A Definition of Proactive Problem Management

A paper describing how problem management can prevent service impacting issues

An itSMF UK White Paper, January 2013

itSMF UK
The IT Service Management Forum

CONTENTS

EXECUTIVE SUMMARY

There is a great deal of discussion about the merits of Proactive Problem Management, but there seems to be some confusion and disagreement regarding the activities involved.

What actually is Proactive Problem Management, and how does it differ from its reactive counterpart? What can it do for an IT department and a business in general?

In this paper we outline a set of activities that make up the practice of Proactive Problem Management and the benefits that these can bring. The information has been gathered from more than 20 problem managers through online forums and telephone discussions.

INTRODUCTION

Purpose of this Paper

This paper is intended to:

- Provoke discussion about Proactive Problem Management
- Define terms and language so that future discussions about the subject share the same terminology
- Provide guidance to those wishing to introduce Proactive Problem Management into their organisation

Objective of this Paper

The objective of this paper is to provide a draft definition of the practice of Proactive Problem Management. We don't consider this to be an end point, but it does represent the thoughts on the subject of a group of experienced IT professionals.

Method

Through online forums, telephone discussions and face-to-face meetings we have sought the opinion of a wide range of people that includes practicing problem managers from many industries, independent consultants and other IT professionals.

Of course, such wide consultation will occasionally lead to contradictory opinions. In such cases we have adopted the view of the majority, tempered with what we[1] believe is realistically achievable.

We have then cross-checked the information against the *Service Operations 2011* manual.[1] Although ITIL is a little vague on the subject of Proactive Problem Management, we have tried to ensure that the guidance we provide here is aligned with the framework.

[1] In this case the term 'we' refers to the itSMF UK Problem Management Special Interest Group.

GENERAL DESCRIPTION

Proactive Problem Management (PPM) means identifying, resolving and preventing problems before they cause service impacting incidents.

Reactive Problem Management deals with the investigation, diagnosis and resolution of problems that have already been detected because they have had a recognised impact on services.

PPM differs from its reactive counterpart by addressing three areas not otherwise covered, namely:

- Proactive Detection – the recognition of patterns of events (service impacting or not) that suggest an underlying problem
- Problem Prevention – identification of opportunities to prevent future problems
- Pre-emptive Action – identification of threats in critical situations
- First-fault Diagnosis – identification of the root cause of a problem upon its first occurrence

Techniques can be learnt, tools can be purchased and procedures developed to assist in the early detection of problems. Because the subsequent investigation and resolution actions will be the same as would be undertaken for any other problem, developing this capability is relatively straightforward.

Developing the ability to identify possible future problems is a greater challenge for a problem management team. The rate of development is dependent on the breadth of technical skills of the team, their relationships with other IT teams and the level of Reactive Problem Management experience. In this area, there is also some overlap with other IT teams, and so some work is needed to define roles and responsibilities.

OBJECTIVE OF PROACTIVE PROBLEM MANAGEMENT

Proactive Problem Management means identifying and resolving issues prior to service disruption, so that we:

- Avoid incidents from occurring in the first place
- Reduce IT support workload caused by repeated low priority incidents

ITIL states that the objectives of PPM are to, "... improve the overall availability and end user satisfaction with IT services". This is not a particularly useful definition as these objectives are the same as those of Reactive Problem Management discipline.

REACTIVE PROBLEM MANAGEMENT

To provide some contrast, it's useful to list the activities that relate to Reactive Problem Management, which are:

- Problem detection
- Problem record logging and management
- Categorisation
- Management of prioritisation
- Management of investigation and diagnosis of production problems
- Management of investigation and diagnosis of problems during pre-production testing
- Managing entries in a KEDB
- Liaising with support teams in the application of fixes and workarounds
- Liaising with the Change Management function
- Management of the resolution of a problem in a live service

- Management of the resolution of a problem in a pre-production service
- Closure of problem records
- Review of closed problems to learn lessons
- Recording lessons learnt in a knowledgebase and CSI initiatives

It's important to note that dealing with Priority 3 and 4 problems is a Reactive Problem Management activity, whereas identifying trends in P3 and P4 incidents to identify problems is a Proactive Problem Management activity. We highlight this particular point as some people consider the resolution of Priority 3 and 4 problems to be a proactive activity.

PROACTIVE DETECTION

Trend Analysis

Incident trend analysis is cited as a PPM activity in ITIL Service Operations. However, three types of trend analysis are common:

- Incident – the analysis of recovered incidents
- Monitoring System – the review of alerts generated by support team monitoring systems
- Knowledge Articles – the review of article usage statistics that may indicate an underlying problem
- Human Detection – the instinctive recognition by technical or service operations people that something isn't quite right

Incident trend analysis should include P3 and P4 incidents, not just high priority events.

