



Root Cause Analysis Concepts and Best Practices for IT Problem Managers

By Mark Hall, Apollo RCA Instructor & Investigator

A version of this article appeared in the April 2010 issue of Industrial Engineer magazine

Many IT departments struggle with the negative business impacts of recurring problems, and many also struggle with how to proceed with formally investigating each major problem. The risks are so significant that IT-related problems are estimated to impact the US economy to the tune of around \$60 billion per year for software errors and around \$140 billion per year for info/data security breaches.

So how are solutions typically being developed for these problems? There seems to be a reliance on the statistical analyses of industry trends to drive IT-related solutions. Information from actual IT problems categorized by type, or assumed cause and solutions, are recommended based on data trends exhibiting the highest percentages or greatest threats.

This approach leads to one of the most common reasons problem solving is often ineffective -- solutions based on categories do not necessarily address the actual causes of a given problem as effectively as solutions that are specifically designed to control the actual causes of the problem. Categories don't cause problems – causes do. Categories are titles created so that similar things can be grouped together, so by nature are generic on many levels. Generic, categorical-based solutions fail at a much higher rate than do solutions that control specific causes of a defined problem.

The problem management component of the Information Technology Infrastructure Library (ITIL)¹ framework sets the stage for each organization to adopt effective problem-solving strategies that will improve the quality of internal and external IT services. Recently, IT organizations have begun to realize that some of the same problem-solving best practices long used by other disciplines such as quality, safety, maintenance and reliability are adaptable and scalable to IT. Successful IT problem-solving organizations are increasingly implementing formal root cause analysis (RCA) within their ITIL problem management structure.

So, What is a Problem?

¹ ITIL – Information Technology Infrastructure Library. <http://www.itil-officialsite.com/home/home.asp>

In ITIL terminology, problems and incidents have unique definitions for IT-related events. An IT 'problem' is the unknown cause of one or more incidents, often identified as a result of multiple similar incidents.² An incident is an unexpected event which negatively impacts the quality of IT service³.

In more simplistic terms a problem is simply a deviation from a goal. An organization must have clear business goals that should be known to everyone inside. Otherwise an organization cannot effectively identify problems and efficiently allocate resources to solve them.

Problems are relative to specific organizational goals. An ad hoc approach to identifying organizational problems can lead to unstructured approaches to solving them, which can cause the organization to waste time, money, resources and potentially experience recurring problems.

Best Practice: Understand and communicate all organizational business goals.

Establish Threshold Criteria to Help Prioritize Problems

RCA is a tool that is used to solve problems that are preventing an organization from achieving its business goals.

Threshold criteria are used to recognize and quickly define problems when they occur so that the overall response can also be right-sized to the problem at hand.

For each major business goal, target, or KPI, it is critical to establish formal written criteria that represent when the magnitude or direction of a deviation is far enough from the original goal that it warrants a formal investigation to prevent its recurrence. These threshold criteria need to represent the point where the cost of investigating and solving a problem are less than the cost of the problem. For example, you may have a goal to prevent unscheduled network outages during business hours that last more than 15 minutes. If an event occurs during business hours that interrupts normal network service for less than 15 minutes, then the service is restored as quickly as possible. However, relative to the stated goal it was not severe enough to warrant investing time and resources for a formal RCA.

Business goals can also be prioritized in order of importance so that the logic of recognizing and defining a problem is aligned with the relative importance of each goal to the organization.

² itlibrary.org

³ http://www.itlibrary.org/index.php?page=Incident_Management

Best Practice: To ensure your formal RCA process is deployed in a timely and cost-effective manner, make sure to use RCA to solve problems that have prevented, or could prevent, achievement of specific business goals.

Customize Threshold Criteria

Within each level of the organization you should have goals that are inter-related so everyone is working toward achieving the same overall corporate goals. At each level, establish threshold criteria to define what is considered a problem at that particular level. Those criteria will trigger RCAs only at the appropriate level of the organization. If the threshold criteria are set slightly lower for each level in the organization, this works toward “compartmentalizing” problems before they become problems at a higher level in the organization.

For example, a threshold for a corporate goal pertaining to cost control on a new infrastructure upgrade project might state that any unplanned cost expenditure greater than \$500,000 is a problem – at a corporate level. At a regional level, though, threshold criteria might be set at \$400,000 for the same goal and \$300,000, or even less, for the business unit, and so on.

