No more surprises?

Linking MRO inventory to maintenance tasks

Mark Horton, August 2014

Why is life full of surprises?

I once witnessed an interesting discussion between a store controller and a maintenance consultant.

"There were far fewer surprises"

The consultant had given a presentation on his experience—which was naturally all good—of implementing Reliability-centred Maintenance (RCM) for part of chemical manufacturing company. His presentation had finished with a rousing list of benefits that his client could expect from its new maintenance schedules. They included higher availability, better safety compliance, lower costs, and an overall move from reactive to proactive maintenance.

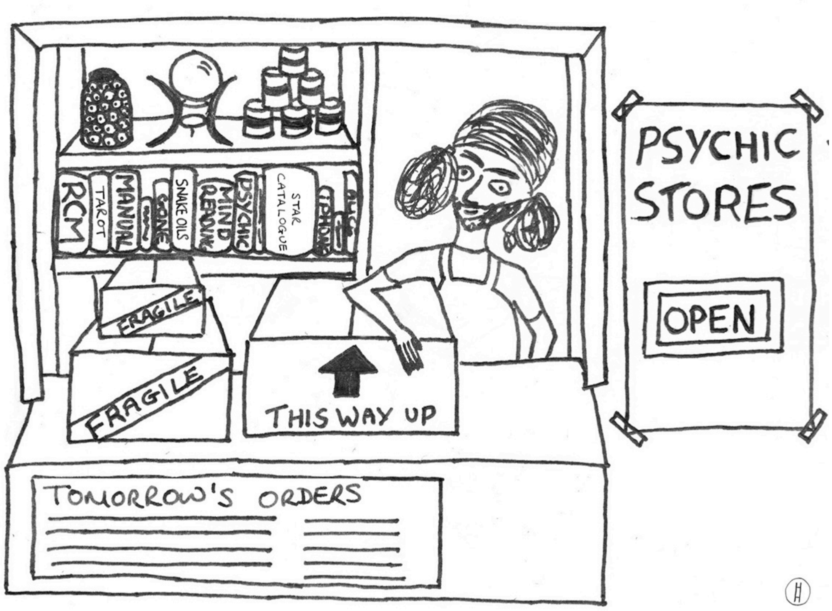

The maintenance organisation was in control, he said. With the old maintenance schedules, most of their work involved fixing unexpected failures that always seemed to happen at the worst time. The operations manager had been on their backs because downtime cost money. Now everyone came to work feeling more relaxed because they knew what had to be done. There were far fewer surprises.

The store controller had worked in a completely unrelated organisation, and he had a different view of the world.

"When my company started a maintenance review a couple of years ago. I went to a senior management briefing that promised us most of the benefits on your last slide.

"There was another presentation when the review was finished and implemented. It said the same things as your last slide again: fewer surprises, more control. That sounded good. Under the old maintenance schedules, any day could bring a demand for critical parts that we didn't have.

"Every few months the financial controller would get steamed up about the value of inventory that wasn't used. Control sounded like something that we desperately needed.

"why am I still not in control of our maintenance spares?"

"So I have just one question for you. If our maintenance team is so confident and relaxed, why am I still not in control of our maintenance spares?"

These were two people looking at the same world and seeing totally different things. The maintainers saw control: work could be planned, managed and executed with very few unexpected failures. The store controller saw chaos and almost nothing but unwanted surprises.

Linking demand and maintenance

How is it possible to make the calm and peace of the new maintenance schedules flow through to the stores organisation? Does the stores organisation have to accept that each demand arrives without any warning, or is there a better way?

Successful inventory management involves choosing the correct reordering policies and optimal stock level for each line item. Just-in-time policies—ordering items so that they are available at specific dates—reduce the capital tied up in inventory and at the same time eliminate the risk that obsolete items will be left on the shelf. At the day-to-day level, good inventory policy selection makes the difference between a smooth, efficient organisation and a chaotic service where every demand is a surprise.

