


Maintenance control in 90 days.

Pick the KPI that hurts most, run one 90-day play, and have a number that moved by day 90. Ten plays, a 5-minute self-assessment, and the cadences that make a fix stick.

REACTIVE WORK, % OF LABOUR HOURS • EXAMPLE PLANT • 13 WEEKS



Maintenance control in 90 days.

A maintenance manager's field guide to getting work under control.

Publisher Fabrico **Published** October 2026

Copyright 2026 Fabrico. All rights reserved. No part of this field guide may be reproduced, stored or transmitted in any form without the publisher's written permission, other than brief quotations in articles and reviews.

It is shared with the understanding that the publisher is not providing engineering, legal or other professional services for your specific plant. Results vary. A plant that reads this and changes nothing will get what it has always gotten.

ON THE SOURCES

Every benchmark, threshold and method in this guide has a citation: SMRP Best Practices, ISO 55000, SAE JA1011 and JA1012, API inspection codes, and the books listed in Part 5. If you disagree with a number, disagree with the source.

ON THE SCREENSHOTS

The product screens in the "How to run this in Fabrico" boxes are real Fabrico screens from a demonstration plant, with sample data. They show the method, not a specific customer.

CONTACT

team@fabrico.io · fabrico.io

WHO THIS IS FOR

- **Plant managers and operations directors** who sign off the maintenance budget and get the Monday question about Line 4.
- **Maintenance managers and reliability engineers** in plants of roughly 50 to 2,000 people, with a CMMS that is only half used.
- **Anyone** who has watched the same asset fail three times and wants the fourth time not to happen.

WHO IT IS NOT FOR

- A plant looking for a tool to buy instead of a system to run. The tool comes last here, on purpose.
- A reader who wants theory. There are better books for that, and they are listed in Part 5.
- Level 5 plants. If you are one, we would like to hear what you do differently.

The plants we meet keep asking the same question. So we wrote the answer down.

Fabrico is a cloud-based MES and OEE platform for manufacturers, with maintenance management (CMMS) built in. Our consultants keep sitting in plants where the PM program exists on paper and the schedule is a fiction by Tuesday. The software is rarely the problem. The system around it is. This field guide describes that system: five stages of work, ten plays, and the benchmarks that tell you whether it is holding.

01 A 5-minute check-in that tells you where to start

Ten statements, one honest score, and the two plays that match it. Page 7.

02 The five-stage work cycle, with the failure patterns named

Identify, plan, schedule, execute, close out. Each stage with what good looks like and a cadence to run it.

03 Ten plays, each with a 90-day roadmap

Pick the one that matches your worst number. Every roadmap starts on day 1, not next quarter.

04 The benchmarks, sourced

Eight KPIs with action, target and world-class thresholds, every one with its SMRP, ISO or SAE citation.

05 How each play runs in Fabrico, on real screens

The actual screens a maintenance manager works in.

Read it in an hour. Run one play for 90 days. If you want to see your own lines in it, the last page tells you how. Until then, it is just a good field guide.

The Fabrico team

fabrico.io



Contents

PART 1 • BEFORE YOU START

Why this playbook exists	5
What this is, what it is not, how to use it	6
The 5-minute Maintenance Control Check-In	7
What your score means, and where to start	8

PART 2 • FOUNDATIONS

Why foundations matter: the work cycle	10
Work identification	11
Planning, and the ready-to-schedule gate	13
Scheduling	15
Work execution	17
Close-out and the feedback loop	19

PART 3 • WHERE ARE YOU

The maturity model	22
The self-assessment	23
The benchmark snapshot	25
Where the numbers come from	26

PART 4 • THE TEN PLAYBOOKS

1 • PM compliance recovery	28
2 • Repeat failure investigation	30
3 • Backlog growth control	32
4 • Reactive work reduction	34
5 • Schedule compliance improvement	36
6 • Inspection compliance recovery	38
7 • Spare parts management	40
8 • Condition-based monitoring	42
9 • Downtime reduction	44
10 • Asset reliability and MTTR	46

PART 5 • TOOLS AND CADENCES

RCA quick guide	49
The 90-day implementation cadence	51
References	53
Where Fabrico fits	54
What happens in the 30 minutes	55

READING ORDER

In a hurry? Page 7 (the check-in), then the two plays it points you to, then page 55. The rest is there when the first play holds.

The gearbox is not the problem. The system that lets it fail every few months is.

Monday morning at a mid-size plant. The operations review runs from 8 to 8:45. Near the end, the plant manager asks about the three unplanned outages on Line 4. The maintenance manager explains. The gearbox that failed was the same one that failed in January. And in August. And last October. The plant manager nods, says something about "getting to the root cause", and the meeting ends.

Nothing changes. The gearbox will fail again in eight to ten weeks.

ONE GEARBOX, LINE 4 • THE FAILURE CALENDAR



This is the shape of the problem. Not one person doing anything wrong. Not one moment of negligence. Just a plant that has quietly accepted that some things fail every few months and nobody is going to stop them. Until one of those things stops the plant for a week.

If you run a mid-size plant, or the maintenance in one, you already know what is broken. The PMs are slipping. The backlog is growing. A specific pump on a specific line has its own failure calendar. You do not need another diagram of the reliability pyramid. You need a way to get the next thing fixed before it shows up on Monday morning for the fourth time.

This playbook gives you that.

PAGE 7

Five minutes, ten statements, one score. It tells you which play to run first.

PAGE 25

Eight KPIs with the thresholds. Mark yours red, yellow or green.

PAGES 28 TO 47

Ten plays. Each roadmap starts on day 1, with something you can do this week.

Concrete steps, not frameworks about frameworks

WHAT THIS IS

- A field guide for the manager who sees the problems but cannot always get the budget to fix them.
- Built on SMRP best practices, ISO 55000, Doc Palmer's planning discipline and RCM fundamentals.
- Written in steps you can start on a Tuesday.

WHAT THIS IS NOT

- A software pitch. The product boxes show how a play runs in one system. The plays work on paper too, slower.
- A guide about "the future of maintenance".
- A list of talking points to justify hiring a consultant.

How to use it

NOT SURE WHERE TO START?

Read Part 2, the plumbing. Every play assumes it works.

KNOW YOUR WORST NUMBER?

Confirm it on page 25, then go to Part 4 for its play.

ALREADY RUNNING A PLAY?

Part 5 has the RCA guide and the sequence for the next one.

ONE THING BEFORE WE START

None of this matters if the line supervisor does not enforce the standard. No CMMS, no playbook, no benchmark fixes a plant where the people running the work do not believe in it. Fix the system first. Then back the people running it. That is the job.

Let's get into it.

The Maintenance Control Check-In

Rate each statement from 1 to 5. One means "not true of my plant", five means "true every week, and I could prove it from the system in ten minutes". Add the ten numbers. The next page tells you where to start.

#	STATEMENT	1 TO 5
01	Operators write the notifications. Not "I'll tell Mike". Every one is triaged within 24 hours.	
02	At least 80% of work is planned before it is scheduled: parts confirmed, job plan attached, hours estimated.	
03	Parts for planned jobs are kitted and staged 48 hours before execution.	
04	Next week's schedule is locked 48 to 72 hours before it starts, and breaking it needs a conversation.	
05	One named person owns schedule compliance as a number, and it is above 85%.	
06	PM compliance is above 90% and I know the top cause of every missed PM last month.	
07	Every closed work order has a failure code, a cause and an action. "Repaired, RTO" (return to operation) is sent back.	
08	An asset with three failures of the same code gets a formal RCA with an owner and a verified fix.	
09	I can pull PM compliance, reactive %, backlog weeks and MTTR for last month in under ten minutes.	
10	Reactive work is below 20% of labour hours, and the trend is down.	
Total (10 to 50)		

Score what is true, not what the org chart says should be true. The point is not to feel good. The point is to know.

Three scores, three starting points

10 to 24

Firefighting, and the fire is winning

Work gets found by failure, planned by nobody and closed with "RTO" (return to operation). The schedule is a list of hopes. Nothing in Part 4 will hold until Part 2 does. Read the five foundation chapters this week. Then run **Play 1 (PM compliance)** and **Play 4 (reactive work)**, in that order.

25 to 37

The PM program exists. Execution discipline does not.

You have a planner, a schedule and a CMMS. You also have break-ins every Monday and a backlog that became wallpaper. Start with **Play 5 (schedule compliance)** and **Play 3 (backlog)**, then **Play 2 (repeat failures)**. This is where most mid-size plants live.

38 to 50

You have earned the right to optimise

The cycle works. Now make the equipment tell you earlier. **Play 9 (downtime)**, **Play 10 (MTTR)** and **Play 8 (condition monitoring)** are yours. Hold the standard where it is inconvenient. That is where it counts.

Two readings of the same score

A 40 made of four 5s and six 3s is not the same plant as a 40 made of ten 4s. The first has a strong planning office and a close-out culture that is quietly broken. Look at your lowest two statements. Those are your first two plays, whatever the total says.

If statement 09 scored below 3, note it. A plant that cannot pull its own numbers in ten minutes has a data hygiene problem, and every other play will be measured by feel.

KEEP THE NUMBER

Write today's score and the date inside the front cover. Run one play for 90 days. Score again. A number that moved is the only proof the plant will believe.

