

An introduction to

RELIABILITY-CENTRED SPARES

*linking spare part holdings
to maintenance and
production needs*



Copyright © 1995-2001 Information Science Consultants Ltd

Copies of this publication may be obtained from:

Information Science Consultants Ltd
3 Oak Street
Lechlade
Gloucestershire GL7 3AX
UK

Telephone (01367) 253739
International +44 1367 253739
Fax +44 8700 543144
Electronic mail info@infoscience.co.uk
WWW www.infoscience.co.uk

RELIABILITY-CENTRED SPARES

Matching spare part holdings to maintenance and production needs

1 Introduction

Modern engineering stores hold a wide range of parts, from cheap consumables used in thousands per year to critical insurance spares costing tens or hundreds of thousands of pounds which may never be used over the entire lifetime of the plant. Up to 50% of inventory value may consist of spare parts which are used at the rate of one per year or less; parts to the value of 10% to 30% of the inventory can sit on a store's shelves for a whole 25 year plant lifetime. Taking a financial view, perhaps these parts should never have been purchased; on the other hand, if they were not available when they were needed, the business could suffer severe downtime consequences.

This paper describes a new technique, derived directly from Reliability-centred Maintenance, which addresses these specific problems for any inventory, whether consumables or slow-moving parts. Substantial savings are usually achieved by applying the method to expensive, slow-moving, critical parts. RCS determines the level of spare parts inventories based not on manufacturer's recommendations, nor on a subjective judgement of service level, but on the requirements of the equipment and maintenance operation that the inventory supports.

1.1 Pressures for Change

Over the middle part of this century industry began to rely on mechanised assets to generate wealth. As those machines became a more important part of industrial production, so equipment downtime became more critical. It now mattered a great deal if equipment failed, and industry responded by developing planned preventive maintenance schedules whose aim was to replace components or overhaul assets before a failure actually happened. At the same time those assets had become more complex and their spare parts had become far more costly

and more difficult to obtain or fabricate. This combination of pressures from high equipment downtime costs on the one hand and long part lead times on the other meant that businesses could no longer afford to rely on third parties to provide those parts when a failure occurred. In this new environment lack of a spare part could cause protracted equipment downtime and lead to consequences which could be far worse than the original failure of the component. Companies therefore acquired large stocks of spare parts to satisfy both planned and unplanned maintenance needs, particularly where the consequences of not having a part on hand could result in massive loss of sales or even bankruptcy.

“In the short term it is possible to cut inventory to almost any extent...”

During the 1990's there has been sustained pressure to increase the return a business makes on capital invested. On one hand, efforts are made to increase production rates and quality; on the other, there are strong pressures to reduce the capital tied up in the business, including high value spare parts. This is the central dilemma of decision-making: in the short term it is possible to cut inventory to almost any extent, but if we do, the consequences may be felt many years later in the form of extended downtime.

Finally, the general return to core business has forced many companies to ask whether they should hold stocks of spare parts at all. Not only are stores contracted out, but maintenance "healthcare" contracts relieve the operating company of both maintenance and spare parts responsibilities. RCS provides the tools to ensure that these contracts represent value for money, and suggests what conditions should be attached to vendor contracts.

1.2 New Techniques and Technologies

Stockless production is sometimes thought of as a new phenomenon, but in fact the pressure to reduce manufacturing and distribution inventories has persisted throughout the industrial era. During the 1920s many companies in the USA became obsessed with minimising stocks and maximising stock turnover (the ratio of the value of stock sold per year to the value of stock held). Without adequate systems, communications and technology to support this ideal, businesses which were unable to cope with manufacturing problems and variations in demand went under. However, over the past two decades, two techniques have been successful in minimising manufacturing inventories:

Materials Requirements Planning

Materials Requirements Planning (MRP) analyses the subassemblies and raw materials needed to manufacture a product, and uses information about the production process and part lead times to determine the number of each product, subassembly and component needed during any period.

Just-in-Time Manufacture

Just in Time Manufacture (JIT), developed by the Japanese automotive industry, has revolutionised manufacturing industry. The goal of the system is to limit the level of buffer stock at each stage of the manufacturing process. Once that level of stock has been produced, production stops until there is further demand. JIT is often termed a "pull" system: the demand for final product stimulates production of subassemblies, which in turn controls production of components, and so on down the production line. Under JIT equipment availability is all-important: downtime at any point in the system quickly halts the entire production process.

“Production may stop for days or weeks if a spare is not available...”