Threshold criteria are critical to ensure expenditures on RCAs are aligned with the overall organization goals. These criteria should be fluid and change with shifting business climates or with individual organizational direction/situations. Formal threshold criteria give direction and purpose to your formal investigation programs.

Best Practice: Create threshold criteria at all levels of the organization.

What is Root Cause Analysis?

RCA is a structured process designed to help people understand the causes of past problems for the purpose of preventing recurrence. It is step-wise and structured so that it can be consistently applied to different problems at different times by different people. Solutions will only be effective if they act on the specific known causes of a defined problem.

Best Practice: Adopt a robust formal RCA method. Avoid simplistic, generic or categorical methods like the 5 Whys’, Fishbone diagrams or “brainstorming” because they are not as effective at identifying the causes of problems as the leading modern RCA methods are. Instead, choose robust, principle-based methodologies that dig deeply into the causes of problems and support known causes with factual evidence.

Root Cause Analysis Steps

Effective problem solving has four primary steps. These steps must be followed in sequence. Jumping around or skipping a step will ultimately lead to failure in solving the problem

1. Define the problem
2. Create a causal understanding of the problem
3. Identify solutions that act on known causes of the problem
4. Implement the best solutions

Best Practice: To ensure your formal RCA process is deployed in a timely and cost-effective manner, make sure to use RCA to solve problems whose impacts meet or exceed your threshold criteria.

Now we'll go through the RCA steps in detail.

Define the Problem

- What is the problem?
- When did it occur?
- Where did it occur?
- What is the significance or impact (\$) of the problem on the business?

Thoroughly defining the problem includes formally pinpointing exactly what the problem is, when and where it happened, and the impact that it had on the business. "Problems" are defined by the most significant consequence of a "major incident," or as the business goal, target, KPI or metric that was not achieved as a result of an unwanted event. Often people confuse defining the problem (i.e., the bad thing you don't want to have happen again) with the causes of the problem. The simplest way to identify the "problem" is to state the threshold criteria that was compromised by the event. For example, "unscheduled network outage greater than 15 minutes during business hours."

Properly defined problems ensure everyone working on the problem is pointed in the same direction. There is no misinterpretation of exactly what the "problem" is or where the search for causes needs to begin. As well, a properly defined problem fully describes the impact it had on an organization so that proper attention and resources are given to the people working on the problem. People often do not account for the full impact (i.e., actual and potential costs) problems have on their organizations. In turn, this results in an insufficient level of priority or resources for a given problem, and ultimately a failure in finding effective ways to prevent the problem from recurring.

Best Practice: Standardize the format for defining problems so everyone does it the same way each time.

Create a Causal Understanding of the Problem

The key elements to fully understanding why a problem occurred include:

- Identifying specific causes of the problem
- Creating certainty by validating causes with evidence such as verifiable, factual data
- Organizing by showing logical relationships between causes

After the problem is properly defined, conducting an RCA requires the analyst to develop an understanding of as many causes as possible. In this step, identify the specific causes of the problem, ensure that you have factual evidence to verify each cause and organize all the causes in a format that shows the logical relationship between them. In other words, what caused what to happen?

When done effectively using a robust formal RCA method, this step produces a visual roadmap showing why the problem occurred. An understanding of all the causes of a defined problem is the platform from which effective solutions are identified.

Best Practice: Present the analysis in a visual format (i.e., cause and effect diagram). Complex relationships between problems are best understood when people can “see” how things happened. Improved understanding will ultimately lead to identifying more and better solutions.

Identify Effective Solutions

To be effective, solutions must:

1. Eliminate or control known causes of the problem
2. Be within the “owner’s” control
3. Be cost effective, feasible and practical. Earn “buy in.”
4. Avoid causing other problems that can be anticipated

Use your understanding of the causes of the problem that were revealed in the previous step to identify effective solutions. Look for ideas that attack each of the known causes of your problem. Evaluate the ideas using the four criteria above and create a short list of solutions ideas.

Best Practice: Reject solution ideas that do not meet all four of these criteria. Only rely on solutions that will be effective. To be effective solutions must meet all 4 solutions criteria.