Score 1: _____ on _____

Score 2: _____ on _____

Statement 09 scored low? Fabrico pulls PM compliance, downtime, MTTR and work order completion from the same records the technicians close.

[Request a demo](#)


PART 2

Foundations: the work management cycle

Most plants try to fix maintenance by buying software, hiring consultants or rolling out a new dashboard. None of it works if the cycle underneath is broken. Five stages, two pages each.


- p. 10 Why foundations matter
- p. 11 Work identification
- p. 13 Planning, and the ready-to-schedule gate
- p. 15 Scheduling
- p. 17 Work execution
- p. 19 Close-out and the feedback loop

Five stages. Very few plants have all five working together.



Most plants have something happening at every stage. When one stage breaks, the next four inherit the mess. A planning office that plans beautifully cannot recover from a scheduling function that ignores parts availability. A tight schedule cannot survive an execution culture that closes every work order with "Repaired. RTO."

This part is not glamorous. It is the plumbing. Everything in Parts 3, 4 and 5 assumes the plumbing works.



A status screen over a painting line: OEE, availability, line speed and the current stop, live.

Identification is the front door. If work does not get captured, nothing downstream matters.

Tuesday morning. An operator sees a small oil leak under a gearbox. He grabs a rag, wipes it down, does not mention it. Shift changes. The next operator sees the same leak, assumes it was already reported. Six weeks later the gearbox seizes mid-shift and takes out two days of production.

Nobody did anything wrong. The work just never got identified.


Most plants over-focus on the PM program and under-focus on identification. PMs catch what is scheduled. Everything else has to be captured by people who were not hired to capture it.

WHAT GOOD LOOKS LIKE

- **Operators write notifications.** Not maintenance, not the supervisor. The person closest to the equipment.
- **Triage within 24 hours** into three buckets: immediate action, planning queue, or reject with a reason.
- **Enough detail to act on.** Not "pump broken". "Pump 43 discharge pressure fluctuating 40 to 60 PSI at morning start-up, bearing whine from the coupling end."

COMMON FAILURES

- **The "I'll tell Mike" culture.** Operators report verbally to their favourite tech. Mike forgets. The notification never exists.
- **Notifications go to a queue nobody owns.** Fifty items, no triage, no reject reason.
- **The form takes 15 minutes.** That is not identification. That is a filter.



The Work Cards board. Everything new lands in the Unassigned column on the left until someone triages it.

Run it: the cadence

DAILY

The coordinator triages the new queue. Zero-inbox rule: every notification has a bucket by the end of the day.

WEEKLY

Operations and maintenance review the reject pile together. A reject without a reason goes back in.

MONTHLY

Pull the numbers. Low notification counts almost never mean a healthy plant. They mean the channel is broken.

FURTHER READING

SMRP Best Practices, Work Management pillar: work identification

Terry Wireman, Benchmarking Best Practices in Maintenance Management (work identification chapter)


ISO 14224:2016, Collection and exchange of reliability and maintenance data

HOW TO RUN THIS IN FABRICO

The notification is a photo and three taps at the machine

The front door has to be easier than telling Mike. In Fabrigo the operator scans the machine's QR code on a phone or tablet, picks the task type, adds a photo and a line of text. The machine, the time and the reporter are logged for them. The responsible role gets a notification, and the item appears in the Unassigned column for the coordinator to triage.

- ✓ Report from any device or by scanning the machine QR code
- ✓ Urgent tasks are flagged with a warning icon and scheduled for the same day by default
- ✓ Every report carries the machine, so the history builds by asset, not by memory



Task type on a phone: Urgent, One-off, Periodic, Stopped by operator.

[See it live in a demo](#)

A planner who is never allowed to plan is a firefighter in a different chair

Every plant has a Senior Maintenance Planner. His name is on the org chart. Everybody likes him. When a breakdown happens at 6 AM he is the first call. When the shutdown gets pushed he is the one who renegotiates with operations. Last month, across four weeks, he planned three jobs. The other six hundred went to schedule without him.

He is not a bad planner. He is a planner who is never allowed to plan.

A real planner does one thing. They take a piece of identified work and make it ready to execute: parts on hand, procedure written, hours estimated, safety reviewed. When a job leaves the planner's desk, a technician should be able to walk to the asset and start. Doc Palmer's research puts one real planner behind twenty to thirty technicians, and a planned job at roughly a third of the labour hours of the same job done unplanned.


WHAT GOOD LOOKS LIKE

- **80% of work planned before it is scheduled.**
Not a goal. A gate.
- **Every planned job has a job plan:** steps, parts list, hour estimate, special tools, safety, drawings where needed.
- **Parts kitted and staged 48 hours before execution.** This one practice kills the number one cause of schedule breakdown.
- **Hour estimates land within 20% of actual** on routine work. Off by 50% every time is not an estimate. It is a wish.

COMMON FAILURES

- **The planner is actually a clerk.** Opens the WO, maybe pulls parts, never builds a real plan.
- **Jobs released without parts.** Every week. Nobody flags it.
- **Job plans are three words** in the description field.
- **The planner gets dragged into every breakdown** because he knows the assets best. The queue never gets touched.

PALMER'S RULE • LABOUR HOURS FOR THE SAME JOB



The ready-to-schedule gate

Before a work order can be scheduled it must have all five. If any is missing, it stays in planning. Hold this one rule and the most common reason a schedule breaks, a job released without parts, goes away.

- ☐ **Parts confirmed available.** In the bin, not on order.
- ☐ **Job plan attached.** Steps, tools, safety requirements, drawings.
- ☐ **Hour estimate and craft.** Who, and for how long.
- ☐ **Asset access window confirmed** with production.
- ☐ **Safety requirements documented.** LOTO on the plan, not improvised on the floor.

DAILY

30 minutes, first thing: the planner processes the notification queue. Nothing else until it is empty.

WEEKLY

Planning meeting with the supervisor and the parts lead. Jobs without parts go back, not forward.

MONTHLY

Audit planned against actual hours. Read five job plans; three-word plans get rewritten with the tech.

FURTHER READING


Doc Palmer, Maintenance Planning and Scheduling Handbook (4th ed.), the bible
SMRP metric 5.3.1, Planned Work; 5.5.2, Craft Worker to Planner Ratio
ISO 55001:2014 Clause 8.1, Operational planning and control

HOW TO RUN THIS IN FABRICO

The gate lives on the task card

Each Fabrico task card carries the machine, the scheduled date, the estimated duration and three sections: files, used materials, performers. A card whose materials section is empty or whose estimate is missing is visibly not ready.

- ☒ Estimated and reported duration side by side on every card, exportable, so the 20% test is one column
- ☒ SOPs, checklists and drawings attached once, opened at the machine



A periodic task card with its estimate, files, materials and performers. [See it live in a demo](#)

This is not a scheduling problem. It is a commitment problem.

The Friday scheduling meeting runs forty minutes over every week. Maintenance reads off next week's schedule. Operations agrees to everything. Monday at 7 AM, production calls about a priority changeover on Line 3 that nobody mentioned on Friday. Two planned jobs get bumped. By Tuesday afternoon the schedule is a fiction nobody references.


A real schedule is a contract. We will make the equipment available in these windows; you will have the right crews with the right parts to do the right work. Both sides sign. Breaking the contract should require a conversation, not a text at 7 AM. The target in Part 3 is 85%. Plants without the contract sit well below it. The gap is not craft. It is discipline.

WHAT GOOD LOOKS LIKE

- **The schedule is locked 48 to 72 hours before the week starts.** This is the frozen zone. No changes unless safety-critical.
- **Operations, maintenance and the storeroom in the same room** for the weekly meeting. No parts, the WO goes back to planning.
- **A two-week look-ahead is published.** Everybody can see what is coming.
- **One person owns schedule compliance as a number.** Named. Reviewed weekly.

COMMON FAILURES

- **The "emergency" that is not.** A changeover known for three weeks shows up as urgent on Monday.
- **A schedule built without confirming labour.** Ten on paper, five on Monday.
- **No frozen zone.** The schedule becomes a list of things maintenance hopes to do.



Month view: red is overdue, green done, orange in progress, grey unplanned.

Run it: the cadence

DAILY

Ten-minute morning huddle.
What broke in, who authorised it, what it displaced.

WEEKLY

Scheduling meeting, 30 minutes. Lock next week.
Review last week's compliance by name, not by department.

MONTHLY

Review schedule breaks by category: failure, management request, parts, absenteeism.
Fix the top one. Only the top one.

FURTHER READING


Doc Palmer, Maintenance Planning and Scheduling Handbook (3rd ed.), weekly scheduling discipline
SMRP metrics 5.4.3 and 5.4.4, Schedule Compliance (hours, work orders)

HOW TO RUN THIS IN FABRIGO

The week is a board, and every person has eight hours on it

On the Work Cards board, drag tasks onto people. Each column shows how many of the day's eight hours are already committed, so overloading a technician is visible before Monday, not after. Locking the week means the board for next week is full by Thursday and nothing moves onto it without a conversation.

- ✓ Drag and drop assignment, the technician is notified on their phone
- ✓ Busy hours per person per day, against the shift
- ✓ The calendar shows missed and not-yet-planned tasks for any week or month



Hours already committed per technician, in each column header. [See it live in a demo](#)

Execution is where the program meets reality. The discipline is capturing what actually happened.