The removal of buffer stocks in a JIT system places heavy emphasis on equipment reliability. The buffer stocks of a traditional manufacturing

environment represent tied up capital, but they do enable minor breakdowns to be tolerated. The same breakdowns under JIT can bring an entire production facility to a halt within minutes. This pressure on reliability applies equally to the availability of engineering spares. Production may stop for days or even weeks if a spare is not available when required.

Maintenance has adapted well to this new environment, but vendors of spare parts are often unwilling or unable to give adequate guarantees of part availability and lead times. Where one company's output is someone else's input, this can lead to production standstills on a much larger scale, and possibly additional losses because of penalty clauses. It is therefore essential to hold adequate stocks of parts to cover both planned and unplanned maintenance: but what stocks?

Service level

The concept of *service level* was developed to give a simple idea of the effectiveness of an inventory. A service level of 90% means that one in every ten of a client's demands will not be met; at 99%, the store will supply ninety-nine out of every hundred demands successfully.

Service level can be used either as a measure of effectiveness or as a goal. For example, an organisation might specify a target service level of 99%, and then translate that goal into allowed minimum and maximum stock levels using appropriate statistics as shown below.

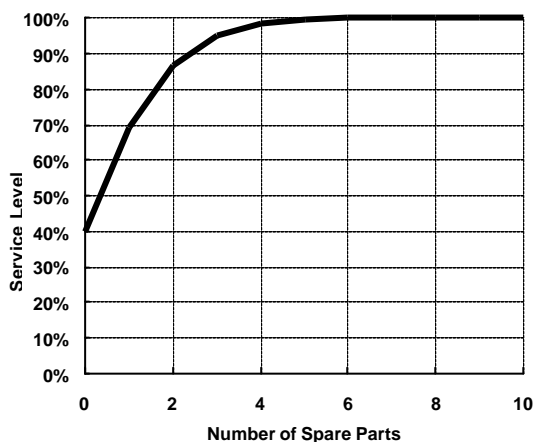
“What does it mean to have a service level of 95% on a motor which is needed only once in ten years?”

The concept of service level is easily applied to fast-moving stock, and relates well to our everyday experience of shops and warehouses as well as the supply of nuts, bolts, gaskets and seals for engineering use. However, service-level driven policies fall down when they are applied to engineering spares.

- The idea does not work well with slow-moving stock. For example, what does it mean to have a service level of 95% on a

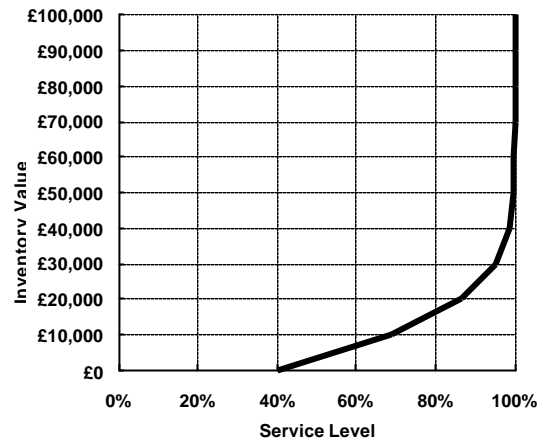
spare motor which is needed, on average, only once every ten years?

- Someone must choose the required service level for each part. Very high service levels, often well above 99%, are needed for critical engineering spares. Here we encounter problems with human perceptions of risk, since we all have problems in estimating numbers which are very close to 1. A 99% service level sounds very good and 99.9% as near perfect as could ever be needed. Compare this with the reality that, because of the consequences of a stockout, the effective service levels needed from critical insurance spares are 99.999% or even higher.



In this example, 4 parts must be held to achieve a service level of 99%

- Finally, the relationship between the service level required and the number of spares needed is not an intuitive one. The graph above illustrates the increase in service level achieved as the number of spares held increases. The picture is nothing like so comforting if we turn the same graph around, and show the *cost* of achieving a given service level. The higher the number of spares held, the lower the return we achieve by buying an additional spare part.



The cost of achieving the service level above if each spare part costs £10,000

Economic Order Quantity

The Economic Order Quantity (EOQ) method works by striking a balance between the cost of holding stock and the cost of ordering. It recognises that every purchase order carries some overhead (estimates range typically from £1 to £60 per order processed). To reduce this cost, we must make fewer, larger orders. On the other hand, larger orders mean more money tied up in stock on average. The economic order quantity is a compromise between these two extremes which minimises the total cost to the business.

ABC Analysis

In the years following the second world war there was growing appreciation that planning and statistical techniques could be applied to limit inventories and optimise inventories of manufactured goods. The growth of stock analysis itself posed a problem: one could spend almost any length of time developing models of demand, supply and logistics, sometimes spending more on stock analysis than could conceivably be saved by it.