In the following subsections we look at these areas in more detail. Trend analysis strictly refers to the change in a metric over time, and some of the guidance below should be described as pattern analysis.

Incident

Technical Causes

Incident trend analysis should be carried out to identify frequent occurrences, common failures and fragile CIs. The analysis should include the trending of low priority incidents as these are often a sign of bigger problems to come. The evidence can be supplemented by studying CI availability figures.

Although it may seem counter-intuitive, users should be encouraged to raise incidents with the service desk. This will ensure that we get a clearer picture of the true state of our services, which helps with PPM. Of course, all teams must be seen to be investigating problems for the users to remain interested in raising incidents.

A review of incidents should be carried out periodically (weekly, monthly or quarterly). Incident analysis can be computerised by reporting:

- Incidents by service – to identify the problematic applications and underpinning infrastructure
- Incidents by location – to identify problems with shared infrastructure
- Incidents by CI make and model – to identify problems with hardware or software
- Incidents by business unit – to identify problems associated with functional groups of systems or particular transactions
- Incidents by date/time – to identify transitory periods of overload
- Incidents by user – to identify possible training issues

It's wise to consult with the appropriate support group before raising a problem ticket.

Poor incident categorisation can be a problem and this issue is dealt with in *Poor Incident Categorisation* starting on page 17.

Development and Transition Issues

If software releases from a particular group of developers always seem to result in an increased number of incidents for associated services, it would be a good idea to flag the issue. This would allow problem managers to review pre-production testing of new releases with Release Management and the development team.

Identification of these types of issues can be achieved by:

- Pattern Analysis – what is the current level of the issue compared to other development teams in the organisation or versus an industry benchmark?
- Trend Analysis – is there are growing issue with a particular team?

Infrastructure Support Issues

Similarly, issues relating to installations, changes and upgrades carried out by technical support teams or suppliers can be analysed in the same manner as those for Development and Transition Issues.

Application Support Issues

The issue here is slightly different to that described under *Development and Transition Issues* in that these relate to ongoing BAU support. There are often multiple application support teams, perhaps aligned with business units. Through analysis of incidents it's possible to check that each team is as effective and accurate as the next. To see how this is possible see *Appendix B – Pattern Analysis Example* on page 24.

Training Issues

Analysis of incidents by application or service where the closure code indicates a user error or related cause can be undertaken to identify user training issues. Other possible solutions are better online help systems and encouragement to self-help through use of Internet resources.

Eliminating training issues will reduce the number of calls to the service desk and increase user productivity.

Monitoring System

Automated monitoring would generally involve tools and systems installed and managed by application and infrastructure technical support teams. These monitoring systems generate two types of event:

- Alert – indicating that a threshold has been crossed
- Alarm – indicating the failure of a component

An alert may be generated where, for instance, a server CPU load has exceeded a threshold, the rate of network errors is excessive or response times are not meeting the target time. The condition that causes the alert will often not result in a call to the service desk. However, alerts are very often an early warning of future problems.

Business systems are often built for high availability by duplicating components and providing a failover mechanism. Therefore an alarm may not impact the service delivered to the user.

Ensuring that the technical support teams act on alerts and alarms is a valid PPM activity. Inevitably this also means stepping on the toes of support managers. For example, if monitoring is in place to alert a server team to a low disk space condition, then proactive work involves following up on repeated alerts where no remedial action is noted such as log file clean up, ordering more disks, etc.

Many technical monitoring systems provide the ability to automatically generate a 'ticket' (Incident or Problem Record) when an alarm or alert is generated. Although this seems ideal, in reality such configurations take a significant amount of management. Common issues are:

- Thresholds are set too low and so a high number of false tickets are generated

- Thresholds are set too high and so no tickets are generated at all
- Planned downtime generates false tickets
- Configuration changes cause false tickets or monitoring fails so no tickets are generated

Therefore, a more realistic approach is to ask each technical support team to produce a consolidated exception report which can then be manually reviewed by the problem management team.

The technical support workload required can be reduced by improved tooling. This is a very rich area at this time with a new generation of performance monitoring, data gathering, symptom matching and automatic fix installation tools becoming available; these in themselves could be considered as proactive problem management tools.

Knowledge Articles

The usage statistics of a self-help system are also a useful source of trend information. By analysing the high usage of an article that describes a workaround or fix it may be possible to identify an underlying problem that should be investigated.