A senior tech opens a centrifugal pump scheduled for a seal replacement. The seal is shot, as expected. But the coupling is also gone, the bearing housing has scoring, and the shaft has a slight bend. The job plan said two hours. He is now looking at sixteen. He does not call planning. He pulls parts from the storeroom himself, finishes at 9 PM, and the WO closes Monday with "Repaired. RTO."

The work got done. And the data is worthless, the storeroom is short a coupling nobody logged, and nobody will know how often the bearing fails too.

WHAT GOOD LOOKS LIKE

- **Safety and LOTO pre-staged.** Lockout procedure on the job plan, locks and tags in the kit. Nobody works out LOTO on the fly.
- **Scope changes trigger a call.** Tech finds more damage? Radio, planner updates the WO, storeroom pulls more parts.
- **Hand-offs have content.** What is done, what is left, what was found.
- **As-found and as-left recorded.** "Coupling worn 80%, shaft TIR 0.012. New coupling, aligned to 0.002." That is what makes an RCA possible six months later.

COMMON FAILURES

- **Scope creep that never reaches the CMMS.** Five things fixed, one documented.
- **LOTO as improv.** Mostly works. Until it does not.
- **The silent hand-off.** The next crew spends an hour rediscovering what the last shift found.
- **"Repaired. RTO."** That is not data. That is a receipt.

EXAMPLE • WHAT THE PLAN SAID, WHAT HAPPENED, WHAT THE RECORD SAYS

PLAN	2 H • SEAL
ACTUAL	16 H • SEAL, COUPLING, BEARING HOUSING, SHAFT, PARTS PULLED UNLOGGED
	RECORD: "REPAIRED. RTO." • FINDINGS CAPTURED: 1 OF 5

Four findings never reached the record: the worn coupling, the scored housing, the bent shaft, the coupling taken from stock. Next month the storeroom is short a part nobody logged, and the next time this pump family whines nobody will know the bearing fails too.

Run it: the cadence

PER JOB

Safety briefing before starting.
Status check at the midpoint on jobs over four hours. A scope change triggers an update the same day, not at close.

DAILY

End-of-shift debrief, five minutes. What was found that the plan did not expect.

WEEKLY

The supervisor pulls five random completed WOs and reads them. Anything closed with "RTO" as the only note goes back to the tech.

FURTHER READING

Ron Moore, Making Common Sense Common Practice (work execution discipline)

SMRP metric 5.4.2, Proactive Work


OSHA 1910.147, The Control of Hazardous Energy (Lockout/Tagout)

HOW TO RUN THIS IN FABRIGO

The record is written at the machine, while the covers are off

The tech works from a phone. The card opens with the procedure and the parts list; photos and notes go on during the repair, not from memory on Monday. The parts used are written off by the tech who used them, so the storeroom count moves the moment the coupling leaves the bin.

- ✓ Photos, notes and files on the card during or after the repair
- ✓ Spare parts written off by the technician, QR scan on the part
- ✓ Labour, parts and failure type logged at the asset, an auditable history



A technician's card list on a phone, with the day's estimates. [See it live in a demo](#)

The equipment is telling you something. The CMMS is not listening.

Six months after the pump job, the same pump fails again. Same symptoms. The reliability engineer pulls the history. Three prior WOs, each closed with a variation of "Repaired. RTO." No failure codes. No as-found or as-left notes. No photos. Nothing.


Close-out is the most neglected stage because it feels like paperwork. It is not. It is the difference between a plant that learns and a plant that keeps relearning the same lesson. Most CMMS implementations have the right fields. What they lack is enforcement. Techs skip them, supervisors sign anyway, and six months later the data is garbage nobody wants to admit is garbage.

WHAT GOOD LOOKS LIKE

- **Every close-out has a failure code, a cause and an action.** Three fields. Non-negotiable.
- **Actual labour hours captured** and compared to the estimate.
- **Photos and notes on anything non-routine.**
- **A monthly audit of ten random closed WOs.** More than two incoherent, the discipline is broken.

COMMON FAILURES

- **Failure codes that do not match the failure.** Tech picks "Other" because the dropdown is 200 options deep.
- **Round-number labour hours only.** Suspiciously tidy. Almost certainly invented.
- **No feedback from close-out to planning.** Captured and ignored. The worst failure mode: it creates the illusion of a learning system.



Status	Nº Task	Name
🔧	7268	проверка лагер
🔧	7253	проверка масло
🔧	7252	Пълно обслужване
🔧	7250	Пълно обслужване
🔧	7248	обслужване
⚠️	6829	Авария с помпата
🔧	6410.0009	източване на филтър 1
🔧	6311.0021	Преглед на комуникацион
🔧	6311.0008	Преглед на комуникацион

Repair history by task, one machine.

HOW TO READ IT ON MONDAY

01 Look for the third occurrence

Three closed cards on one machine with the same symptom is not bad luck. It is the trigger for a formal RCA (Play 2).

02 Read the type icon first

Urgent, one-off, periodic or stopped by operator shows which work was planned and which was a reaction. A run of urgent cards on one asset points at a gap in its PM.

03 Check what the close-out says

A card closed with "Repaired. RTO." gives the next review nothing to read. Send it back and the history starts to mean something.

Run it: the cadence

PER JOB

Close out at completion, not at the end of the week. The tech who did the work writes the note.

WEEKLY

The reliability engineer reads every closed WO on critical assets. Looks for the third occurrence of anything.

MONTHLY

PM adjustments from actual labour and failure data. Close-out quality score published on the board, by crew.

FURTHER READING


John Moubray, Reliability-Centered Maintenance II, the feedback loop
SMRP Best Practices, Equipment Reliability pillar: defect elimination
ISO 14224:2016, Failure data collection standards

HOW TO RUN THIS IN FABRIGO

The feedback loop is a report you open every Monday

Repair history by machine and by task, downtime by machine, the mean recovery time after failure, and the "frequent offenders" view that ranks machines by breakdown count and duration. The data is only there because the tech closed the card at the machine with the type, the parts and the photos. That is the loop: capture at close, read on Monday, change the PM.

- ✓ Monday: open the frequent offenders, take the machine with the third occurrence, change its PM task or frequency the same day
- ✓ Month end: export the closed cards, read ten at random, publish the close-out score by crew



Analytics: total downtime for the period, split by task type. [See it live in a demo](#)

Want the five stages in one system? Fabrigo is a cloud-based MES and OEE platform for manufacturers, with maintenance management (CMMS) built in.

[Request a demo](#)

PART 3

Where are you


A maturity model, a five-question self-assessment and a benchmark table. They tell you honestly where you sit, not where last year's audit said you sat.

- p. 22 The maturity model
- p. 23 The self-assessment
- p. 25 The benchmark snapshot
- p. 26 Where the numbers come from

Five levels. Every plant sits at one. Most do not know which.

Walk into Plant A at 7 AM. The whiteboard is full of work orders stamped URGENT. The planner is on the phone with a vendor. Walk into Plant B at 7 AM. A technician is taking a vibration reading on a critical pump. The planner is reviewing next week's kitted jobs.

Same industry. Same equipment. Different maturity level.



The bands are working bands, built from the SMRP targets on page 25. Use them to place yourself, not to argue.

MOVING UP A LEVEL

You cannot skip levels. A Level 1 plant cannot leapfrog to Level 3 by buying vibration analysers. Condition monitoring on top of broken PM compliance produces nothing but expensive alerts nobody acts on. Each level is built on the one below. The plays in Part 4 are tagged by level. Start with the ones that match where you actually are.

Five questions. Answer honestly.

Circle one answer in each block. The lowest level you circle is the plant's level.

Q1 • PM COMPLIANCE

What percent of scheduled PMs are completed on time, measured weekly?

Below 50%	Level 1
50 to 70%	Level 1 or 2
70 to 90%	Level 2
90 to 95%	Level 3
Above 95%	Level 4 potential

Q2 • REACTIVE WORK

What share of labour hours is reactive: breakdowns, emergencies, break-ins?

Above 50%	Level 1
30 to 50%	Level 1 or 2
20 to 30%	Level 2
10 to 20%	Level 3
Below 10%	Level 4 potential

Q3 • BACKLOG MANAGEMENT

How is the backlog tracked, and how many weeks of planned work is it?

No formal backlog	Level 1
Tracked but growing	Level 1 or 2
Tracked, 4 to 6 weeks, stable	Level 2
Tracked, 2 to 4 weeks, actively purged	Level 3+

Q4 • CONDITION MONITORING

Which critical assets are monitored by condition, and where does the data go?

None	Level 1 or 2
Informal or ad hoc	Level 2
On most critical assets, in the CMMS	Level 3
With trending and predictive models	Level 4

"We send oil samples once a year" is not a program. It is a gesture.

Q5 • ROOT CAUSE ANALYSIS


What happens after the second failure of the same asset with the same code?

Same repair as last time, nobody asks why	Level 1 or 2
The supervisor looks at it informally	Level 2
Formal RCA (5 Whys, Fishbone, FMEA)	Level 3
RCAs tracked, closed, corrective actions verified	Level 3+

Take the lowest level across the five. That is where you are.

Not the average. Maturity is bounded by your weakest stage. A plant with great PM compliance but no RCA process is not a Level 3 plant. It is a Level 2 plant with an unusually strong PM program.