In response to this problem, the ABC classification system was devised by General Electric in the 1950s, based on the observation that a small percentage of items accounted for a large proportion of total sales value, and conversely that a large number of items stocked accounts for only a small proportion of sales. The ABC system divides stores holdings into three categories, A, B and C parts.

Classification	Proportion of Items	Proportion of Sales Value
A	5%-20%	60%
B	20%-30%	20%-40%
C	50%-75%	5%-25%

Whatever techniques are to be used to set stock levels for each item, it is clear that far more care and effort should be taken on the analysis of A items than on B or C classifications.

Some attempts have been made to apply the ABC system to engineering inventory, since there is little point in optimising stocks of a £100 seal if £20,000 bearings are massively overstocked. Traditional ABC analysis does not apply directly to slow-moving parts, but it will be shown later that a similar principle provides a valuable way of achieving fast payback on a stock analysis programme.

Operations Research

Operations Research (OR) encompasses a broad variety of mathematical techniques which are applicable to business decisions, and many commercial software systems are available which are applicable to slow and fast moving items. Far greater success has been achieved in this area with fast moving stock than with slow moving insurance spares. In fact, despite the promise of some queuing theory models, these techniques can recommend *increased* stock levels which go against all common sense.

While it is possible to argue that the assumptions made in these models are unrealistic (for example, the use of a fixed holding cost per item), mathematical tools fail not so much because of their underlying models and assumptions, but usually because the blind application of techniques without understanding the underlying causes of demand: the operation and maintenance of equipment.

New maintenance techniques

During the 1980s and into the 1990s the process of change in industry has continued to accelerate. The industrial climate has demanded improved availability and reliability, greater safety and environmental integrity together with ever higher levels of cost effectiveness.

“We must make certain that stores respond to changes in maintenance policy”

In response to these pressures, maintenance moved away from the principle of fixed interval planned overhaul or replacement towards a *reliability-centred* approach, where maintenance is tailored to the requirements of each item of equipment in its own operating context. The result is the widespread use of condition monitoring equipment to detect problems before a failure occurs, as well as the recognition in some cases that *it is simply not cost-effective to do anything to prevent failure*.

However, if the function of our engineering inventory is to support maintenance, we must make certain that our stores respond to changes in maintenance policy.

2 The Way Forward

What is needed is not a new purchasing system or new procedures, but an auditable method to ensure that inventory fully supports operations and maintenance. The new method described here is an extension of Reliability-centred Maintenance to cover spare parts and inventory services. It is applicable to any engineering inventory, whether fast-moving consumables or slow-moving insurance spares. In practice the greatest return is usually achieved by a detailed analysis of slow-moving stock.

However, before analysing spares requirements, it is essential to ensure that we understand *what* we are supporting:

Step 1: *Ensure that the maintenance requirements of the asset(s) are clearly understood*

No spare parts policy or inventory system can do anything to make an ineffective maintenance schedule more effective. An efficient inventory system may even *increase* costs if it simply supplies parts more quickly to a poorly planned maintenance programme.

Spare parts support both planned, preventive maintenance programmes and breakdown

maintenance (No Scheduled Maintenance in RCM terms). An RCM analysis of maintenance requirements is invaluable because it identifies both planned requirements (which may be predictable) and breakdown requirements (which are inherently unpredictable).

Step 2: Determine the spares requirements of the plant in the context of its operations and maintenance

This is the subject of this paper.

Step 3: Ensure that the resources, procedures and systems are in place to deliver the requirements determined in step 2

Once we know what the spares requirements are, the company's administration procedures, computer purchasing, maintenance and inventory systems must be capable of providing the required response times, ordering and part reservation facilities.

3 The RCS Process

The Reliability-centred Spares method consists of a series of questions, starting with the ways in which equipment can fail (failure modes), moving through the effects of failure and the effects of a stockout (part unavailability) to setting the correct stocking policy for each spare part.

RCS: Five Basic Questions

What are the maintenance requirements of the equipment?

What happens if no spare part is available?

Can the spares requirement be anticipated?

What stock holding of the spare is needed?

What if the maintenance requirements cannot be met?

The first question is answered as part of a Reliability-centred Maintenance (RCM) analysis. The final four questions ensure that spares inventories and systems match the needs of operations and maintenance.

3.1 What Happens if No Spare is Available?

RCS bases the stockholding decision not on manufacturer's recommendations, or on engineering judgement, but on *what happens if no part is available*. This step in the process makes it possible to decide whether the stockout matters, and hence what resources are needed to reduce the risk of the stockout occurring.