Human Detection

Technical or other service operations people (such as incident managers) also have a role to play in Proactive Problem Management. Such people may spot an issue that happens repeatedly in the course of their daily work, and this should trigger Proactive Problem Management by asking for a problem ticket to be opened.

Where an expert has a 'hunch' or feels instinctively that something is going on, this should be followed up and should be a valid trigger for Proactive Problem Management. Experts' hunches have been shown by psychological research to be the subconscious gathering together of seemingly unrelated information, from which highly sophisticated expertise distils the information without the expert consciously realising it is going on.[2]

Human detection also extends to a customer's hunch that may be expressed during a service review or other such gathering. All hunches will need to be assessed objectively through some form of triage process. In the case of a customer's hunch, the problem management team will need to be sure that the issue relates to a problem with an existing system and not a concern relating to the lack of a desired capability. Issues identified by this type of activity would therefore either be investigated in the normal way, or may be fed into the CSI process.

A further trigger in this category for proactive activity is as simple as a request from somebody senior to investigate something.

PROBLEM PREVENTION

Retrofitting Fixes

Future problems can be prevented by retrofitting corrective actions identified from one failure (as a part of RCA process) to CIs of the same type. If we find the root cause of a problem and implement a fix, there will be other CIs where we can implement the same fix to prevent a similar issue before it occurs. Even though the CIs may not be causing problems at this time, changes in workload pattern or use could trigger the problem.

Resource Management

Consider the situation where an incident has been caused by low free disk space on a server. We can set in place a regular housekeeping task to check the free space on a server and so avoid future incidents.

Although the monitoring of resources such as free disk space will be the responsibility of a capacity management team, or the technical support team that owns the CI, Problem Management has a role to play in the identification of this specific resource as a known cause of problems, and ensuring that steps are taken to mitigate the risk of related downtime.

Risk Register Monitoring

A Risk Register can provide a good means to identify potential problems. Periodic reports can be run to check for risks that could cause IT service impairment or failures.

Related to the Risk Register is the identification of Fragile Artefacts, those being CIs with support issues such as an OS that is no longer supported, or a piece of hardware that is out of warranty and for which no service contract has been arranged.

Project Issues Log

A project manager will often maintain a log of issues that have arisen during the project. Some (probably not many) problem management teams treat these issues alongside other reactive activities. The Issues Log should be reviewed by Problem Management to ensure that problems that exist in pre-production are not migrated into the production environment.

Vendor-driven Proactive Notification

It would be possible to avoid problems by working with your vendor to get proactive notifications based on your infrastructure's hardware profile. This is a service that has been offered by suppliers such as mainframe and storage vendors for around 30 years, and is typically based on a 'call home' mechanism. Wider availability would be a benefit to the IT industry.

Activities could include swapping out hardware when it's at 80% of its life expectancy; or when it's detected that redundant hardware is failing. It would then be possible to work with the vendor and swap out all hardware of the same type with similar manufacturing dates.

PRE-EMPTIVE ACTION

Techniques can be used to identify threats to the organisation without a specific incident having yet happened. In critical situations where;

- there is a stress event coming up which needs to go well
- detection of an incident would come too late for anything to be done
- if an incident happens it is important to have planned the restoration activities in advance

then potential problems and likely causes can be generated from the knowledge and experience of technical and business staff.

The synthesis of subject matter expert experiences may highlight potential problems that have not ever happened before, but could do given a new circumstance or starting condition. These likely causes can have pre-emptive actions taken to reduce the threats, and contingent actions planned against the potential problems.

FIRST-FAULT DIAGNOSIS

The concept of first-fault diagnosis is quite simple; the ICT systems are instrumented to gather enough information so that the root cause of any problem can be determined the first time it occurs.[3]

This may seem like an unachievable utopia. In reality, many components of an ICT system already have this capability. For more than 40 years, some operating systems and software applications have included a constantly running trace table that can be used to pinpoint the cause of a crash. More recently we have seen the availability of a new generation of Application Performance Monitoring tools that have a goal to provide first-fault diagnosis capability.

A one-off incident may be recovered and yet the IT team may be unable to determine the technical root cause. Examples of such issues are a potential source of justification for first-fault capabilities.

Working with technical or performance management teams to build a business case for the deployment of first-fault diagnosis capabilities is a PPM activity. The problem management team should have available the evidence that shows how a first-fault tool could significantly reduce the frequency of a certain type of problem, and so is in the best position to justify any investment.