EXAMPLE PLANT • FIVE ANSWERS, ONE LEVEL



Where to go from here

SCORED LEVEL 1

Play 1 (PM compliance), then **Play 4 (reactive work)** from month 2. The foundation plays. Nothing else holds until they do.

SCORED LEVEL 2

Play 5 (schedule compliance), then **Play 3 (backlog)**, then **Play 2 (repeat failures)**. The PM program exists; the discipline around it is what moves next.

SCORED LEVEL 3

Play 9 (downtime), then **Play 10 (MTTR)**, then **Play 8 (condition monitoring)**. The same order as the sequence on page 52.

THE TWO TESTS, SIDE BY SIDE

The check-in on page 7 scores your **discipline**: what the plant actually does every week. This assessment scores your **maturity**: what the plant is capable of. A high check-in with a low self-assessed level usually means one weak stage, backlog or condition monitoring, is holding the level down; fix that stage first. A low check-in with a high self-assessed level is a plant that is lying to itself.

Eight KPIs every mid-size plant should track weekly

These are the industry standards. Not aspirations. The numbers that separate a well-run program from a struggling one. Pull your own for the last 30 days and mark each red, yellow or green.

KPI	UNIT	ACTION	TARGET	WORLD CLASS	YOURS
PM compliance	%	<85	90	95	
Schedule compliance	%	<75	85	90	
On-time completion	%	<80	90	95	
Reactive work	%	>30	20	10	
Backlog	weeks	>8	4 to 6	2 to 3	
Inspection compliance	%	<85	90	95	
Availability	%	<90	95	97	
MTTR	hours	>8	<4	<2	

Red: past the action threshold. Yellow: between action and target. Green: at target or better. Targets follow the SMRP Best Practices metrics named on each play in Part 4; the action and world-class bands are working bands, from practice. ISO 55000 and SAE JA1012 govern the method, not the numbers.

How to read the colour pattern

Three reds is a Level 1 plant pretending to be Level 2. All green is a real Level 3 plant, and there are not many. The pattern tells you where to focus faster than any single number.

A WORD ON MEASUREMENT

Cannot pull these eight in ten minutes? That is statement 09 on page 7, and it comes first. Start with PM compliance and reactive work. The others follow.

Eight definitions, so two people never argue about the same number

Same period, same scope, every time. Where a result is a percentage, multiply by 100.

PM COMPLIANCE

%

$$= \frac{\text{PMs completed on schedule}}{\text{PMs due}} \times 100$$

SCHEDULE COMPLIANCE

%

$$= \frac{\text{WOs executed on the day scheduled}}{\text{WOs scheduled that day}} \times 100$$

REACTIVE WORK

%

$$= \frac{\text{Reactive labour hours (breakdowns, emergencies, break-ins)}}{\text{Total maintenance labour hours}} \times 100$$

The total includes PM and planned hours.

BACKLOG

WEEKS

$$= \frac{\text{Open planned hours}}{\text{Available weekly crew hours}}$$

AVAILABILITY

%

$$= \frac{\text{Total time minus unplanned downtime}}{\text{Total time}} \times 100$$

MTTR

HOURS

$$= \frac{\text{Total repair time on unplanned outages}}{\text{Number of unplanned outages}}$$

ON-TIME COMPLETION

%

$$= \frac{\text{WOs closed by their due date}}{\text{WOs due in the period}} \times 100$$

INSPECTION COMPLIANCE

%

$$= \frac{\text{Inspections completed on schedule}}{\text{Inspections due}} \times 100$$

HOW TO RUN THIS IN FABRICO

Most of the numbers are a by-product of closing the card properly

Six of these eight fall out of the records the technicians close: downtime per machine, mean recovery time after failure, task completion by type and status, PM and inspection compliance. Backlog in weeks still needs the planner's own count of hours. OEE availability per line comes from the machine signal, in the same system.

- ✓ Task distribution by type and status, so the reactive share is one chart
- ✓ OEE availability, performance and quality live per line and shift, from the machine signal

PART 4

The ten playbooks

Ten specific plays for ten specific KPI failures. Each has a trigger, a 90-day roadmap, a benchmark with its source, and the screen to run it from. Pick the one that matches your worst number right now.

- p. 28 PM compliance recovery · Level 2
- p. 30 Repeat failure investigation · Level 3
- p. 32 Backlog growth control · Level 2
- p. 34 Reactive work reduction · Level 2
- p. 36 Schedule compliance improvement · Level 2
- p. 38 Inspection compliance recovery · Level 2
- p. 40 Spare parts management · Level 3
- p. 42 Condition-based monitoring · Level 3
- p. 44 Downtime reduction · Level 2
- p. 46 Asset reliability and MTTR · Level 3

MATURITY LEVEL 2


1. PM compliance recovery

TRIGGER PM compliance below 85% for two or more consecutive periods.

Monday morning. The supervisor opens the PM dashboard. Seventy-three percent. Same as last week. Same as the week before. He sends the number up with a brief note: "working on it". Nothing changes. Two weeks later it is at 68. Then 65.

PM compliance does not collapse all at once. It erodes one missed job at a time.

EXAMPLE • PM COMPLIANCE, SIX WEEKS



Why it matters

A parts delay here, a break-in there. Every individual decision feels reasonable. The pattern is not. Each unexecuted PM lets a failure mode develop silently, and the asset comes back and bites you on a Tuesday afternoon six weeks later, as corrective work that costs several times the PM it replaced. The action line sits at 85%, five points under the SMRP target, because below it the PM program has stopped protecting the plant and started documenting its decline.

PM COMPLIANCE, %

ACTION • TARGET • WORLD CLASS

Below 85

Above 90

Above 95

Source: SMRP metric 5.4.14, PM and PdM Work Order Compliance (target); the action and world-class bands are working bands

THE FOUR MOVES

- ☐ Triage and clear the backlog by asset criticality
- ☐ PM kitting 48 hours before execution
- ☐ A frozen zone protecting PMs from displacement
- ☐ An automated reminder and escalation workflow

FURTHER READING

SAE JA1011, Evaluation Criteria for RCM Processes
ISO 55001:2014, Asset management systems, Requirements

The 90-day roadmap


Days 1 to 3	Triage the PM backlog. Export overdue and upcoming. Categorise by asset criticality (A, B, C). Critical assets first.
Days 3 to 5	Root-cause the delays. Interview planners and technicians. The top causes are always parts, scheduling, labour or procedures.
Week 2	PM kitting and staging. Pre-kit all parts 48 hours before the scheduled date.
Weeks 2 to 3	Protect PM windows. Frozen zone inside 48 hours. No break-in displaces a due PM unless safety-critical.
Week 3	Reminders and escalation. Three days before, day of, plus one day to the supervisor.
Weeks 4 to 6	Review PM frequencies and procedures. Remove tasks that do not address a failure mode.
Ongoing	Weekly reporting. Target 90% in 8 weeks, 95% in 16.

Fabrico HOW TO RUN THIS IN FABRICO

The PM is a periodic task with a frequency and a tolerance

Every PM in Fabrico is a repeatable task: first date, frequency, period type, and a tolerance in days. Due tasks appear on the calendar and on the technician's own card; overdue ones turn red and stay red until closed. Reminders go to the role, and unresolved urgent tasks escalate to the next level. The annual preventive maintenance plan is generated from the recurring tasks with one click.

- ✓ Recurring tasks by calendar, run hours or cycle counts
- ✓ Annual PM plan per machine and month, generated from the recurring tasks
- ✓ PM compliance rate visible across plants and teams



Repeatable settings

First date *	Group end date	Tolerance (in Days)	Period type
03/01/2024	03/02/2024	1	week
Start of periodic tasks	Will not appear in annual plans after this date		
Frequency (in weeks) *	Day of the week		
1	Tuesday		

Repeat settings for a periodic task: first date, tolerance, frequency. [See it live in a demo](#)

MATURITY LEVEL 3

2. Repeat failure investigation

TRIGGER The same asset has three or more work orders with the same failure code in the period.

A centrifugal pump on Line 3 has failed four times this year. Each time: same symptom, same diagnosis, same repair. Replace the mechanical seal, return to operation, close the WO with "Repaired. RTO." On the fifth failure the reliability engineer finally pulls the history and discovers the asset has eaten seven seals in eighteen months.

Nobody had asked why until now.

EXAMPLE • ONE PUMP, 18 MONTHS, SEVEN SEALS



Why it matters

Repeat failures are the loudest signal the equipment sends you. They tell you the repair is treating symptoms, not causes, and the fix-fail-fix cycle consumes labour the plant could have spent once on a permanent fix. A plant above the action threshold has a reliability culture problem, not a craft problem. The techs are not bad at repairing. The system is bad at asking questions.

