RPR Problem Diagnosis

A problem investigation and diagnosis method for IT people

August 2009

Advance Seven Limited
Melville House
High Street
Dunmow
Essex CM6 1AF
01371 876805

Table of Contents

Table of Figures

Introduction

The increasing business need for sophisticated IT solutions means that many systems have become very complex. This is leading to a growth in the number of IT problems where it is difficult to identify the causing technology – the so called “grey problem”. Investigation of such problems is often haphazard and progress is slow.

Rapid Problem Resolution (RPR) is a structured problem diagnosis method that is both fast and reliable, and ideally suited to the investigation of grey problems. The method is IT-specific, built upon tools and techniques that are readily available in the corporate IT environment. This makes RPR a natural progression in the development of IT problem diagnosis and easily adoptable by support teams.

This paper outlines the need for RPR, how it works, when it should be used, how it fits into the IT department and the results that can be achieved.

Terminology

For readability purposes, the word ‘problem’ is used throughout this paper as a general term rather than as per the ITIL definition. A capitalised first letter has been used where it is necessary to emphasis the ITIL meaning of ‘Problem’.

Target Audience

This paper is intended to help IT operations managers understand what can be achieved through a structured approach to problem diagnosis. It should be particularly useful to problem managers who are looking for ways to deliver proactive problem management, and IT support people who need to widen their involvement beyond that of their own technical silo.

Version

This paper is based on v2.02 of the manual *RPR Problem Diagnosis Method: Definition and use.*

The Grey Problem Issue

The majority of problems that are passed to 2nd and 3rd line technical support teams are investigated in a straightforward manner. The nature of the problem or an indication from a monitoring system identifies the failing component and the problem is allocated to the correct technical support team. The following figure places these problems in quadrant Q1, and the bulk of support work falls in this area.

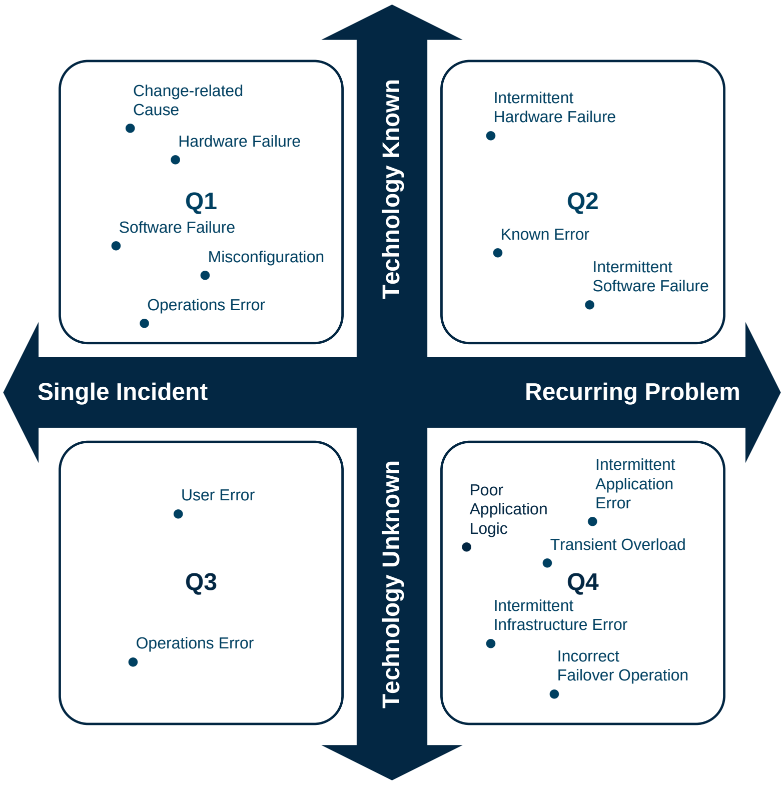

Figure 1 Problem categories

An intermittent response-time or error problem is not so easily dealt with due to its transient nature. Not only does the cause sneak under the radar of monitoring systems, but investigation often starts after the problem has passed making it impossible to use many of the tools available. The result is a recurring problem where the causing technology is unknown (Q4 in Figure 1); the grey problem.