This is an advanced form of PPM, and one that deserves its own white paper. We have included first-fault diagnosis here for completeness.

INHIBITORS

There are issues that work against Proactive Problem Management.

Incoming Data Quality

People

If there are inexperienced people working the Incident and Problem Management processes, then there is a lower chance of getting the quality of incoming information required to deliver PPM.

Process

Incident Records:
The quality of the Incident Records directly affects the ability to implement PPM – this is a primary source of data about the state of the IT estate, and if it is poorly completed, full of irrelevant noise or incomplete, the ability to undertake the PPM described here will be impaired. The Incident Management process will need to ensure that the quality of incident records is consistently good.

Problem Records:
The records of problems are a core area for research, and again they need to be an accurate representation of the problems that have been handled. The Problem Management process will need to ensure that the quality of problem records is consistently good.

Continual Service Improvement:
The activity of the CSI team should build on the discoveries and work of the PPM team. A poorly implemented or non-existent CSI team will diminish the benefits of PPM.

Risk Register:
If there is no Risk Register process, then there is an inherent risk of hardware or software becoming production and service risks without anyone noticing.

Product

The products which are involved in Proactive Problem Management will provide alerts and alarms for the incident managers. The alarms clearly kick off activity to understand the alarm and manage it – the alerts will give an early warning of problems to come. Out of date or non-existent monitoring products will impair proactive activities.

Poor Incident Categorisation

Overcoming Poor Categorisation

Due to problems with controlling the quality of incident records, it's sometimes necessary to scan the free text description within them to spot trends. Such analysis improves with experience and the technical capabilities of the problem manager.

Proactively Managing Poor Categorisation

Since the value of Proactive Problem Management can only be realised when there is good quality data to work on, it is worth first ensuring that the data around each incident is an accurate representation of the incident.

Incident and Problem Management Maturity

There are some prerequisites to effective Proactive Problem Management:

- Mature Incident Management is important to allow successful incident trend and pattern analysis
- Strong assimilation of Reactive Problem Management activities into the IT department is needed so that problem management principles are thoroughly understood

On the issue of the maturity of Reactive Problem Management, an honest assessment of the level of activity and capabilities relating to the activities defined in *Reactive Problem Management* on page 7 would be a good starting point.

Resourcing is a big problem. At the time of writing, many problem management teams are totally consumed with the review of major incidents and high priority problems. They have no time to investigate low priority problems, let alone start a Proactive Problem Management initiative.

One way to overcome the resourcing issue is to ring-fence a group of problem managers that will only deal with PPM. Successful ring-fencing in this way requires a great deal of discipline, an unwavering commitment from senior managers and a clear understanding of the value proposition of PPM.

A variation on this approach is to temporarily allocate one or more problem managers to PPM activity, tackle obvious opportunities, and then return the people to the reactive team. By repeating this process a problem management team might progressively reduce the number of problems occurring and so get to spend more time on PPM.

Hero Culture

An organisation may have a very strong 'hero' culture, in which case staff just love a good incident to get all excited about. The proactive activity is viewed as boring and of little value. A high-profile incident provides a very visible opportunity to demonstrate technical skill.

Some large organisations remain focused on technology rather than service, and so troubleshooting success is measured by the technical support team's ability to recover from incidents, even if the incidents are recurring. Problem management in such an environment is difficult; Proactive Problem Management would probably have little success.

APPENDIX A - IMPLEMENTATION

The following table provides some guidance regarding the order in which each PPM activity should be introduced. There is a large list of possible activities, and trying to tackle them all could swamp a problem management team. A pragmatic way to deal with this issue is to spread the activity in phases throughout the calendar. For example, in Q1 problem managers could look at trends associated with Development and Transition, in Q2 those associated with Infrastructure Support, and so on.

The benefit derived from these activities will vary from organisation to organisation. Nevertheless, we have provided some guidance here based on the knowledge and experiences of the authors.

Step	Activity	Level of Benefit	Prerequisites
1	Trend Analysis: Incident – Technical Causes	H	Accurate incident management categorisation.
2	Trend Analysis: Incident – Development and Transition Issues	M	Logging of development and transition incident records in the ITSM.

Step	Activity	Level of Benefit	Prerequisites
3	Trend Analysis: Incident – Application Support Issues	H	Participation by application support staff in the incident management process, at least use of incident records.
4	Trend Analysis: Incident – Infrastructure Support Issues	H	Participation by infrastructure support staff in the incident management process, at least use of incident records.
5	Trend Analysis: Incident – Training Support Issues	M	Clear categorisation of user errors by the Service Desk, and accurate service references.
6	Trend Analysis: Monitoring System	H	Well maintained monitoring systems that produce reliable alerts and alarms.
7	Trend Analysis: Human Detection	H	The identification of a group of technical staff whose judgement in the status of the IT environment has been found to be reliable.