Implement the Best Solutions

Your organization must now implement the solutions identified in the analysis. The risk of a problem recurring is not reduced until the solutions are implemented. The most critical factors for successfully implementing new solutions include:

- Identify the problem owner
- Assign solution implementation to specific people with due dates
- Track solution implementation and effectiveness

Best Practice: Implement multiple solutions for each problem. Try to control as many causes of the problem as possible. The more solutions implemented, the greater the certainty of preventing recurrence if one or two solutions fail to work properly.

Causes of Ineffective Problem Solving

There are numerous reasons that people fail to effectively prevent problems from recurring. Robust RCA methods will prevent analysts from falling into these common traps:

- No problem definition
- Jumping to solutions
- Solutions not controlling actual causes
- Categorization
- Focusing on blame
- Stopping the analysis too soon
- Searching for a single “root cause”
- Relying too heavily on software, historical databases and tools

Best Practice: Adopt an RCA method that prevents ineffective problem solving

So What is the Root Cause of a Problem?

The burning question – how do I know what the root cause is? Shouldn’t an RCA tell the analyst what the root cause was so it can be fixed? An RCA method is not a computer model that tells you which among all the known causes of the

problem is the one thing that needs to be fixed. A good problem solver must seek to understand all the causes of a defined problem first, validate the authenticity of those causes with factual evidence, and then find ways to control as many of the actual causes of a defined problem as possible given time and money constraints. Every known cause that has a unique solution acting on it is, by definition, a root cause of the problem at hand.

Identifying the root causes of a problem should be done after you choose your best solutions, not the other way round. This does not mean you pick your solutions first then work backwards to identify the causes and root causes. If the step-wise sequence of RCA is not followed ultimately the problem solver will miss seeing something critical to successfully preventing problem recurrence. Root cause labeling is a product of a completed RCA, not the driver. The most robust RCA methods allow the analyst to identify the causes his/her organization has the best chance of controlling, rather than the “method” telling the analyst what the “root cause” was.

With this approach, investigation teams and their managers can have discussions about the best solution options given a suite of equally effective opportunities, level of risk, and potential consequences of a repeat event.

Best Practice: Focus on implementing effective solutions not labeling a root cause.

To envision these best practices in practical application, following is a real-life client example.

Case study: World Class Results in Enterprise Problem Management

The Boeing Enterprise Problem Management (EPM) program is a world-class RCA effort. The EPM Group has 100 percent level of effectiveness in preventing problems after all solutions are implemented. Their RCA program has helped reduce the overall cost of delivering IT services within the company. Because reducing the time to close an RCA is critical in reducing organizational vulnerability and threats, it is significant that they reduced the time to close the RCA of major incidents by 42 percent in their first five years.

Best Practices for Enterprise Problem Management

RCA should be thought of as program, not a “tool.” A program requires a plan that is implemented and regularly improved. Without a plan, running an effective RCA effort is like trying to score a touchdown without knowing how to play football.

Two key elements to successful program start-up and implementation include:

- 1) Doing the “right things,” which we call critical factors for success or program efficiency (i.e., elements of the program).
- 2) Doing those things “well,” which we call results (i.e., program effectiveness).

Best practice: Implement an RCA program as early as possible in your ITIL implementation and create specific goals for the program so its purpose is aligned with the goals of your organization.

Best Practices for Enterprise Problem Management

The Boeing Enterprise Problem Management program established important KPIs for itself in order to measure its IT RCA performance. The key KPIs include: the time it took to complete each RCA, the number of RCAs completed by the team per year, the percent effectiveness of the solutions, and the cumulative savings created by the program based on the prevention of past problems.

Time-To-Close Metric

For Boeing, the Time-To-Close metric is one of its most important KPIs, measuring the time it took to get solutions implemented after a major incident occurred. The length of time that it takes to implement solutions is the period the company was vulnerable to the impacts of problem recurrence – so the focus on this metric ensures that everyone was focused on finding effective solutions and getting them in place as promptly as possible.

Example: Aligning RCA Performance Metrics With Business Goals

Here is an example of aligning your RCA program goals and KPIs with business goals.