REPEAT FAILURE RATE, %

ACTION • TARGET • WORLD CLASS

Above 15

Below 10

Below 5

Source: working bands from practice; SMRP publishes no repeat-failure metric

THE FOUR MOVES

- ☐ Rank repeat combinations by frequency times consequence
- ☐ Run a formal RCA: 5 Whys, Fishbone or FMEA
- ☐ Update PM procedures from the root cause findings
- ☐ Install condition monitoring on the highest-consequence assets

FURTHER READING

SAE JA1012, Guide to the RCM Standard
IEC 60812, Failure modes and effects analysis

The 90-day roadmap


- Days 1 to 2** **Pull WO history by asset and failure code.** Rank by frequency times consequence. The top three to five get the focus.
- Days 3 to 10** **Formal RCA on each.** 5 Whys for the simple ones, Fishbone for multi-factor, FMEA for critical assets (Part 5).
- Week 2** **Corrective action per root cause.** An owner and a deadline. PM update, design change, training, or condition monitoring.
- Weeks 2 to 3** **Update PM procedures.** If the failure is detectable by inspection, add or change the PM task.
- Weeks 3 to 6** **Condition monitoring** on the highest-consequence repeat failures.
- Weeks 6 to 12** **Verify and close.** Monitor for two to three MTBF intervals. A recurrence escalates to engineering.

HOW TO RUN THIS IN FABRICO

The frequent-offenders list is a report, not a hunch

The repair history by machine lists every closed card on the asset with its type, number and date. The analytics rank machines by breakdown count and total duration. When the third card of the same kind shows up, the RCA starts from the record, not from memory.

- ✓ Repair history by machine and by task, exportable
- ✓ Top machines by breakdown frequency and duration
- ✓ Photos and notes on every card, the as-found evidence an RCA needs



Status	Nº Task	Name
	7268	проверка лагер
	7253	проверка масло
	7252	Пълно обслужване

Repair history by task on one machine.

[See it live in a demo](#)

Line	Stoppage Reason	Type	Group	Start Date	End Date	Duration	
1 Line S - Packaging & Filling (LN-PF)	Lack of material	■ Unplanned	operational	29 Jul 15:05	29 Jul 15:08	3m	⋮
1 Line S - Packaging & Filling (LN-PF)	packet setting	■ Planned	cleaning	29 Jul 14:53	29 Jul 14:58	5m	⋮
1 Line S - Packaging & Filling (LN-PF)	packet setting	■ Planned	cleaning	29 Jul 14:22	29 Jul 14:26	3m	⋮
1 Line S - Packaging & Filling (LN-PF)	Cleaning after spill	■ Unplanned		29 Jul 14:05	29 Jul 14:10	5m	⋮

The stoppage list for one line and one afternoon: "packet setting" twice in an hour, "maintenance" twice.

MATURITY LEVEL 2


3. Backlog growth control

TRIGGER The backlog has been growing for three or more consecutive periods.

The Friday scheduling meeting. Someone reads out "backlog is now at 9.2 weeks". Nobody responds. Last week it was 8.8. Three months ago it was 6.4. The number has been climbing for twelve straight weeks.

It has become wallpaper.

EXAMPLE • BACKLOG IN WEEKS OF PLANNED WORK, TWELVE WEEKS



Why it matters

A healthy backlog is four to six weeks of planned work; world class runs two to three. It is a buffer, not a graveyard. When it grows week over week, the plant is creating work faster than it can complete it, and that is a leading indicator of what your reactive percentage will look like in three months. Above six weeks, urgent work cannot find a clean slot. By eight weeks routine work takes two months, and technicians stop documenting small defects because the WO will sit there anyway.

BACKLOG, WEEKS OF PLANNED WORK

ACTION • TARGET • WORLD CLASS

Above 8

4 to 6

2 to 3

Source: SMRP metrics 5.4.9 and 5.4.10, Ready and Total Backlog; working bands from practice

THE FOUR MOVES

- ☐ Conduct a backlog purge to remove stale items
- ☐ Categorise and prioritise: must, should, could
- ☐ Evaluate staffing against the current workload
- ☐ Implement a weekly backlog review cadence

FURTHER READING

Doc Palmer, Maintenance Planning and Scheduling Handbook (4th ed.), backlog management
 ISO 55001:2014, Asset management systems, Requirements

The 90-day roadmap


- Days 1 to 3** **Backlog review and purge.** Close or cancel stale WOs: duplicates, already fixed, superseded. Expect a real share of the list to go.
- Days 3 to 5** **Categorise and prioritise.** Must-do (safety, compliance), should-do (reliability), could-do (nice to have).
- Week 1** **Staffing against workload.** Backlog hours versus available crew hours. Above six weeks and growing: contractor support or overtime, with an end date.
- Week 2** **Weekly backlog review cadence.** 30 minutes with the maintenance leads. Weekly completion targets.
- Weeks 2 to 4** **Root-cause the growth.** WO creation climbing? Completion falling? Parts delays? Target the primary driver, only that one.
- Weeks 5 to 12** **Hold.** Weekly review, weekly completion target, a second purge at week 8. The number stable for four weeks counts as a win.

HOW TO RUN THIS IN FABRICO

The backlog is a month of red dots, and it has a length

Every open card that is not on a person is in one column. The month view shows which weeks are overloaded and which tasks were missed or never planned. For the purge, filter the board by performer, role and location and the calendar by period, and close the duplicates from the card itself. For the staffing check, each person's column already carries the hours committed that day.

- ✓ Missed and not-yet-planned tasks visible by week or month
- ✓ Hours committed per person per day



The month view: red is overdue, and a month of red is a backlog you can count.

[See it live in a demo](#)

MATURITY LEVEL 2

4. Reactive work reduction

TRIGGER Reactive work exceeds 30% of maintenance labour hours.

The shift started with eight planned jobs. At 7:30 a gearbox fails on Line 2. Two techs pulled. At 9:15 a conveyor in packaging seizes. Two more pulled. At 11 the plant manager calls about a critical pump warning on the boiler. The shift ends with two planned jobs done and six displaced.


Some plants run this way for years and call it "being responsive". What it actually is: an expensive way to lose money.

EXAMPLE • ONE SHIFT, EIGHT PLANNED JOBS



Why it matters

Reactive work costs several times what the same job costs planned (Palmer's one-third rule, read in reverse), and the displacement cascades: the bumped jobs become next week's breakdowns. Every hour spent reacting is an hour not spent preventing. The loop does not break by working harder. It breaks by working differently.



Source: SMRP metric 5.4.1, Reactive Work; working bands from practice

THE FOUR MOVES

- ☐ Pareto the top ten failure modes driving unplanned work
- ☐ Map failure modes to PM coverage and close the gaps
- ☐ Install condition monitoring on critical equipment
- ☐ Implement operator care and anomaly reporting

FURTHER READING

SAE JA1011, RCM Evaluation Criteria
Ron Moore, Making Common Sense Common Practice

The 90-day roadmap

Days 1 to 3 **Pareto the top ten failure modes** driving unplanned work.

Weeks 1 to 2 **Map failure modes to PM coverage.** No PM? Create one. PM failing to prevent? Revise the task or the frequency.

Weeks 2 to 4 **Operator-driven reliability.** Train operators on inspection, cleaning, lubrication and anomaly reporting.

Week 3 **Review the emergency priority criteria.** Most plants over-classify routine work as urgent. Tighten it.

Weeks 4 to 8 **Condition monitoring** on high-value assets with repeat failures.


Ongoing **Weekly planned versus reactive ratio.** Target a 5-point reduction per month until below 20.

HOW TO RUN THIS IN FABRICO

Reactive share is one chart, and the Pareto is the stoppage list

The analytics split tasks by type: urgent, one-off, periodic, stopped by operator. Reactive work is the labour hours on unplanned jobs (breakdowns, emergencies, break-ins) over total maintenance labour hours. Reported hours on urgent tasks over reported hours on all tasks gives that percentage from the closed cards. For the Pareto, open the stoppage list, where every stop carries the reason the operator tagged, the machine and the duration, then sort it by duration. Operators report anomalies from the same app they tag stops in, so operator-driven reliability has a front door.

- ✓ Task distribution by type and status for any period
- ✓ Every stop with its reason, type, group, start, end and duration
- ✓ Anomaly reporting from any device or the machine QR code



Line	Stoppage Reason	Type	Group	Start Date	End Date	Duration	
1 Line S - Packaging & Filling (LN-PF)	Lack of material	Unplanned	operational	29 Jul 15:05	29 Jul 15:08	3m	⋮
1 Line S - Packaging & Filling (LN-PF)	packet setting	Planned	cleaning	29 Jul 14:53	29 Jul 14:58	5m	⋮
1 Line S - Packaging & Filling (LN-PF)	packet setting	Planned	cleaning	29 Jul 14:22	29 Jul 14:26	3m	⋮
1 Line S - Packaging & Filling (LN-PF)	Cleaning after spill	Unplanned		29 Jul 14:05	29 Jul 14:10	5m	⋮
1 Line S - Packaging & Filling (LN-PF)	At shift start	Planned	cleaning	29 Jul 13:49	29 Jul 13:54	5m	⋮
1 Line S - Packaging & Filling (LN-PF)	maintenance	Planned	maintenance	29 Jul 13:34	29 Jul 13:38	3m	⋮
1 Line S - Packaging & Filling (LN-PF)	10 min break	Planned		29 Jul 13:01	29 Jul 13:06	5m	⋮
1 Line S - Packaging & Filling (LN-PF)	Start breaking	Planned		29 Jul 12:46	29 Jul 12:50	3m	⋮

Stoppages with reason, type and duration per line. [See it live in a demo](#)

MATURITY LEVEL 2


5. Schedule compliance improvement

TRIGGER Schedule compliance below 75%.