Step	Activity	Level of Benefit	Prerequisites
8	Retrofitting Fixes	L	Mature change and release management. A CMDB so that appropriate target CIs can be identified. Good pre-deployment test procedures.
9	Resource Management	M	Strong interpersonal relationships with the technical support and capacity management teams.
10	Risk Register Monitoring	L	The risks included in the Risk Register must include quite detailed entries relating the IT systems and components. The tool used to maintain the register must be capable of categorising the risks.
11	Project Issues Log	L	A project issues log must be maintained and the tool used to maintain it must be capable of categorising the issues.
12	Vendor-driven Proactive Notification	L	A vendor's capability to provide the required notification.

Step	Activity	Level of Benefit	Prerequisites
13	Pre-emptive Problem Management	L	Knowledge of the techniques needed to think of potential threats in a structured manner.
14	First-fault Diagnosis	H	Availability of suitable tools and systems. Instrumentation of applications.

Level of Benefit – Expected number of issues avoided relative to the mean of all of the steps above.

APPENDIX B – PATTERN ANALYSIS EXAMPLE

The following example was provided by Michael Hall:

[Let's imagine we] *have three application teams supporting three different businesses. Each has a set of applications, let's say three each to make it simple. Incident data for each is tracked and reviewed regularly by Problem Management. Over the last three quarters:*

- *Team 1 is averaging 3 sev 1 incidents, 15 sev 2 incidents, 100 sev 3 incidents*
- *Team 2 is averaging 2 sev 1 incidents, 25 sev 2 incidents, 100 sev 3 incidents*
- *Team 3 is averaging 0 sev 1 incidents, 50 sev 2 incidents, 100 sev 3 incidents*

So what is going on? Why does Team 3 have so many sev 2 incidents? This is where Problem Management can raise a proactive problem to go and do some digging, applying their analysis skills to get to the bottom of it. It's worthwhile, because, if they could get the number of sev 2 incidents down to the average of the other two groups (20), [it would prevent] *30 incidents per quarter, or up to 120 each year.*

AUTHORS

Paul Offord

Paul Offord has had a 35-year career in the IT industry that includes roles in hardware engineering, software engineering and network management. Prior to founding Advance7 in 1989, he worked for IBM, National Semiconductor and Hitachi Data Systems. Paul is now the Development Director at Advance7.

Paul is a Certified IT Professional and a Fellow of the British Computer Society.

Steve White

Steve White began his career in Computer Systems Maintenance in 1989 in various roles that have included Hardware on-site support, software Incident and Problem Management, backline software support and the deployment of RCA processes into a variety of companies and industries. He is currently a Senior Consultant with Kepner Tregoe, specialising in the integration of clear thinking processes in IT Support.

ACKNOWLEDGEMENTS

We are greatly indebted to members of the ITIL Problem Management forum on LinkedIn for their input, particularly:

Peter Dodd
Pramesh Iyer
Patrick Collings
Prasanna Lakshmi
Eric Sambraus
Kundan Yadav
Manish Singh
Rahul Sahasrabudhe
Tim Jobson
Martina Todorova
Carl Bindels
Pradeep Bains
Boitumelo Machogo
Michael Hall
Dan Skwire
Lorie Wilson
Shannon Mcilroy
Swapnonil Bose
Peter Senese
Mark Dickinson
Dahud Khan
Bruce Burghgraef
Joe Gallagher
Stephen Maclean

ADDITIONAL INFORMATION

ITIL definition of Problem Management:

1. *David Cannon and David Wheeldon (2007),* ITIL Service Operation*, The Stationery Office, ISBN 978-0-1133-1046-3*

2. *Gary Klein (2004),* The Power of Intuition: How to Use Your Gut Feelings to Make Better Decisions at Work, *Currency, ISBN 0-385-50289-3*

3. Daniel Skwire et al. (2009), First Fault Software Problem Solving: A Guide for Engineers, Managers and Users*, Opentask*, *ISBN 978-1-9067-1742-1*

itSMF UK
150 Wharfedale Road
Winnersh Triangle
Wokingham
RG41 5RB
United Kingdom

Phone: +44 (0) 118 918 6500
Web: www.itsmf.co.uk

Revision 1.08