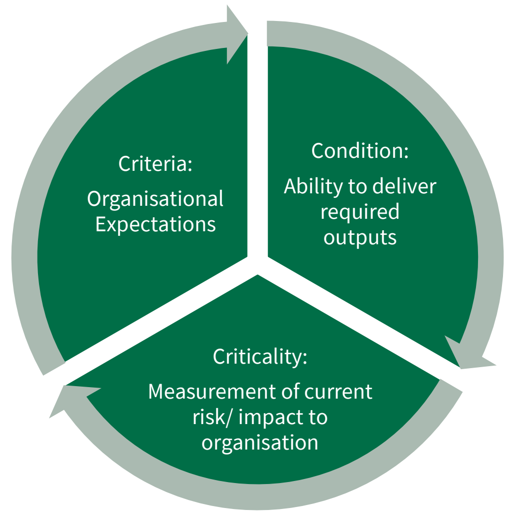

This slide is about a fundamental shift in how we think about equipment criticality. Traditionally, we've assigned criticality based on worst-case scenarios—fixed definitions that don't change. But in today's dynamic operating environments, that approach is no longer sufficient.

Criticality is now context-dependent. It changes based on market conditions, weather, demand, and even the operational status of other equipment. For example, a backup pump might be low criticality—until the primary pump fails. Suddenly, it becomes essential.

We're also seeing the condition of equipment becoming 'live'. With better monitoring and data availability, we can assess asset health in real time. That means we're no longer guessing—we're responding to actual conditions.

This dynamic view also helps us better understand so-called 'dumb' assets—those without built-in intelligence. With external sensors and pattern recognition, we can now monitor these assets more effectively and factor them into our planning.

So what does this mean for maintenance? It means we need to redefine criticality not as a static label, but as a fluid measurement of current risk and impact. It must reflect the likelihood of an asset performing under dynamic conditions—not just its theoretical importance.

The takeaway here is: if your criticality model doesn't change with your plant's status, it’s not helping you make good decisions. Dynamic maintenance requires dynamic thinking—and that starts with how we assess and respond to risk.

Which equipment do we consider "Medium Criticality"?

Indicative quantities	% of items	Cumulative
C5	50%	50%
C4	25%	75%
C3	15%	90%
C2	7%	93%
C1 (top)	3%	100%

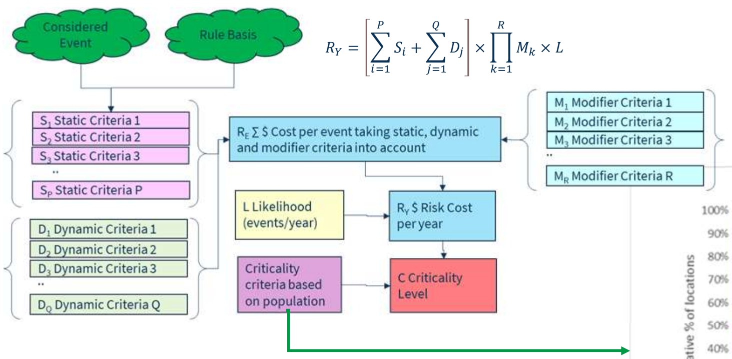

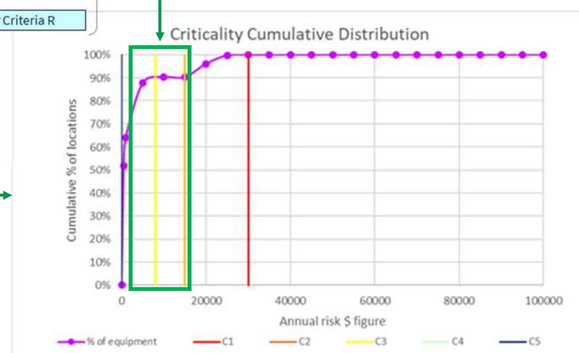

Focus on Criticality 2 and 3 items
Remember: All Criticality does is put things in 5 "buckets"
- Anything better than 20% accuracy a waste of effort

Now let's talk about how we define and identify medium criticality equipment—because this is where most of our maintenance effort and decision-making actually happens.