Understanding demand sources is the key to moving from reactive to proactive stock management.

Far from being a threat to the engineering inventory, the development of a formally-derived maintenance schedule provides a real opportunity: to link inventory policy, line item by line item, to an audited list of demands.

One specific problem experienced by modern stores is the shift from "hard time" task to condition-based maintenance.

A typical traditional maintenance schedule includes a large proportion of fixed-interval overhaul and replacement tasks. While they are often not the most effective way to manage failures, and sometimes even cause more problems than they prevent, these tasks are good news for inventory planners for one reason: maintenance demands are regular and entirely predictable. In contrast, modern schedules contain a majority of condition-based tasks because they are cheaper, more effective, and they maximise useful asset life.

But what is good news for the maintenance budget is bad news for the inventory planner: a predictable demand calendar has been replaced by a series of unpredictable demands, often with little advance warning. Stock policy has become disconnected from the root cause of demands.

There are four steps in the process of ensuring that spares policy matches maintenance requirements.

1	Identify the maintenance tasks that can give rise to demands for spare parts or maintenance materials
2	Classify each maintenance task as described below
3	Determine whether it is possible to predict when a demand will occur
4	If it is not possible to predict a demand, determine how many spare parts should be held in stock

Failure management policies

The first step in bringing stores and maintenance together is to understand the maintenance actions that trigger demands. Demands arise both from planned tasks and from unpredictable breakdowns, so it is vital to consider all *failure management strategies*. These include regular, proactive maintenance and management strategies that allow failures to occur: the list of demand sources is not limited just to the list of planned maintenance tasks.

Maintenance can respond in a number of ways to possible equipment failures. Ignoring the possibility of a one-off change such as redesigning the equipment in some way, there are five ways in which failures can be managed. Every maintenance task should fit into one of these categories, whether or not your organisation has implemented Reliability-centred Maintenance.

Sometimes the maintenance tasks in an ERP or maintenance management system may be a combination of two or more tasks, like this one:

> Every week: clean the filter by back flushing with alcohol. Change the filter if the mesh is corroded.

This task could be picked apart into two separate activities:

> Every week: clean the filter by back flushing it with alcohol.
>
> Every week: check the filter mesh and change it if corrosion is visible.

Scheduled Discard

Also known as	Scheduled replacement
What it does	Parts are replaced with new items before they fail. The part is replaced at fixed intervals regardless of the part's condition at the time.
Examples	Replace car brake fluid every two years Replace rubber hydraulic hoses every five years

Scheduled Refurbishment

Also known as	Scheduled overhaul
What it does	Parts are overhauled or refurbished before they fail. The part is refurbished at fixed intervals regardless of the part's condition at the time.
Examples	Clean the mesh filter every two weeks Lubricate pivot pin every day before use

On Condition

Also known as	Condition-based maintenance (CBM) Detective maintenance Scheduled inspection
What it does	Parts are inspected at defined intervals and they are repaired or replaced if they fail to meet minimum requirements.
Examples	Check each truck tyre every week and replace the tyre if its tread depth is less than 1.5 mm Measure the turbine bearing's vibration monthly and replace the bearing if its vibration level exceeds specified limits Measure the mesh filter's differential pressure every week and clean the filter if the pressure exceeds 0.1 bar

Failure-finding

Also known as	Scheduled testing
What it does	A protective device is tested at specified intervals and the device is repaired or replaced if it does not function correctly
Examples	Test the building's fire alarm weekly and schedule repair or replace of any failed detectors, sirens or other parts Test an oil pipeline's pressure relief valves every three months and repair or replace any valves that fail to lift at the defined pressure

Corrective Maintenance

Also known as	Run-to-failure Breakdown maintenance
What it does	No action is taken to prevent or predict a failure.
Examples	Replace a building's fluorescent lights when they fail Replace an industrial logic controller's circuit boards when they fail

Failure management and planned demand

When each failure management policy has been identified, it is time to work out whether any store demands that it generates can be anticipated and whether parts can be ordered to meet the demand ("just-in-time") rather than held in stock waiting for the demand to occur ("just-in-case").