The weekly scheduling meeting locks next week at 60 hours. Monday, production asks to move a changeover up two days. Tuesday, two techs call out. Wednesday, operations swaps Line 3 priorities. Thursday, half the kitted parts did not arrive. Friday, 32 of 60 hours got executed. Everyone shrugs and rebuilds for next week, the same way.


Below 75%, the schedule exists and nobody is defending it.

EXAMPLE • HOURS LOCKED ON FRIDAY VERSUS HOURS EXECUTED



Why it matters

Below the threshold, technicians context-switch, wait for parts and re-plan, and labour productivity falls with every break. Schedule breaks drive backlog growth, which drives reactive work, which drives more schedule breaks. The loop feeds itself until someone names an owner for the number.



Source: SMRP metrics 5.4.3 and 5.4.4, Schedule Compliance; the 75 action line is a working band

THE FOUR MOVES

- ☐ Pareto the break-in causes and prevent the top one
- ☐ Install a 48 to 72 hour frozen schedule window
- ☐ Enforce the parts-available rule before scheduling
- ☐ Implement a gatekeeper for emergency reclassifications

FURTHER READING

Doc Palmer, Maintenance Planning and Scheduling Handbook (4th ed.), weekly scheduling

The 90-day roadmap


- Days 1 to 3** **Analyse break-in patterns.** Pareto by cause: failure, management request, parts, absenteeism.
- Weeks 1 to 2** **Frozen schedule window.** Lock 48 to 72 hours before the work week. No changes inside the zone unless safety-critical.
- Week 2** **Strengthen the weekly meeting.** Operations, maintenance, storeroom in the same room. Lock the week.
- Weeks 2 to 3** **Parts and material planning.** Kit and stage before the work day. No parts, no schedule.
- Weeks 3 to 4** **Tighten emergency criteria.** A gatekeeper role. Only the supervisor authorises a break.
- Weeks 5 to 12** **Hold.** Compliance counted every Friday by name. Target above 85% by day 90.

HOW TO RUN THIS IN FABRICO

The locked week is the board nobody edits after Thursday

The week's cards sit on named people with their hours. A break-in is a new urgent card, flagged, that displaces something visible. On Friday, count what moved. The work schedules define who is on shift, so a schedule built on ten people when five are rostered shows up before Monday, not on it.

- ✓ Work schedules per team, so capacity is real
- ✓ Urgent cards flagged and scheduled same day, visibly displacing planned work
- ✓ Weekly and monthly task views for the compliance count



The week as a board: cards on people, hours per day in each header. [See it live in a demo](#)

MATURITY LEVEL 2

6. Inspection compliance recovery

TRIGGER Inspection overdue rate above 15%.

The external inspector shows up on Tuesday for the quarterly audit. The plant manager sits in, confident. Twenty minutes in, the inspector asks for the pressure vessel inspection records. The last API 510 inspection was 62 days past due. The one before that, 45. The inspector notes it in the report.


The room gets very quiet.

EXAMPLE • API 510 INSPECTIONS, DAYS PAST DUE



Why it matters

Inspections are the earliest detection opportunity. Each missed one extends the window in which a developing failure goes unseen, and the misses compound: a small corrosion defect missed at month six is a large defect at month nine and a through-wall leak at month twelve. For regulatory inspections the exposure is not internal. It is legal.



Source: working bands, aligned with SMRP metric 5.4.14; API 510, 570 and 653 set the regulatory intervals

THE FOUR MOVES

- ☐ Regulatory and safety-critical inspections first
- ☐ Mobile inspection checklists
- ☐ An automated reminder and escalation workflow
- ☐ Audit inspection frequencies against actual condition data

FURTHER READING

API 510, 570 and 653, inspection codes for vessels, piping and tanks
ISO 55001:2014, Asset management systems, Requirements

The 90-day roadmap


Days 1 to 3	Prioritise overdue inspections. Regulatory, then safety-critical, then condition-based. Clear regulatory first.
Days 3 to 5	Find the bottleneck. Conflicts with PMs? Inspector capacity? Narrow access windows?
Week 1	Reminders and escalation on the regulatory clock. Three days before, day of, plus one to the supervisor, plus two to the manager.
Weeks 2 to 3	Mobile inspection checklists. A tablet checklist at the asset, with the evidence the auditor asks for captured on the spot.
Weeks 3 to 4	Review inspection frequencies against the API codes and the condition data. Some are too frequent (waste). Others too infrequent (risk).
Weeks 5 to 12	Hold. Zero regulatory inspections overdue; overdue rate below 10% by day 90.

HOW TO RUN THIS IN FABRICO


The inspection is a periodic task with the checklist on the phone

An inspection in Fabrico is a recurring task on the asset, with the checklist attached and a tolerance in days. The inspector scans the vessel's QR code, the checklist opens, photos go on the card, the close-out is time-stamped. Notification rules decide who hears about an overdue one, by role and channel.

- ✓ Checklists, SOPs and drawings on the task, opened by QR at the asset
- ✓ Notifications by role and channel: e-mail, app, notification list
- ✓ Time-stamped close-out and export for the auditor



Notification settings for a role: an urgent task reaches e-mail, the notification list or the mobile app. [See it live in a demo](#)



Two weeks of the calendar: the red dots that stay are overdue tasks, inspections included. The grey ones were never planned.

MATURITY LEVEL 3

7. Spare parts management

TRIGGER Average WO completion time exceeds twice the estimate, with three or more completed WOs affected.

The storeroom has 42,000 line items. The maintenance manager can recite this fact. What he cannot do is explain why, on a given Tuesday, a technician is waiting three days for a specific bearing on a specific pump. The bearing exists in the inventory system. Nobody knows if any are actually in the bin until someone walks over and looks.


Parts delays are one of the most chronic causes of MTTR inflation. And it is rarely about spending too little on inventory. Most plants have the wrong inventory, kept badly.

EXAMPLE • THE STOREROOM IN TWO NUMBERS



Why it matters

Emergency procurement costs a multiple of the normal price, waiting stretches every repair, and a stock-out on the one critical spare turns a breakdown of hours into one of days. The fix is not a bigger storeroom. It is knowing which parts matter, where they are, and when they drop below the minimum. Parts-related delays above 10% of work orders is the same signal in KPI form.



Source: working bands from practice; the nearest SMRP metrics are the 5.6 MRO stores group (turns, stock-outs)

THE FOUR MOVES

- ☐ Analyse WO delays for parts-related bottlenecks
- ☐ ABC classification for spare parts
- ☐ Reorder points and safety stock for critical spares
- ☐ A PM parts kitting process

FURTHER READING

Terry Wireman, Benchmarking Best Practices in Maintenance Management, inventory chapter
ISO 55001:2014, Asset management systems, Requirements

The 90-day roadmap


- Days 1 to 5** **Analyse parts-related WO delays.** Review the WOs where actual exceeded estimate. Interview the techs.
- Weeks 1 to 2** **ABC classification.** A: critical, high consequence. B: important, substitutable. C: low cost.
- Weeks 2 to 3** **Reorder points for critical spares.** $ROP = (\text{average daily usage} \times \text{lead time}) + \text{safety stock}$.
- Weeks 3 to 4** **PM parts kitting.** A bill of materials for scheduled PMs. Kit and stage 48 hours before.
- Weeks 4 to 6** **Vendor partnerships** for fast-turnaround items. Consignment stock or emergency delivery agreements.
- Weeks 7 to 12** **Hold.** Count parts-related delays on every closed WO. Below 5% by day 90.

HOW TO RUN THIS IN FABRICO

The part belongs to a machine, has a bin, and has a minimum

In the inventory each spare is tied to the machines it fits and to its storage location, so the tech knows where to walk. The minimum quantity triggers a notification the moment a write-off takes the count below it. Because parts are written off from the card by the person who used them, the count in the system moves the moment the coupling leaves the bin, which is as close as a storeroom gets.

- ✓ QR scan on the part for the write-off; accounting integration for the costs



A low-stock notification over the parts list.
[See it live in a demo](#)

The reorder point, worked through

EXAMPLE

$$ROP = (\text{average daily usage} \times \text{lead time}) + \text{safety stock}$$

A bearing used 0.2 a day on average, 14 days from the vendor, with two in reserve: $0.2 \times 14 = 2.8$, plus 2 = 4.8. Set the reorder point at 5. When the write-off takes the bin to 4, the notification fires and the order goes out, not the technician.

MATURITY LEVEL 3

8. Condition-based monitoring

TRIGGER MTBF below 500 hours on critical assets with repeat failures.

A critical compressor fails roughly every 400 operating hours. Same failure mode every time: bearing degradation, then seal failure. The PM program says inspect the bearings quarterly. Quarterly inspections are taking quarterly photographs of a failure that runs its course in weeks.

The plant is catching the failure, not preventing it.

EXAMPLE • FAILURE CYCLE VERSUS INSPECTION INTERVAL, HOURS



Why it matters

Time-based PMs cannot prevent failures whose timing is driven by operating condition, not the calendar. That is what condition monitoring exists for. Vibration would show that bearing degrading two weeks before failure; oil analysis would show metal particles climbing. The data is there. The plant is not collecting it. CBM is the transition from Level 2 to Level 3, and it only pays on assets where the failure mode is detectable and the consequence is high.