The chart here shows a typical criticality distribution, where equipment is grouped into five buckets—C1 through C5—based on cumulative risk or impact. The top 3% of assets, the C1 group, are your most critical. These are the ones that, if they fail, could shut down your plant or cause major safety or compliance issues.

But what's really interesting is that Criticality 2 and 3—C2 and C3—make up just 22% of the asset base, yet they represent the sweet spot for proactive maintenance. These are the assets that are important enough to matter, but not so critical that they're already getting constant attention.

When we talk about 'medium criticality,' we're really talking about this middle band—the 7% to 25% of assets that can quietly become high risk if neglected. These are often overlooked because they're not screaming for attention—until they fail.

The key point here is that criticality is a tool for prioritization, not perfection. As the slide says, anything better than 20% accuracy is probably a waste of effort. What matters is that we're using a structured approach to focus our resources where they'll have the most impact.

So when you’re building your maintenance strategy, focus on C2 and C3 in addition to the work already done on C1. That’s where you’ll get the best return on effort—and where you can prevent the most surprises.

Scenarios where medium criticality equipment become critical

All pictures from Brady Heywood report C4 failure

- Commissioning works increases the expectation of normally 2nd tier equipment
  - Callide C4 incident 2021: Unit Battery Charger, normally considered a medium criticality device, became super critical during switching operations.
  - Industry wide naming convention error: “Battery Charger” as opposed to “Unit Main DC supply”
  - Importance of risk assessment during testing and commissioning operations.
- Protection systems – such as fire systems:
  - When no fire, the system is low criticality
  - In case of fire, it becomes the most important system on site.
  - How do you manage maintenance impairment of these systems?
- Redundant equipment out of service
  - Many plant designed with N+1 redundancy.
  - When e.g. one pump out of service, the “lower criticality” resulting from redundancy disappears.
  - Redundancy and the likelihood of failure needs to be taken into account when considering criticality.

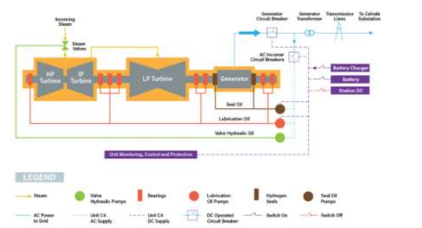

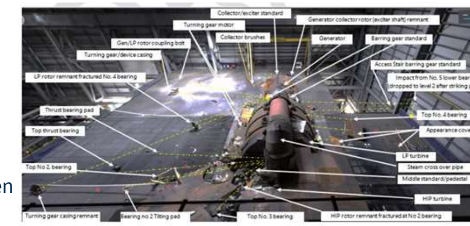

This slide highlights a key challenge in maintenance: medium criticality equipment doesn’t always stay medium. Its importance can spike depending on the situation—and if we’re not ready, the consequences can be serious.

A great example is the Callide C4 incident in 2021. A unit battery charger—normally considered medium criticality—became absolutely critical during switching operations. The issue? It was mislabelled as a ‘battery charger’ instead of its true role as the main DC supply. That naming error led to a misjudgement in risk—and ultimately, a major failure.

Another scenario is fire protection systems. When there’s no fire, they’re low criticality. But in an emergency, they instantly become the most important system on site. So how do we manage maintenance when the criticality is conditional?

Redundancy is another factor. Many plants are designed with N+1 redundancy—meaning one spare unit is available. But if that spare is out of service, the remaining unit’s criticality jumps. Suddenly, what was a backup becomes a single point of failure.

The takeaway here is that criticality is dynamic. It’s not just about the asset—it’s about the context. That’s why risk assessments during commissioning, testing, and operational changes are so important. We need to be asking: What happens if this fails right now? Not just What does this usually do?

When planning maintenance, especially for medium criticality assets, we need to think ahead, monitor context, and be ready to reassess—because the stakes can change in an instant.

When to intervene in medium criticality equipment?`

- Traditionally:
  - Preventive Maintenance (PM): For known wear patterns (e.g. pumps pumping slurry)
  - Condition-Based Maintenance (CBM): For detectable degradation (e.g. vibration, temperature)
  - Breakdown Maintenance: Cost benefit position – consider redundancy and spares availability in decision

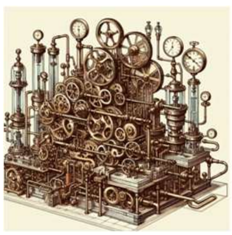

- Consider evolution of technology
  - Can we use AI and drones to monitor previously “not worth it to monitor” assets?
  - Structural inspections previously expensive through accurate scanning.
  - Where there are measurements already in place can these be monitored by performance models, digital twins or AI to detect emerging failures?
- How can we employ Agentic AI to look after our equipment on site?