Scheduled Discard and Scheduled Refurbishment

These are the simplest maintenance policies to supply because the tasks and their demands occur at fixed calendar intervals.

A maintenance review selects one of these policies if both of these conditions are satisfied:

1. There is an identifiable age after which the probability of failure increases rapidly
2. Most (preferably all) of the parts survive to that age

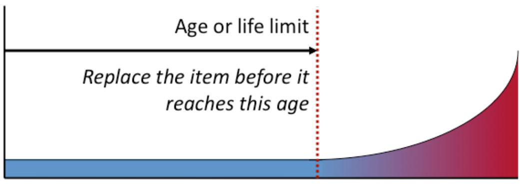

The inventory planner just needs to ensure that the parts are ordered, delivered and ready for use before each task is carried out. The maintenance task schedule translates directly into an ordering calendar: add up the item lead time, order and delivery times, any delays and contingencies, and allow for any uncertainties. Then subtract the time from the task date and arrange the orders.

Proactive maintainers sometimes take advantage of an early maintenance opportunity

Before committing to a pure just-in-time, dependent demand policy, remember to check that there are unlikely to be any early life failures before the scheduled discard or replacement task. It is possible that the maintenance review group was aware of the possibility of early life failures, and assumed that spare parts would be available for those breakdowns as well: it may be necessary to procure stocks specifically to support the initial high failure rate.

Finally, it is important to be aware of any changes to the planned maintenance schedule. A delayed maintenance task does not matter much—at least not to the inventory operation—provided that the parts can be stored somewhere between delivery and use. But proactive maintainers sometimes want to carry out maintenance earlier than planned because of an unexpected opportunity, perhaps as the result of a failure in related equipment or changing operational requirements. Not holding a part in stock may be a very unwelcome constraint and reduce the maintenance planner's flexibility.

On-Condition

This type of maintenance task isn't as simple to manage as scheduled discard and refurbishment. To see why, you need to understand how a maintenance review sets the on-condition task's interval.

An on-condition task looks for signs of deterioration showing that a failure is developing. These signs are called a *potential failure* condition. When the part has completely failed, it has experienced a *functional failure*. It is important to understand the distinction between the two states, so here are a few examples.

Part	**Functional failure**	**Potential failure(s)**
Filter	Filter blocked	High differential pressure
Switch	No current	High contact temperature
Pump impeller	Flow is too low for process	Flow is reduced

An effective on-condition task has to detect a potential failure condition and provide sufficient warning to avoid a functional failure. The time period between the earliest signs of a potential failure and the ultimate functional failure is called the P-F interval. This interval determines how often the on-condition task has to be carried out.

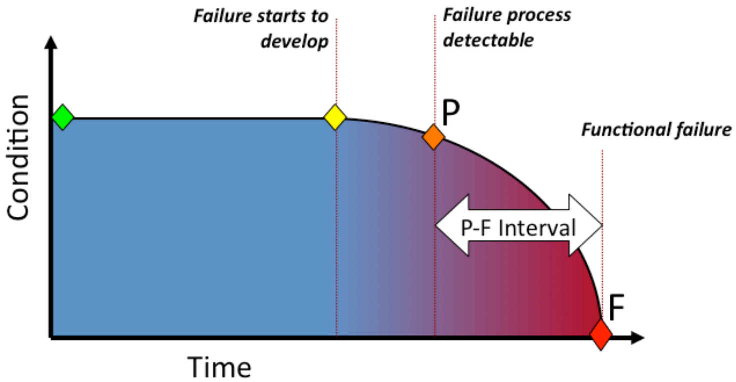

Deterioration in condition over time.
Deterioration can be detected at P and failure occurs at F

The P-F graph might look like this for a helicopter drive shaft. Remember that "condition" is on the vertical axis, so upwards means "better condition" or "smaller and fewer cracks".