Grey Problem Characteristics

Because the causing technology is unknown, a grey problem will bounce between Technical Support Teams as each in turn produces evidence (often in the form of a health check) to prove that their technology is not to blame.

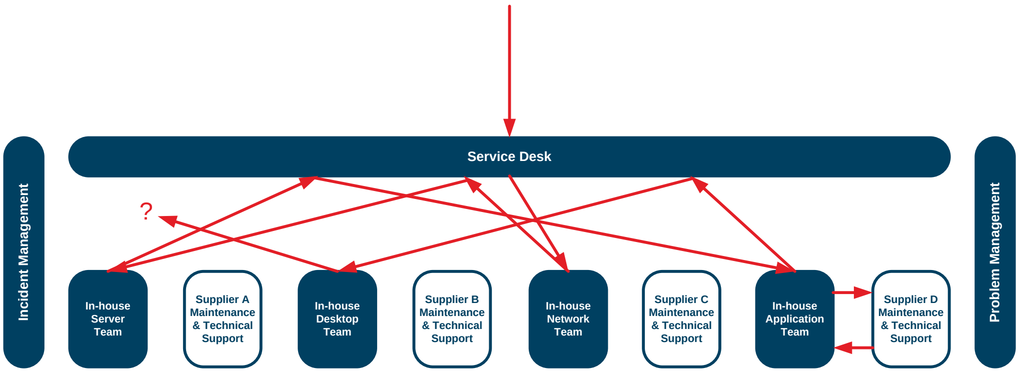
Service Desk
Incident Management
Problem Management
?
In-house Server Team
Supplier A Maintenance & Technical Support
In-house Desktop Team
Supplier B Maintenance & Technical Support
In-house Network Team
Supplier C Maintenance & Technical Support
In-house Application Team
Supplier D Maintenance & Technical Support

Figure 2 Routing of grey problems

There is a link between recurring, grey and non-urgent problems in that:

- Many recurring problems are grey
- Many recurring problems are at the priority 3 & 4 end of the list
- The difficulty in tackling grey problems and their perceived low priority means that they often remain open for months maybe years[1]

Other typical characteristics of a grey problem situation are:

- An ever-growing number of people becoming involved in the problem
- Long meetings or conference calls to discuss what might be the cause of the problem
- Support people shying away from becoming involved
- Repeated changes with no clear reason or objective

[1] Advance7 has recently been helping a customer with a problem that they have had for more than 10 years

Grey Problem Consequences

Because many grey problems are treated as mid to low priority they are often put on the "too difficult to deal with" pile. The consequences for IT are serious as failure to deal with grey problems causes:

- An ever growing backlog of problems
- A "fog" that hinders the investigation of other more urgent problems
- A growing pool of problems that escalate into Major Incidents as patterns of use and business priorities change
- A heavy workload for hard-pressed IT support people and managers
- Wasted IT budget as money is spent on poorly targeted upgrades
- Barriers to integration due to concerns about the stability of component systems
- Loss of confidence and satisfaction with the IT department
- Executive pressure to outsource IT services

The consequences for the business are no less serious with:

- Loss of actual and potential sales
- Reduced customer satisfaction
- Higher costs as the business adjusts to accommodate the problem
- Higher IT staffing costs
- Impaired share price

Method-based Troubleshooting

The State of the Art

There is a growing interest in problem resolution methods, and there are lots to choose from. It's long been recognised that the fastest and most reliable way to solve a difficult problem is to determine root cause first, and most of the methods available follow this principle.