When should we intervene in medium criticality equipment? Traditionally, we’ve had three main approaches:
First, Preventive Maintenance—used when we know the wear patterns, like pumps handling slurry. We schedule interventions based on expected degradation.
Second, Condition-Based Maintenance—this is for assets where we can detect degradation through vibration, temperature, or other sensor data. It’s more responsive and data-driven.
And third, Breakdown Maintenance—where we accept the risk of failure because the cost of intervention outweighs the impact. This is often used when redundancy or spare parts are readily available.

But here’s the shift: technology is changing what’s worth monitoring. AI, drones, and digital twins are making it feasible to track assets that were previously too costly or complex to monitor. For example, structural inspections that used to require scaffolding and shutdowns can now be done with high-resolution drone scans.
And where we already have measurements in place, we can use performance models and AI to detect emerging failures before they become visible. This opens the door to Agentic AI—systems that can autonomously monitor, assess, and even

initiate maintenance actions.
The key message is: intervention timing is no longer just about cost—it's about capability. With smarter tools, we can intervene earlier, more precisely, and with better justification. And that's especially important for medium criticality equipment, where the risk is real but not always obvious.

Principles of "nursing" equipment till opportunity

Key Concept	Description	Why It Matters
1. Process Safety	Ensure the equipment can continue operating without compromising safety systems or personnel.	Prevents escalation into hazardous events or regulatory breaches.
2. Exclusion Zones & Controls	Implement physical or procedural barriers to limit access or interaction with compromised assets. Is load reduction an option?	Reduces risk of injury or accidental interference during degraded operation.
3. Monitoring & Early Warning	Increase frequency or sophistication of monitoring (e.g. sensors, manual checks).	Detects worsening conditions before failure, enabling timely intervention.
4. Business Impact Assessment	Evaluate the cost of downtime vs. the risk of continued operation.	Supports informed decisions on whether to delay or expedite repairs.
5. Secondary Damage & Cost	Consider how continued operation may cause further damage to surrounding systems.	Helps avoid compounding repair costs and extended outages.

Do we know what is happening?

Do we trust the assessment?

Who carries the risk?

Who makes the decision?

Consequential damage?

What is lead time for parts?

What is the financial impact?

Sometimes, we can't fix everything right away. Whether it's due to access, cost, or operational constraints, we often need to 'nurse' equipment along until the right opportunity arises. But that doesn't mean we do nothing—there are clear principles we can follow to manage the risk.

First, we start with Process Safety. The number one priority is to ensure the equipment can continue operating without compromising safety systems or personnel. If there's any risk of escalation into a hazardous event, we need to act immediately.

Second, we implement Exclusion Zones and Controls. If the equipment is degraded, we reduce the risk of accidental interference by limiting access—physically or procedurally. Sometimes, even load reduction can buy us time.

Third, we ramp up Monitoring and Early Warning. This could mean more frequent manual checks or deploying sensors to detect worsening conditions. The goal is to catch any signs of deterioration before failure occurs.

Fourth, we conduct a Business Impact Assessment. We weigh the cost of downtime against the risk of continued operation. This helps us decide whether to delay or expedite repairs.

And finally, we consider Secondary Damage and Cost. Sometimes, letting a component run to failure can cause cascading issues. So we ask: What else could be

affected? What’s the lead time for parts? What’s the financial impact?
These principles help us make informed, risk-based decisions. Because in dynamic maintenance, it’s not just about fixing things—it’s about managing uncertainty intelligently.