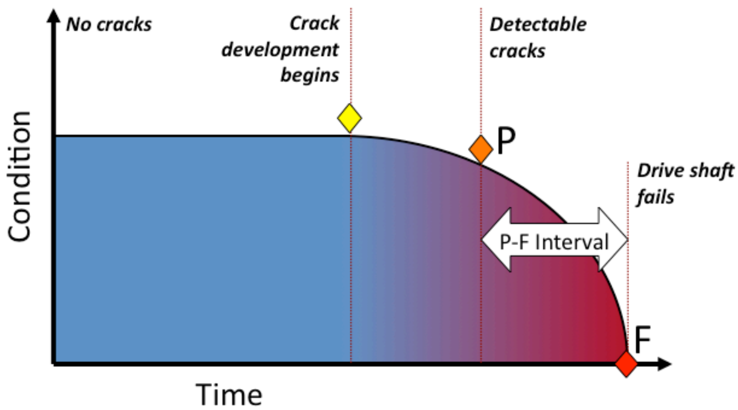

Deterioration of a drive shaft's condition as determined by non-destructive testing

On-condition tasks consist of two parts: the inspection, usually carried out at fixed intervals; and the remedial action that is carried out if the item fails inspection.

> Inspect the helicopter gearbox rotor drive shaft every 1000 flying hours and schedule replacement if cracks larger than limit size are found.

In this example, the inspection can only be carried out by removing an access plate which is sealed by a gasket and a set of bolts. A kit of these parts is needed every time the inspection task is carried out (1000 flying hours). The replacement drive shaft—along with many other parts—is needed only if a crack is found, and it is not possible to know when those demands will occur.

Although the failure's MTBF tells you something about the average demand rate on the store, it does not determine whether or not you need to hold parts in stock

The materials needed to carry out the inspection can be safely ordered in the same way as for scheduled restoration and discard tasks. However, the monitoring period (the time between tasks) is not related to the part's failure rate or mean time between failures (MTBF). Although the failure's MTBF tells you something about the average demand rate on the store, it does not determine whether or not you need to hold parts in stock: that decision depends on the P-F interval and the chosen condition monitoring interval.

Imagine that the shaft's P-F interval (the interval between detectable cracks and failure) is 2000 hours and that the maintenance review group decided to inspect the drive shaft every 2000 hours. This could give 2000 hours' notice of failure, but if the maintainer were unlucky, the inspection task could just miss a failure that is in the process of developing. Next time round, the maintainer could get almost no warning of failure at all.

Is it possible to order a new shaft when the task detects cracks? It might be if the overall lead time is less than the calendar equivalent of 2000 flying hours. But sometimes the warning period is very short, so there is no guarantee that there would be time to order a new shaft when the first cracks are detected.

Fortunately most maintenance review groups set the task interval at something less than the P-F interval. With a P-F interval of 2000 flying hours and a monitoring interval of 1000 hours, there would be at least 1000 hours' warning of a demand. The warning period increases to 1500 hours with 500-hour monitoring, and to 2000 hours if some way to monitor the shaft continuously can be found. If the total time taken to obtain a part from ordering to receipt into the warehouse is less than the warning period, it may be possible to avoid holding the drive shaft in stock to support this failure mode.