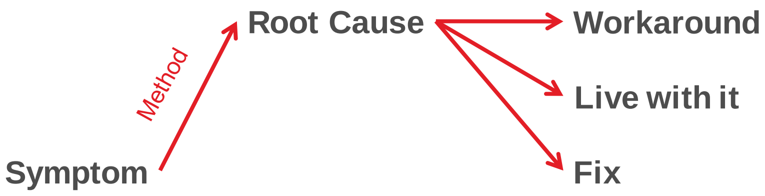

Figure 3 Steps to problem resolution

Some methods have sprung from the IT industry, but many of the front-runners are adaptations of manufacturing and business problem resolution methods. These methods have some common issues:

- Most methods are generic and so are not designed to take advantage of the logic and tools available to us in the IT industry
- Many of the “methods” are actually just processes with no supporting IT techniques making it difficult for IT people to run the process
- Many methods require that the root cause is already known as one of a list of possible causes, which makes them heavily reliant on experience and product knowledge
- Some methods are based on the recreation of the problem in a lab environment, which is time consuming, can be expensive and rarely works for grey problems
- Many methods rely on statistical analysis which often fails with grey problems due to their intermittent nature and transient causes
- Some methods rely on trial and error

Pattern Searching Methods

Perhaps the biggest problem is that most methods use a pattern-searching technique to identify root cause. As the name suggests, pattern-searching relies on establishing patterns to link pieces of information. For example, if a CRM system is slow only on a Monday morning it must be caused by an event that only happens on a Monday morning. The success of pattern-searching methods is highly dependent on perfect information. Unfortunately, in a busy complicated IT department it’s rarely possible to get perfect information and so in practice pattern-searching methods can suffer reliability problems.

Continuing with the example, it may only be after months of investigation that it’s discovered that actually the CRM system is slow at other times it’s just that those events haven’t been recorded.

A phrase that is often heard when using a pattern searching method is “.....that doesn’t explain why.........”.

Continuing again with the example, possible reasons might be dismissed with the explanation, “It can’t be synchronisation with the email system because that happens every hour so it doesn’t explain why we only get the problem on Monday mornings”.

The “that doesn’t explain why” issue can also make pattern searching slow as much time is spent trying to determine the link between symptoms that may or may not have the same root cause.

RPR Overview

RPR was designed from the outset to solve IT problems and is heavily influenced by software engineering techniques. From this starting point RPR avoids the shortcomings suffered by other methods because:

- RPR is evidence-based and so is completely reliable
- RPR makes full use of the IT tools that are available in every business
- It’s a mature method with a clearly defined core process and supporting techniques making for easy integration into the IT department
- RPR requires no pre-conceived idea of the cause of the problem, in fact such thoughts are positively discouraged
- The method predominantly uses non-disruptive techniques and so there is minimal business impact
- RPR doesn’t require recreation in a lab environment, or even testing outside of normal working hours
- RPR is a mature method that has been developed and refined over a period of 20 years and is fully documented in a 150-page manual.

RPR Principles

RPR is built on four key principles:

- Identification of root cause is best achieved by focusing on one symptom
- Reliable root cause identification (RCI) requires diagnostic data as proof
- Successful RCI requires diagnostic data capture in a live environment
- RCI requires a precise match of diagnostic evidence with a user’s experience

The first point is of particular interest as traditional problem investigation is often based on the idea that by identifying a collection of “related” symptoms it’s easier to find root cause (a pattern-searching approach), whereas more often than not the complete opposite is true.

RPR Limitations

The RPR principles dictate three important factors:

- The use of diagnostic data captured at the precise time of a single instance of the problem makes RPR extremely reliable
- The requirement to capture diagnostic data at the time of the problem means that RPR can only be used for ongoing and recurring problems
- RPR’s focus on diagnostic data means that it can only determine the technical root cause

The second point illustrates why other non evidence-based methods, including those based on pattern-searching, remain important tools in problem investigation. The final point means that RPR can't be used to determine causes such as management process, training or contractual issues which can be tackled by other methods.

The RPR Method

Process

RPR has two parts, the first of which is a process.