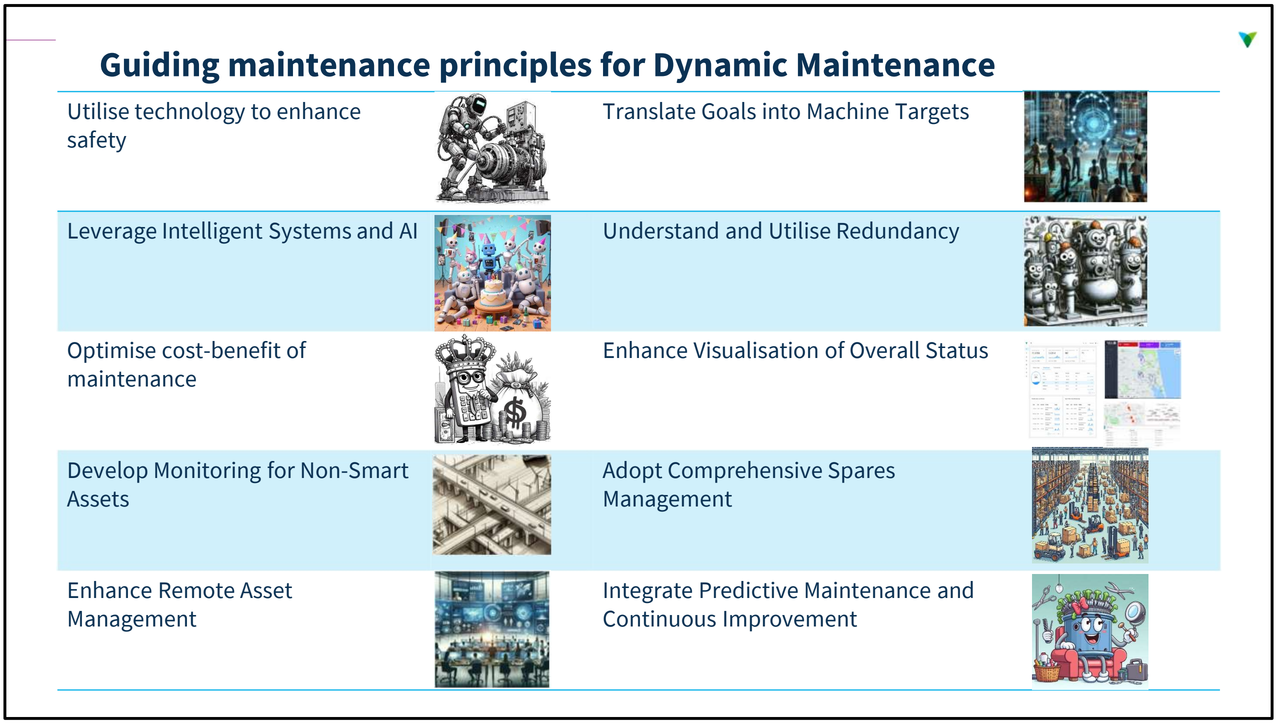

To wrap up, this final slide brings together the key principles that underpin dynamic maintenance—especially when it comes to managing medium criticality equipment.

Throughout this presentation, we've seen how equipment design is changing, how reliability categories are evolving, and how criticality itself is no longer static.

Medium criticality assets sit in that crucial middle ground—they're not the most urgent, but they're often the ones that surprise us when things go wrong.

So how do we manage them effectively? We start by translating business goals into machine targets—aligning asset performance with what the organisation actually needs. Then we leverage technology—AI, remote monitoring, drones—to enhance safety and visibility, especially for assets that were previously too costly or complex to monitor.

We also need to understand and utilise redundancy. Medium criticality equipment often plays a backup role, but when redundancy is compromised, its importance skyrockets. That's why we must track not just condition, but context.

Intelligent systems and predictive models help us anticipate failures before they happen. And by integrating these with comprehensive spares management, we reduce downtime and avoid scrambling for parts.

Finally, we focus on continuous improvement—using data to refine our strategies,

and ensuring that even non-smart assets are part of the picture.
The big takeaway is this: dynamic maintenance isn't just a strategy—it's a mindset. It's about being responsive, informed, and proactive. And when applied to medium criticality equipment, it's the difference between smooth operations and unexpected disruptions.

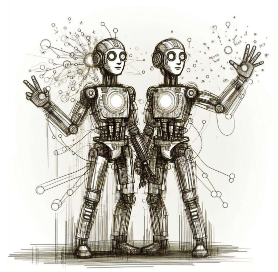

Questions?