MTBF ON MONITORED ASSETS, HOURS

ACTION • TARGET • WORLD CLASS

Below 500

Above 1,000

Above 2,000

Source: SMRP metric 3.5.1, Mean Time Between Failures; the hour bands are working targets for the compressor class

THE FOUR MOVES

- ☐ Three to five critical low-MTBF assets as pilots
- ☐ The monitoring technique that matches each failure mode
- ☐ Baseline readings and alert thresholds
- ☐ Techs trained on collection and basic analysis

FURTHER READING

SAE JA1012, Guide to the RCM Standard
 ISO 17359:2018, Condition Monitoring and Diagnostics
 ISO 20816 (formerly ISO 10816), Mechanical Vibration Evaluation

The 90-day roadmap

Days 1 to 5 Identify CBM candidates. High criticality, MTBF below 500 hours, repeat failures. Pick three to five for a pilot.

Weeks 1 to 2 Select the technique per failure mode. Vibration (rotating), thermography (electrical), oil analysis (lubricated), ultrasound (air, bearings).

Weeks 2 to 3 Baseline readings on healthy assets.

Weeks 3 to 4 Alert and alarm thresholds per ISO 20816 (formerly ISO 10816) for vibration and NETA for thermography.

Weeks 4 to 6 Integrate with the CMMS workflow. An alert generates a WO with the condition data attached.


Weeks 4 to 8 Build internal expertise. Train one or two techs. Certifications for the long term.

HOW TO RUN THIS IN FABRICO

Fabrigo is not a vibration analyser. It is where the alert becomes work.

The condition data comes from your sensors and your analysts. What Fabrigo adds is the machine state from the PLC, an IoT sensor or a camera, so every stop and micro-stop sits on the timeline with its cause, and the technician opens the work card from the stop, on the same asset, with the stop data attached. The run-hour and cycle counters that trigger PMs come from the same signal.

- ✓ Machine state and counts from PLC, IoT sensor or AI camera
- ✓ From the stop on the timeline to a task on the asset, in two taps
- ✓ PM triggers by run hours or cycle counts, not only the calendar
- ✓ What it does not do: vibration, oil or thermography analysis. Bring those tools; this is where their findings turn into work



A stop on the timeline: edit the reason or add a task on that machine. [See it live in a demo](#)

MATURITY LEVEL 2


9. Downtime reduction

TRIGGER Availability below 95%, or unplanned downtime above target.

The morning production meeting. The production manager reads off last week's numbers. Availability on Line 4: 91 percent. Line 2: 93. The numbers have been in the low 90s for six months. Maintenance says "we're working on it". Production says "okay".

Nobody asks what "working on it" actually means.

EXAMPLE • AVAILABILITY, SIX MONTHS, TWO LINES



Why it matters

Availability is the KPI that executives understand. It maps straight to throughput, to revenue, to the numbers on the monthly report. Unplanned downtime costs a multiple of planned downtime through emergency response, overtime and expedited parts. When availability is low, everything else in maintenance lives downstream of that conversation. Get it above 95 and your relationship with operations changes.

AVAILABILITY, %

ACTION • TARGET • WORLD CLASS

Below 90

Above 95

Above 97

Source: SMRP metric 2.2, Availability; 95 and 97 are working target and world-class bands

THE FOUR MOVES

- ☐ Pareto the top five downtime-contributing assets
- ☐ RCA on the worst offenders
- ☐ Review PM procedures for high-downtime equipment
- ☐ Condition monitoring on critical assets that lack it

FURTHER READING

SMRP metric 2.1.1, Overall Equipment Effectiveness
ISO 55001:2014, Asset management systems, Requirements

The 90-day roadmap


- Days 1 to 2** **Pull downtime by asset for 90 days.** Rank by hours. The top five to ten assets usually carry most of the hours: the Pareto rule in practice.
- Days 2 to 4** **Break it down by reason.** Unplanned, planned, material, operator, set-up, quality. Find the dominant one.
- Weeks 1 to 2** **RCA on the top three to five assets.** 5 Whys or Fishbone.
- Weeks 2 to 3** **Improve PM coverage** on high-downtime assets: add or change tasks, align frequencies with MTBF.
- Weeks 3 to 4** **Reduce MTTR through preparation.** Pre-stage spares. Standardise job plans. Tools and manuals accessible.
- Ongoing** **Weekly monitoring** by asset group and shift. Target one to two points of availability per month to 97.

HOW TO RUN THIS IN FABRICO

Availability is live, per line and per shift, and every stop has a cause

The OEE overview shows availability, performance and quality for each line while the shift runs, with the count and duration of planned and unplanned stops. Downtime is logged from the machine signal, so the short stops nobody writes down are counted too, and the operator tags the cause at the line. The morning meeting starts from yesterday's biggest causes per line, not from collecting them.

- ✓ OEE, availability, performance, quality per line and shift, live
- ✓ Stops and micro-stops from the signal, reason from the operator
- ✓ Downtime by cause, machine and shift, for any period



OEE overview, one card per line. [See it live in a demo](#)

MATURITY LEVEL 3

10. Asset reliability and MTTR

TRIGGER MTTR above eight hours, or critical assets with recurring downtime events.

A critical motor fails at 2 AM. The night tech diagnoses the failure in fifteen minutes: bearing. He then spends four and a quarter hours trying to find the right bearing in the storeroom. He finds it at 6:30. Installs it in ninety minutes. The WO records sixteen hours of downtime.

The actual repair time was under two. Everything else was waiting.

EXAMPLE • 16 HOURS OF DOWNTIME, WHERE THEY WENT

0.25 H DIAGNOSE




REPAIR: 1.5 H • RECORDED MTTR: 16 H

Why it matters

MTTR is the sum of all the delays between "broken" and "running". A high MTTR means the plant has not prepared for the failures it knows will happen. World-class plants hold MTTR under two hours on critical assets; plants that have not prepared often run several times that. The gap is rarely about tech skill. It is about what was staged before the breakdown: the spare, the job plan, the second person who can do the job.

MTTR, HOURS

ACTION • TARGET • WORLD CLASS



Source: SMRP metric 3.5.2, Mean Time To Repair; working bands from practice

THE FOUR MOVES

- ☐ Stage critical spares for high-MTTR assets
- ☐ Standard job plans for common failure modes
- ☐ Cross-train technicians on critical equipment
- ☐ Repair versus replace for chronic failures

FURTHER READING

SAE JA1012, Guide to the RCM Standard
ISO 55001:2014, Asset management systems, Requirements

The 90-day roadmap


- Days 1 to 5** **MTTR by asset and failure mode, then the delay causes.** Which combinations drive the longest repairs? Parts waiting? Diagnosis? Access? A specialist needed?
- Weeks 1 to 2** **Spare parts staging.** Critical spares accessible 24/7. Emergency kits for the top ten scenarios.
- Weeks 2 to 3** **Standardise repair procedures.** Job plans for common failures: steps, tools, LOTO and expected time.
- Weeks 3 to 6** **Cross-train technicians.** At least two per shift capable on each critical asset.
- Weeks 4 to 12** **Repair-versus-replace for chronic failures** (repair above half of replacement in two years says replace), then hold: MTTR by asset weekly, below four hours on critical assets by day 90.

HOW TO RUN THIS IN FABRICO

Mean recovery time after failure is a number on the dashboard, with a trend

The analytics show total downtime and the mean recovery time after failure against the previous period. The urgent task reaches the technician's phone with the machine, the photo and the procedure; the spare's bin location is on the part.

- ✓ Urgent tasks notified to the role at once, escalated if unresolved
- ✓ Job plans and part locations on the card, in the tech's pocket



Analytics cards: total downtime and mean recovery time after failure.

[See it live in a demo](#)

MTTR, worked through from the story

EXAMPLE

$$\text{MTTR} = \text{total repair time on unplanned outages} \div \text{number of unplanned outages}$$

One outage, 16 hours from "broken" to "running", 1.5 of them repair. Cut the 4.25-hour search to a walk and the recorded MTTR falls from 16 to about 12 hours, roughly a quarter, before anyone repairs faster.

Ten plays, one system. See a stop become a work card, in a 30-minute demo.

[Request a demo](#)

PART 5

Tools and cadences

The RCA quick guide, the twelve-month sequence of plays by maturity level, the references, and where Fabrigo fits in all of it.

- p. 49 RCA quick guide: 5 Whys, Fishbone, FMEA
- p. 50 Which one, when
- p. 51 The sequence of plays, by maturity level
- p. 53 References
- p. 54 Where Fabrigo fits

Root cause analysis is not one method. It is three, each with a use.

Most plants pick one, usually 5 Whys because it is the easiest, and apply it to everything. A simple failure needs a simple tool. A multi-factor failure needs structure. A critical asset needs a risk-ranked, proactive approach. The fishbone below is the pump from Play 2.

METHOD 1

5 Whys

Best for: simple, single-cause failures. Linear cause-and-effect chains. Done in under a day.

Method: define the problem. Ask "why?" five times. Verify the root cause with evidence. Define a corrective action that addresses the root, not the symptom.

METHOD 2

Fishbone (Ishikawa)

Best for: multi-factor failures. Team-based. Recurring problems where the single-cause assumption already failed.