Like RCM, RCS recognises five categories of consequences:

Hidden (Increased Risk)

The failure (for RCM) or stockout (for RCS) itself has no direct consequences, but we are exposed to an increased risk of the consequences of another failure

Safety

The failure or stockout itself has direct consequences which could hurt or kill someone

Environmental

The failure or stockout itself has direct consequences which could lead to the breach of an environmental standard or regulation (in practice, stockouts in safety and environmental categories are rare)

Operational

The failure or stockout itself leads to a loss of production or other economic loss to the business

Non-operational

The effect of the failure or stockout is limited to the expense of repair and obtaining parts.

The RCS decision diagram leads from the analysis of stockout consequences to an appropriate stocking policy for that part.


3.2 Can the Requirement be Anticipated?

Some spare parts requirements, such as those arising from breakdown maintenance, are inherently unplannable: parts fail at random, without any obvious signs that a breakdown is

about to occur. However, some requirements can be anticipated:

- Parts needed for planned overhaul or replacement routines which occur at regular intervals regardless of the equipment's condition
- Parts which are subject to condition monitoring, where components or equipment are checked and replaced if a failure is about to occur

Spare parts usage which can be anticipated is often known as a *dependent demand*.



Parts are replaced or overhauled at fixed intervals if there is some characteristic *life* after which their reliability deteriorates rapidly. Planned preventive maintenance is scheduled to replace or overhaul the component regardless of its condition at regular intervals which are determined by the life. If intervals are based on convenient calendar time intervals, so part requirements can be planned even if the time between requirements is shorter than the part lead time.

“One of the most significant changes is a move to on-condition tasks”

One of the most significant changes brought about by RCM is a move away from the planned, preventive replacements of a second generation maintenance system towards on-condition tasks. These involve checking the condition of a component and overhauling or replacing it *only* if its condition is unacceptable. This causes problems for procurement, since we do not know whether a part will be needed until the results of monitoring are available. Nevertheless, using the RCS rules it is often possible to avoid holding stocks on site.

3.3 What Stock Holding is Needed?

If it is not possible to anticipate a spares requirement (and therefore avoid holding parts), RCS then asks how many parts must be held to support maintenance and production. RCS recognises that 100% availability is an unattainable ideal. Before calculating the stock requirement, the RCS analyst needs to specify a performance standard which depends on the consequences of a stockout:

Category	Performance Standard
Increased Risk	<i>Minimum availability of hidden function</i>
Safety/ Environmental	<i>Maximum rate of stockout (stockouts per year)</i>
Operational	<i>Minimum through-life costs</i>
Non-operational	<i>Service level</i>

In many cases the stockout has a direct impact on operations (this is generally true even if the equipment failure has safety or environmental consequences). RCS uses the technique of *through-life costing* to determine the spares needed.

3.4 Through-Life Costing

A stockout has operational consequences if lack of a spare part leads to increased costs over and above the cost of obtaining a spare. In this case it is possible to find a balance between the cost of holding spare parts and the losses incurred if a spare is not available when it is needed. A stockout may cost money in several ways, including:

- Extended downtime or reduced output leading directly to lost sales
- Penalty clauses for late delivery
- Cost of overtime to make up lost production
- Lower process efficiency or higher raw material costs
- Poor product quality, leading to returns, rework and poor customer impression

On the other hand, there are costs associated with holding spare parts. As well as the cost of purchasing the initial stock there are continuing expenses while the parts are held:

- Purchasing, systems and administration costs
- Deterioration (shelf life)
- Maintenance and repair while parts are in the storeroom

Traditionally these expenses have been bundled into a single *holding cost* which is a fixed percentage of the part's purchase price. The idea is to spread the stores administration costs over all the lines in stock. It works well for fast-moving parts, but the costs of holding slow-moving items vary widely depending on their physical size, shelf life and maintenance requirements. The optimum spares level is a balance between total holding costs and the cost of stockouts: high downtime costs are incurred if stocks are too low, but holding the spare parts is expensive if the level is too high.

A second and more serious disadvantage of the holding cost method is that it always recommends a single spares holding without taking any account of the fact that, while the benefits of an increased spares holding are felt over a period of time, the cost of purchasing stock is felt immediately. The problems of this approach can be seen by considering three examples:

Example 1: Pre-commissioning

A new chemical process relies on a magnetically driven pump. If the pump fails, the process is shut down at a cost of about £500 per hour. The manufacturers estimate that the pump will fail catastrophically about once every three years, but it is not possible to predict when it will fail. If a spare were not available, the manufacturer could supply a new unit within five days, but nevertheless they recommend that one spare pump should be held on site at a cost of £40,000. Are they right?