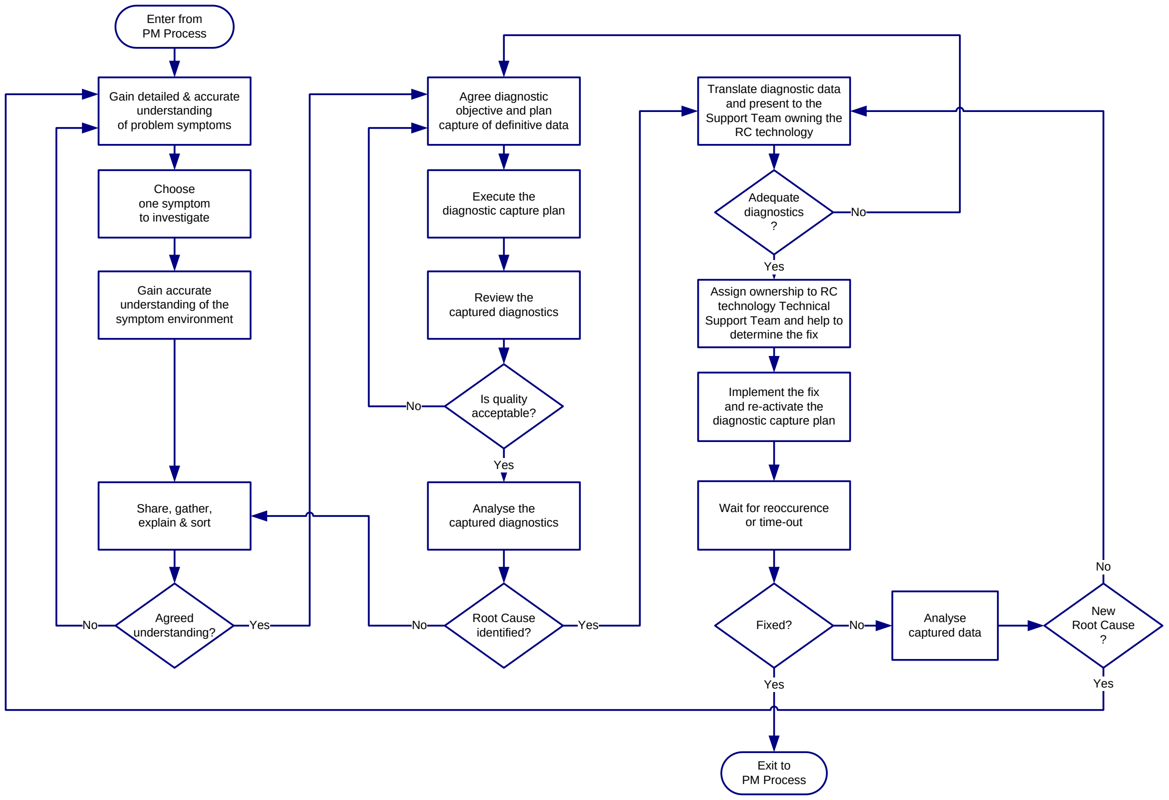
Enter from PM Process
Gain detailed & accurate understanding of problem symptoms
Choose one symptom to investigate
Gain accurate understanding of the symptom environment
Share, gather, explain & sort
Agreed understanding?
No
Yes
Agree diagnostic objective and plan capture of definitive data
Execute the diagnostic capture plan
Review the captured diagnostics
Is quality acceptable?
No
Yes
Analyse the captured diagnostics
Root Cause identified?
No
Yes
Translate diagnostic data and present to the Support Team owning the RC technology
Adequate diagnostics ?
No
Yes
Assign ownership to RC technology Technical Support Team and help to determine the fix
Implement the fix and re-activate the diagnostic capture plan
Wait for reoccurence or time-out
Fixed?
No
Yes
Analyse captured data
New Root Cause ?
No
Yes
Exit to PM Process

Figure 4 RPR Core Process

Each step of the process guides the practitioner on matters of what to do and when. The RPR Manual provides a whole section on each process step, and each section contains:

- Purpose
- Objective
- Activities
- Supporting Techniques
- Checklist
- Next Step

The checklist is particularly important as it acts as a set of quality gates to ensure that no activities are skipped.