Method: problem statement at the fish head. Six branches: man, machine, method, material, measurement, environment. Brainstorm causes, ask why to drill, rank by likelihood, verify with data.

METHOD 3


FMEA

Best for: proactive analysis of critical assets. New equipment. Design or process change review.

Method: list potential failure modes. Rate severity, occurrence and detection from 1 to 10. $RPN = S \times O \times D$, or the Action Priority tables in the 2019+ editions. Act on the highest first. Re-score after implementation.

FISHBONE • SIX BRANCHES INTO ONE PROBLEM STATEMENT

EXAMPLE




Which one, when

SITUATION	USE
Single failure, non-critical asset, one day or less	5 Whys
Repeat failure, multi-factor, team-based	Fishbone
Critical asset, proactive, before the failure	FMEA

THE LADDER GOES UP

Stack them. A 5 Whys that does not resolve escalates to a Fishbone; a Fishbone that finds a high-consequence mode escalates to an FMEA.



FURTHER READING

Taiichi Ohno, Toyota Production System, the origin of 5 Whys
 Kaoru Ishikawa, Guide to Quality Control, the original Fishbone text
 SAE J1739 (2021) and IEC 60812, FMEA

Fabrico HOW TO RUN THIS IN FABRICO

The RCA starts from the card, not from memory

Whatever the method, the input is the same: when it failed, what was found, what was fitted, how long it took, what the operator tagged. Repair history gives the sequence, the stoppage list the durations, the card photos the as-found condition.

- ✓ Repair history by machine and by task
- ✓ Stop reasons and durations per line and machine
- ✓ Corrective actions as tasks with owners, dates and close-out

The corrective action is a task on the asset with a performer and a scheduled date. Verification in week 12 is a search for it. [See it live in a demo](#)

The ten playbooks are not a buffet. Two at a time, in this order.

Each play keeps its own 90-day curve; the sequence stacks them over a year. Each one consumes planning capacity, technician attention and supervisor enforcement, so running three at once is hard and five will break. Start where the Part 3 self-assessment put you.

LEVEL 1 PLANT • REACTIVE

Months 1 to 3	PM compliance recovery (Play 1) Get the PM program off the floor. The single highest-leverage play for a Level 1 plant.
Months 2 to 4	Reactive work reduction (Play 4) In parallel from month 2. Tightly coupled with PM recovery.
Months 4 to 6	Schedule compliance (Play 5) With PMs flowing and reactive shrinking, schedule discipline is the next bottleneck.
Months 6 to 9	Backlog growth control (Play 3) By month 6 the backlog shape has shifted. Purge, prioritise, stabilise. You are now a credible Level 2 plant.

LEVEL 2 PLANT • PREVENTIVE

Months 1 to 2	Schedule compliance (Play 5) If not already running. Level 2 plants almost always have this gap.
Months 2 to 3	Backlog growth control (Play 3) Purge, prioritise, stabilise, now that the week holds.
Months 3 to 5	Repeat failure investigation (Play 2) Start the RCA discipline on your top three to five repeat failures.
Months 5 to 7	Spare parts management (Play 7) ABC classification and PM kitting.
Months 7 to 9	Inspection compliance (Play 6) If the overdue rate is elevated. Regulatory inspections are non-negotiable.
Months 9 to 12	CBM pilot (Play 8) Three to five critical assets. The transition play to Level 3.

LEVEL 3 PLANT • CONDITION-BASED

Months 1 to 3	Downtime reduction (Play 9) If availability is below 95%.
Months 2 to 4	Asset reliability and MTTR (Play 10) MTTR reduction: staged spares, standard job plans, cross-training.
Months 4 to 9	CBM program expansion (Play 8, expanded) From five pilot assets to twenty-plus critical assets with integrated alerts and trending.
Months 9 to 12	Repeat failure elimination at scale (Play 2, at scale) Apply RCA systematically, asset family by asset family.

WHAT A WIN LOOKS LIKE • BEND THE CURVE, THEN HOLD IT



A NOTE ON PACING

You are not trying to finish a play in ninety days. You are trying to bend its curve, then hold the improvement before reaching for the next one.

BEFORE THE FIRST PLAY: PRICE THE GAP

The OEE calculator at fabrico.io/calculator-oeo turns the distance between your OEE and your target into lost units per line. Write the result next to the Availability row on page 25. It is the number that gets the play its budget.

References

Organised by type. The methods and the targets have a citation; where a band is a working rule, the page says so. If you disagree with a number, disagree with the source, then write your own target in the Yours column on page 25 and hold it.

SMRP BEST PRACTICES

2.1.1 Overall Equipment Effectiveness

2.2 Availability

3.5.1 Mean Time Between Failures

3.5.2 Mean Time To Repair

5.3.1 Planned Work

5.4.1 Reactive Work

5.4.2 Proactive Work

5.4.3 Schedule Compliance, hours

5.4.4 Schedule Compliance, work orders

5.4.9 Ready Backlog

5.4.10 Total Backlog

5.4.14 PM and PdM Work Order Compliance

5.5.2 Craft Worker to Planner Ratio

5.6 MRO stores metrics (turns, stock-outs)

Where there is no SMRP metric (repeat failures, inspections, parts delays) the bands are working bands and say so.

ISO STANDARDS

ISO 55000:2014 Asset Management Overview

ISO 55001:2014 System Requirements

ISO 55002:2018 Application Guidelines

ISO 14224:2016 Reliability Data

ISO 17359:2018 Condition Monitoring

ISO 20816 (formerly 10816) Vibration Evaluation

SAE STANDARDS

SAE JA1011 RCM Evaluation Criteria

SAE JA1012 Guide to the RCM Standard

SAE J1739 FMEA in Design and Manufacturing

INSPECTION CODES (REGULATORY)

API 510 Pressure Vessel Inspection

API 570 Piping Inspection

API 653 Tank Inspection

OTHER STANDARDS

IEC 60812 Failure modes and effects analysis

MIL-STD-1629A FMEA (historical, cancelled 1998)

OSHA 1910.147 Lockout/Tagout

ANSI Z244.1 Lockout/Tagout

NETA MTS Maintenance Testing Specifications

ORGANISATIONS

SMRP, smrp.org

Reliabilityweb and Uptime Magazine

IMC, International Maintenance Conference

Maintenance World journal

BOOKS

Doc Palmer, Maintenance Planning and Scheduling Handbook (4th ed.)

Ron Moore, Making Common Sense Common Practice (4th ed.)

John Moubray, Reliability-Centered Maintenance II

Terry Wireman, Benchmarking Best Practices in Maintenance Management



Kaoru Ishikawa, Guide to Quality Control

Taiichi Ohno, Toyota Production System

Robert Abernethy, The New Weibull Handbook (5th ed.)

Know why your lines stop. Fix it in the same system.

Fabrico is a cloud-based MES and OEE platform for manufacturers, with maintenance management (CMMS) built in. The machine data, the stop with its cause, the work card, the spare part and the report live in one place, so the cycle in Part 2 runs without a spreadsheet between the stages.



Fabrico on web and mobile: the OEE overview and the hourly line timeline.

Sample data

- ✓ **Machine data** for availability, performance and quality, per line and shift (Plays 9 and 10)
- ✓ **Every stop with its cause**, from operators or AI computer vision (Plays 2 and 4)
- ✓ **A stop can become a work card** in the same system: work orders, PM plans, spare parts, QR codes (Plays 1, 3, 5 to 7)
- ✓ **Web, iOS and Android apps.** Interface in English, German, French, Polish and Bulgarian
- ✓ **ISO 27001 certified (DNV).** Hosted on AWS in an EU region

WHY THIS GUIDE IS FREE

We make money when a plant decides to run these plays in Fabrico. Most readers will run them on paper and a CMMS they already own, and that is fine: a plant that fixes its system is a better customer later, or a better plant, full stop.

Bring your worst number. Leave with the play mapped onto your lines.

A 30-minute demo with a Fabrico consultant. No commitment. Here is exactly how it runs, so you can decide whether it is worth half an hour.

0 to 5 min

Your number

You tell us the KPI from the benchmark table that hurts most, and the asset behind it. We pick the play.

5 to 15 min

The stop becomes a work card

On a line like yours: the signal, the stop, the operator's reason, the card to the technician, the part written off, the report.

15 to 25 min

The play, on the screens

The roadmap from Part 4 for your KPI, step by step, in the product. You see where your data would come from on day one.

25 to 30 min

Your questions

What Fabrico does not do (page 43), what connecting your machines involves, and what a rollout on your lines would look like.

WHO YOU TALK TO

A Fabrico consultant who has run this on real lines.

WHAT IT COSTS

Thirty minutes. Nothing else. Nothing to install, nothing to set up.

WHAT TO BRING

The KPI from page 25 that hurts most, the asset behind it, and last month's number.

The next page has the link and the QR code.

NEXT STEP

See your worst KPI on the screens, in 30 minutes

A 30-minute demo with a Fabrigo consultant, on a line like yours, with your questions. No commitment. The gearbox on page 5 fails again in eight to ten weeks; a play started this week has a result before it does.

[Request a demo](#)

[Contact us](#)



Scan to request a demo

Or e-mail team@fabrigo.io with the subject line "Maintenance control" and the number you want to move. A consultant replies to set up the 30 minutes.

fabrigo.io · ISO 27001 (DNV) · AWS, EU region