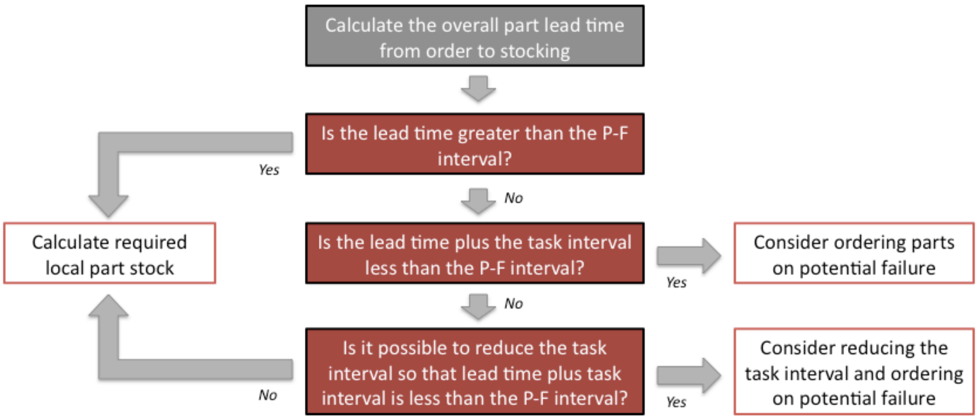

On-condition task stores decision logic

Failure-finding

A failure-finding task tests a protective device to determine whether it has failed. This type of task can only be applied to equipment that provides a *hidden function*, in other words one that does not operate in normal circumstances. The devices involved are usually alarms, trips and standby equipment, including the following.

- Fire alarms
- Standby pumping equipment that operates if the duty pump fails
- High level trips in a liquid storage tank
- Fuses
- Residual current circuit breakers
- Pressure relief valves
- Gas alarms

Failure-finding tasks are only applicable to hidden functions, but remember that it is possible to manage hidden failures using scheduled restoration, scheduled discard and on-condition tasks as well. Don't assume that every task applied to an alarm or trip system is a failure-finding task just because the device provides a hidden function.

Like on-condition maintenance, failure-finding tasks consist of two parts:

- The testing task, usually carried out at regular intervals
- The remedial task which is only carried out if the device or system is found to be in a failed state

Parts that support the regular testing task may be ordered in the same way as for scheduled refurbishment and discard tasks so that they are available on the task dates.

It is impossible to know whether the device is working without carrying out a test, so the requirement for remedial parts (the parts needed to repair the failure) cannot be predicted. As a result it is usually necessary to consider holding stock of the remedial parts. Even so, there are some circumstances in which it may be possible to order parts when the demand occurs.

1 It is acceptable to operate the equipment without the protective device until a new part has been obtained.
2 The time taken to prepare for replacement of the failed part is longer than the part's lead time. This may happen if it takes an extended period to obtain permits, shut down the equipment, gain access and to prepare maintenance equipment such as scaffolding.

Corrective

Corrective maintenance work demands spare parts and maintenance material shortly after a failure occurs. Since failures usually occur with no predictable pattern, it is not possible to order parts before they are needed. Consideration should be given to holding part stocks as discussed in a separate paper.

Even though the store room has no warning of a demand before it occurs, there are some unusual circumstances in which it may still be possible to order parts when the demand occurs. These occur if the time taken to prepare for the repair, including access time, preparation, gaining access and so on, is longer than the part's lead time. In these circumstances it may be possible to order parts on a just-in-time basis.

One of the advantages for the store planner of a formal maintenance review process like RCM is that it does not just tell you about the planned maintenance tasks that the group chose to select. It also lists all the failures that the group knew could happen, but where there was either no maintenance task that could prevent them, or where carrying out maintenance was more expensive than letting the equipment fail.

This information is a gold mine for the warehouse: it means that you can plan to support both scheduled tasks and breakdown requirements effectively. The days of guessing what failures could happen are over!

Is a Just-in-Time Policy Really Feasible?

Finally, before committing to a just-in-time policy, consider whether there are any factors that could affect your decision. Some of these are listed below.

Administration

Does your maintenance organisation have the discipline to plan ahead (sometimes years ahead) and does it have the processes in place to allow maintainers to place and manage forward orders?

Have you allowed for all possible administrative delays?

Are you certain that your EAM or purchasing system specifies the right part and variant?

If parts are obtained from abroad, have you considered possible customs requirements?

Maintenance policy

How certain are you that the maintenance policy has been correctly derived and audited?