Step 3 – Understand Symptom Environment		
Checklist	**✓**	**Notes**
1. Major application components identified		
2. Major infrastructure components identified		
3. Data flows determined and noted		
4. Diagnostic data sources identified and noted		
5. Purpose of application understood		
6. All user contact information gathered and recorded		
7. Time symptom occurs recorded		
8. Function of transaction showing the symptom noted		

Figure 5 Example of a Step Checklist

This means that RPR meets an IT Service Management goal of a standardised process that is:

- Repeatable
- Scalable
- Controllable
- Measurable

Six Sigma practitioners will also readily recognise that RPR is quite closely aligned with DMAIC[2] principles.

Supporting Techniques

A complaint often heard regarding any process is that it defines **what** should be done and **when** but does not show **how**. RPR deals with this issue by providing a set of supporting techniques that show how the objectives of each step of the process can be achieved. The supporting techniques are very practical and relate directly to current technologies. This is one of the reasons why RPR is constantly evolving; although the process rarely changes new supporting techniques are introduced with each new release.

[2] DMAIC – Define, Measure, Analyse, Improve and Control

When to Use RPR

RPR is not intended for initial investigation of a problem but for use when traditional techniques have been exhausted.

Escalation

3rd Line	• Check for overload • Investigate with diagnostic tools • Liaise with supplier technical support • Review design
2nd Line	• Check monitoring systems • Review recent changes • Check for misconfiguration • Check supplier knowledgebase
1st Line	• Check for misuse • Check for known error
Service Desk	• Check for user error • Check for PC fault

Figure 6 Typical problem investigation

RPR should be used when the standard investigation tasks (such as those above) have been completed and have been unsuccessful in resolving the problem.

There are also two tell-tale phrases that indicate that the IT team is struggling with a grey problem:

"We are just going to try one more thing"

"We made some changes and things have improved a bit
but we still have some work to do"

At this point a structured approach should be adopted. Whether RPR is chosen or another offering, a method-based approach will be faster and more reliable.

RPR in Use

Most organisations have recognised the need for a function to take ownership of a problem from end-to-end i.e. a responsibility that spans the technical teams. In an ITIL environment that function is Problem Management which is responsible for coordinating the problem investigation activities of the individual technical support teams.

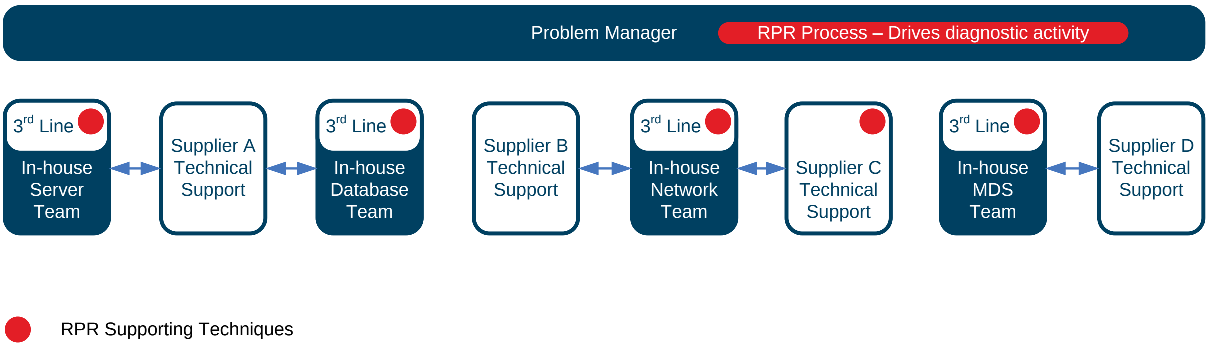
Problem Manager
RPR Process – Drives diagnostic activity
3rd Line In-house Server Team
Supplier A Technical Support
3rd Line In-house Database Team
Supplier B Technical Support
3rd Line In-house Network Team
Supplier C Technical Support
3rd Line In-house MDS Team
Supplier D Technical Support
● RPR Supporting Techniques

Figure 7 Integrating RPR into the IT department

RPR fits neatly into this model, providing a process to help the Problem Manager drive the problem investigation and supporting techniques to help the technical support teams achieve the goals of the process steps.