- 1) Business goal: maintain production levels about 70 percent
KPI: Number of days production fell below 70%
- 2) RCA program goal: Reduce production delays related to IT service interruptions by 50%
KPI: Number of production delays related to interruptions in IT service

Improving Time-To-Close Metric

Because the company is vulnerable to the impacts of recurrence after major incidents occur, Boeing undertook a number of initiatives to improve its ability to close out an RCA and reduce its exposure to risk. These included:

- Problem Resolution Board - A board of senior people that can support the RCA effort from a high level, approve expenditures and remove road blocks that are delaying solution implementation.

- Short- vs. long-term solutions – Separate solutions that can be implemented quickly from solutions that will take longer to implement. Implementation of all short-term solutions defines when an RCA is closed. Seeing long-term solutions through to completion becomes the responsibility of the problem board or committee.
- Virtual RCA – Use electronic submissions for evidence gathering, along with virtual, team-based tools that unite RCA team members from different regions in a cost-effective manner.
- Robust evidence collection – Ensure all evidence and relevant data are collected and catalogued as soon as possible after the recognition of a triggering event, including server logs and other data that is sensitive to time destruction.
- Action item tracking – Track all interim RCA-related action items in a formal tracking database.
- Participation / escalation process – Create a formal process to report when people who should be assisting the RCA effort are not. Direct this information to the problem board/committee for resolution.
- LEAN Value Stream Mapping – Use LEAN as a tool to improve the efficiency of your RCA program.
- Searchable RCAs – Format and store completed electronic RCAs so that other investigators can search. After a triggering event occurs, an effective initial step is to search the records for an RCA that may have already been completed on the same or very similar event. This may allow a team to quickly identify known causes and effective solutions.

Reactive vs. Proactive Problem Solving

While there's a common conception that RCA is only used to react to problems that have already occurred, RCA is also proven to be effective in proactively mitigating risk. And while reactive problem solving is critically important, evolving from reactive to proactive problem solving is a logical progression for an organization that has matured in its use of RCA.

Reactive problem solving involves analyzing problems from the past and implementing solutions that prevent their recurrence, which is accomplished with each RCA that is completed. Proactive problem solving can also involve implementing solutions for potential or hypothetical problems and from trend analysis of multiple RCAs.

RCA can be used to plan how a "system" or "process" should ideally function. RCA can also be utilized to enhance continuous improvement processes, demonstrate due diligence, and analyze positive events so they can be repeated.

Using multiple event analysis -- analysis of the common causes that played a role in multiple problems -- organizations can identify links between problems that might otherwise be undetectable using other analytical tools. This is a feature of a solid RCA method.

To summarize, essential best practices when implementing an RCA method for IT problem management include:

1. Based on your organization's business goals or scorecard metrics, develop threshold criteria to differentiate between incidents that do not require full investigations and major problems that do.
2. Precisely define each major problem and quantify the business impact.
3. Allocate adequate time and resources for each RCA that are commensurate with the impact and risk associated with each major problem.
4. Complete each RCA consistently using the same method and investigation process. This helps to ensure the RCA program is defensible and ready for independent auditing.
5. Don't categorize problems or their causes and then rely on those categories to drive the solutions. This inevitably leads to the identification of ineffective solutions.
6. Institute a rigorous validation process that uses evidence to verify causes and weed out uncertainty.
7. Based on a clear understanding of the causes of major problems, utilize the talents of the people who use, maintain and deliver IT services to help identify the best solutions. This is more effective than relying on an automated "tool" that prescribes solutions in a formulaic fashion.
8. Prioritize solutions based on criteria like cost, payback and ease of implementation so they can be rationalized against the business impact statement.
9. Develop solutions that are clear and descriptive enough to be successfully implemented by a third party, and to be effectively monitored.
10. Focus monitoring metrics on implementation timing and effectiveness of the solutions, and report regularly on your RCA program's successes.

As an Apollo Root Cause Analysis expert, Mark has particular experience serving clients in the IT field, as well as the natural resources sector, including oil and gas, engineering and forestry. Mark's thought leadership on IT problem solving has been published in BizTech magazine and showcased in presentations at itSMF US Fusion 2009, itSMF Canada 2008, and Business Analyst World Conference 2009. His best practices recommendations on risk management have been published in Professional Safety magazine. Mark is also a selected speaker for the itSMF US Fusion 2010 conference.