How applicable is your organisation's maintenance to the way that you operate the equipment?

Is it possible that the maintenance policy will change within the planning period? If it does, how will you find out?

Consistency of the supply chain

Will the manufacturer and distributor of the part be able to maintain the same performance over your planning period?

Are there alternative suppliers who could provide the same items if required?

Is it possible to put in place pooling arrangements within your organisation?

Equipment obsolescence

If the equipment you are supporting will approach the end of its life within your planning period, are you sure that your suppliers will be able to provide the same response as they have in the past?

Do you need to consider holding local stock in case the manufacturer of the part stops production or changes its production process?

Are there other suppliers or alternative parts that could provide a fallback if the original parts are not available?

Are there opportunities to pool obsolescent parts with other users?

Multiple assets

Some parts are common items that support several assets (sometimes whole fleets) and multiple maintenance tasks. If the maintenance tasks are not carried out at the same time, it may be easier to provide parts through your ERP's standard stock-based reorder cycle than by setting up a large number of just-in-time orders.

How do I implement a just-in-time policy?

Effective and reliable implementation of a just-in-time policy is critical, and it depends on your organisation's purchasing systems and on effective communication between maintenance, operations and inventory functions. In a separate paper we look at how to set up and administer an effective mix of stock policies using the policy and reorder options available in modern ERPs.

What if I need to hold stock?

Sometimes it isn't possible to plan for a demand. In a separate paper we look at the factors that need to be taken into account to ensure that your inventory can support these maintenance tasks while at the same time as minimising your investment in stock and taking account of obsolescent equipment.