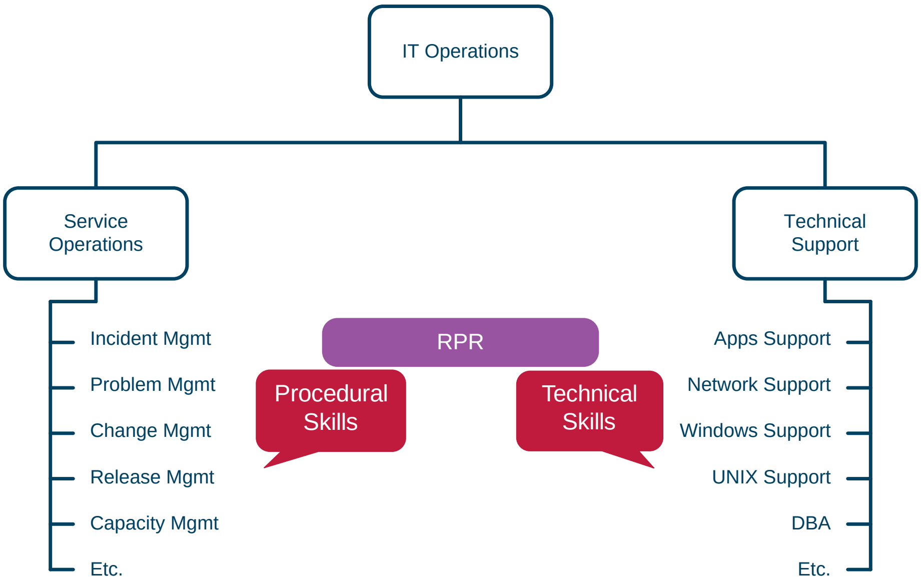

Figure 8 Bridging the gap, unifying the team

IT departments organised around the ITIL framework can find it difficult to create a cohesive team to investigate problems. Service-focused people tend to be driven by process, procedures and SLAs, whereas technology-focused people tend to be more pragmatic but operate only within their technology silos. This creates a gap between the Service Operations and Technical Support groups.

By providing a common framework for problem investigation, RPR bridges the gap to bring the two sides together to produce a cohesive team with clearly defined objectives, and ultimately faster problem resolution.

Results

Since 1990, Advance7 has delivered an RPR-based problem diagnosis service called REACT. By analysing one hundred of the most recent REACT projects information regarding savings through use of the RPR has been determined. The REACT projects were undertaken for organisations from 1,500 to 200,000 users. The majority of the problems were performance related and this is typical for grey problems.

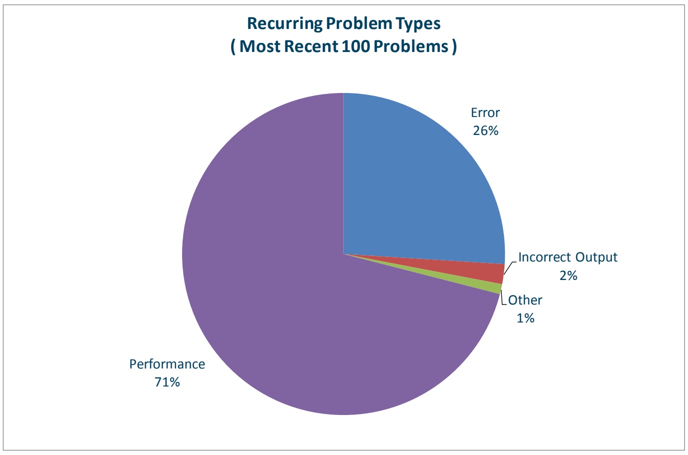

The figures to the right show how RPR cut the number of man days of effort used by IT support to determine the root cause of a recurring problem. The categories along the x-axis relate to the percentage reduction in IT support man days achieved i.e. 0 to 10%, 10 to 20%, etc. The values on the y axis show how many of the 100 problems fall into each category.