Example 2: Stock Review

A cost review in steel plant has valued the engineering spares inventory at £50M, of which parts worth £10M have never been used over the plant's 10 year life. Two spare gearboxes, priced at £30000 each, are part of the non-moving stock. The consultants carrying out the review recommend that they should be sold as scrap. Are they right?

Example 3: End of Plant Life

An offshore installation uses water injection to maintain pressure in an oil reservoir. The pump bearing costs £5000, and current policy is to hold 2

spares because of the 3 month part lead time. One spare has just been used, making a total of 5 over the 18 years of operation. The computerised purchasing system recommends that a new spare should be ordered, but the platform has a remaining life of only 3 years: should the part be ordered to bring the stock back to 2, or should the stock be allowed to run down?

In the first example the decision is whether or not to invest in a spare pump. If we buy a spare, it will cost £40,000 *immediately*. On the other hand, if we do not, the business will suffer because of production downtime *in the future*. So the overall decision does not just involve a balance between costs and benefits, but also depends on the *timing* of these costs.

The question is different in the second case: *is it worth disposing of spares that we already have, or should they be retained in case they are needed?* We may recoup a little of the parts' value if we sell them; however, if the parts are then eventually needed, the cost of lost production could be far greater than the scrap value.

Finally, example 3 looks at another aspect of the spare parts problem. There is clearly no point in reaching the end of plant life with a full complement of spare parts (although this is exactly what the traditional approach would suggest). The question here is: *how should stocks be run down toward the end of plant life?*

“The traditional approach fails because it does not answer the most fundamental question: is it worth buying a spare part?”

In all of these cases, the traditional approach fails because it does not answer the most fundamental question of all: *is it worth buying a spare part, and if so, how many should be purchased?*

The through-life method answers the stock analyst's real question: should we spend money now (by buying stock) in order to secure lower downtime costs in the future?

This suggests that we can look at the inventory decision as an investment. If we buy a spare,

we immediately incur purchase costs. During each year we incur further expenditure as a result of purchase and repair of spare parts, maintenance of parts in stock and a risk of process downtime waiting for spares. We can add together all of this expenditure over the plant lifetime to determine our total expenditure. Initial outlay is lower if we buy no spares, but the costs of downtime are higher. Buying more spares increases the initial investment, but may reduce subsequent costs. The spares holding which gives lowest overall expenditure will be chosen. The versatility of this method means that it can be applied to a wide variety of decisions including stock disposal and situations where plant lifetime is limited. Since these methods relate directly to questions of investment and disposal, we have found that they are particularly powerful in justifying these decisions to managers who are responsible for purchase authorisation.

4 Vendor Stocking

The total number of spares needed to support, say, eight identical installations is generally less than eight times the requirement for each individual installation.

“We do not transfer risk”

Vendor stocking, or any other centralised inventory, should be a "win-win" for both parties involved because they can split the benefits of this economy of scale. But how do you know whether a vendor stocking contract is worthwhile? If we hold parts ourselves, our company holds the parts and bears the risk of substantial downtime if parts are not available. Vendor stocking contracts transfer responsibility for part stocks, but when evaluating the contract it is essential to remember that we do not usually transfer the *risk*.

Using the through-life model, the answer is to evaluate the vendor contract against the local inventory alternative. Again, the alternative which results in a lower net present value is preferred.

5 Implementing RCS

A Reliability-centred Spares review can be undertaken either before buying spare parts for a new asset or when an asset has been in service

for some time. In either case, the review team should include the RCS analyst and representatives of maintenance and production functions.

A review of a complete engineering inventory — perhaps consisting of several thousand items — would be a long and costly exercise. Not all items have equal value, nor are the effects of a stockout the same for each line. It follows that a review of an existing inventory should begin with the most significant items.

These are:

- Where a stockout has serious operational consequences
- High value items
- Where safety or environmental integrity could be compromised

Pareto's principle is borne out here: a small number of items are responsible for a high proportion of inventory value and could result in particularly large production losses if a part were not available when needed. It follows that, if we are concerned with financial consequences, the fastest payback is ensured by analysing the most significant items first.

6 The Benefits of RCS

The immediate and most obvious benefit of applying RCS to critical spares is that stock levels are based directly on the requirements of maintenance and operations. Because the method is driven by consequence analysis, it achieves these requirements with the optimum investment in spare parts, typically saving 30% to 60% on inventory value while complying with production, safety and environmental requirements.

The method has human benefits in addition to the technical and financial improvements achieved:

- Improved communications between engineering, production and inventory personnel
- Improved understanding of inventory and maintenance systems requirements
- Clearer and more profitable relationships with suppliers