Summary

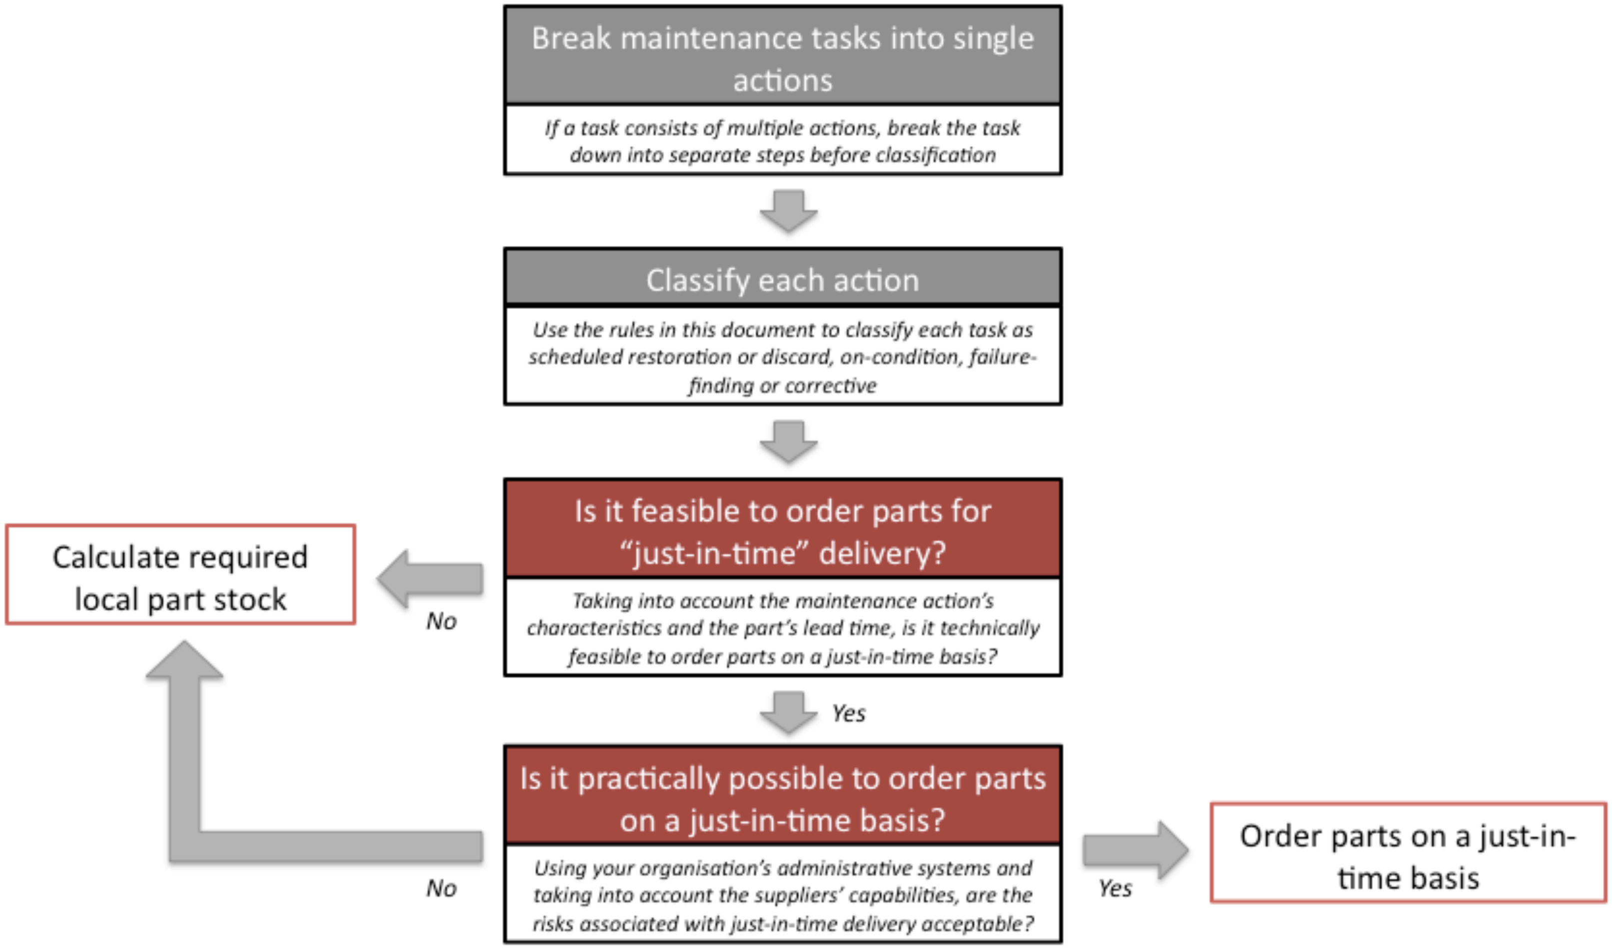

Policy selection flowchart

Management Policy	Need to Know	Stock Policy	Comments
Scheduled discard	Maintenance task dates Lead time Other logistical delays	Calculate the overall part lead time Order parts to ensure that they are available on the task dates	Check for possible early life failures that mean you have to hold stock locally Allow for variations in lead time and task dates
Scheduled refurbishment	Maintenance task dates Lead time Other logistical delays	Calculate the overall part lead time Order parts to ensure that they are available on the task dates	Check for possible early life failures that mean you have to hold stock locally Allow for variations in lead time and task dates
On-condition	P-F Interval Task interval Lead time Other logistical delays	Order parts needed for the *inspection* task so that they are available on the task dates If the task interval plus lead time is less than the P-F interval, order remedial parts when the potential failure is detected. Otherwise hold local stock of remedial parts	If the lead time is less than the P-F interval but the task interval is too long, consider reducing the monitoring interval to enable just-in-time ordering.
Failure-finding	Failure rate Required system availability	Order parts needed for the testing task so that they are available on the task dates Consider holding local stock of remedial parts	See separate stock level paper
Corrective maintenance	Lead time Other logistical delays Maintenance setup time	Consider holding local stock	See separate stock level paper

Terms of use and Copyright