The projected minimum reduction in IT workload is 47%.

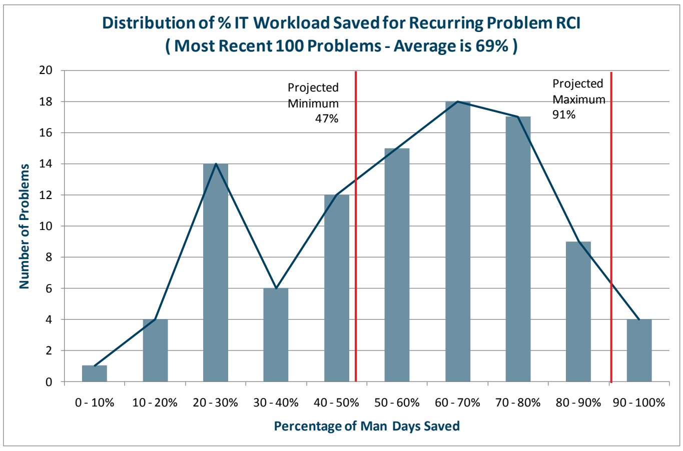

The figures to the right show how RPR improved productivity. Loss of productivity occurs when an IT service is slow or unreliable, and cutting the time taken to resolve recurring problems improves productivity across the business.

The chart above shows the distribution of reductions in problem duration achieved.

The projected minimum reduction in problem duration is 64%.

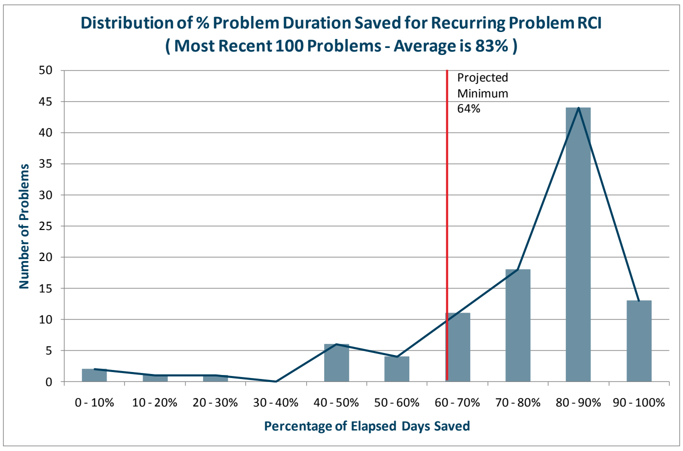

By determining the proportion of these problems that required an upgrade to resolve them it is possible to indicate the expected rate of capital expenditure for recurring problems.

Only 10% of problems required an upgrade to resolve them. To allow for an element of variation between projects the norm for upgrades is taken as double the value shown in the chart i.e. 20%.

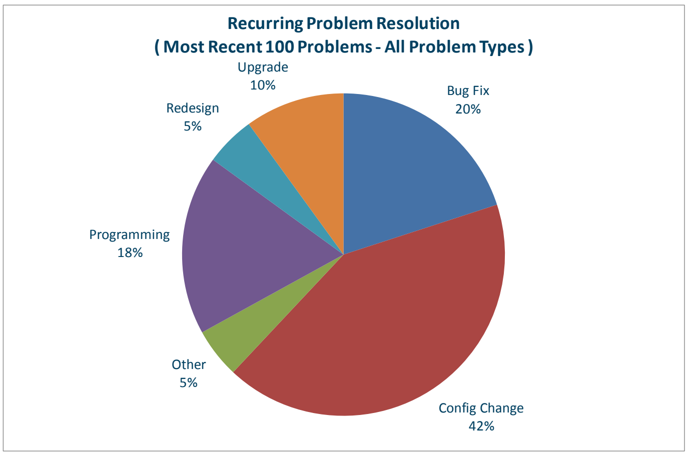
Recurring Problem Resolution
(Most Recent 100 Problems - All Problem Types)
Upgrade 10%
Bug Fix 20%
Redesign 5%
Programming 18%
Other 5%
Config Change 42%