

Oracle On Demand Best Practice: Root Cause Analysis

An Oracle White Paper

January 2006

Oracle On Demand Best Practice: Root Cause Analysis

Intended Audience	3
Introduction	3
Using RCA In Oracle On Demand Services	3
Business Impact	4
Adapting RCA For The Oracle On Demand Business	4
Reactive Management	5
Proactive Management.....	6
Oracle On Demand RCA in Practice	7
Proactive Knowledge Transfer	7
Oracle On Demand RCA Example.....	8
Conclusion.....	10
Terminology.....	11
Appendix	13

Oracle On Demand Best Practice: Root Cause Analysis

Oracle best practices are based upon the IT Infrastructure Library (ITIL) framework, a de facto industry standard for IT service management. Root cause analysis (RCA) is a component of the ITIL Problem Management process.

INTENDED AUDIENCE

This white paper is intended for IT management and senior staff responsible for managing Oracle technology and applications.

INTRODUCTION

Oracle has established the global leadership position in on demand solutions. Oracle On Demand takes responsibility for managing the infrastructure, software, security, service levels, and IT governance; and offers customers unparalleled flexibility and choice. Delivery models range from optimized on site, to Internet based shared solutions, to highly tailored private environments. Regardless of the customers preference for delivery models, Oracle makes best practices available to Oracle customers by either advising and assisting them or by taking responsibility for deploying and managing their software.

Oracle best practices are based upon the IT Infrastructure Library (ITIL) framework, a de facto industry standard for IT service management. Root cause analysis (RCA) is a component of the ITIL Problem Management process.

This paper describes Oracle On Demand's adaptation of RCA to further improve on the ITIL best practices to provide customers a superior ownership experience.

USING RCA IN ORACLE ON DEMAND SERVICES

Oracle's unique scalable grid capability provides dynamic services for application provisioning and hosting. Oracle On Demand also champions the use of RCA to continue to improve service availability and quality. While RCA has been a standard practice in the manufacturing sector, its applicability to software management has been growing. Over the past few years, Oracle On Demand has been adopting RCA to improve service levels and provide an increasingly better customer experience over time. Enterprise software is very complex and comprehensive, deployed and used in a variety of configurations. At the present time, problems do still occur, either due to the software itself, or the use thereof. While continuing to improve the software, Oracle On Demand has implemented processes, or best practices to better manage failures to ensure that:

- They do not happen again, and that,
- New solutions do not have known or related failures.

Oracle On Demand’s extensive use of RCA as a part of its Problem Management process delivers on that promise.

Business Impact

Preventing repeated failures not only improves system availability, stability, and the overall end user experience, it also reduces the total cost of ownership. Every preventable problem causes unnecessary downtime, business loss, and additional overhead to request technical assistance. It is well known in the IT industry that most problems are rediscoveries. A rediscovery is where the same problem occurs more than once, and hence must be fixed repeatedly. As such, these problems are great candidates for “find one, fix many” opportunities. RCA is the process within Oracle On Demand’s offering that finds problems once, fixes them for many customers, and prevents them from occurring again to affect a positive business impact.

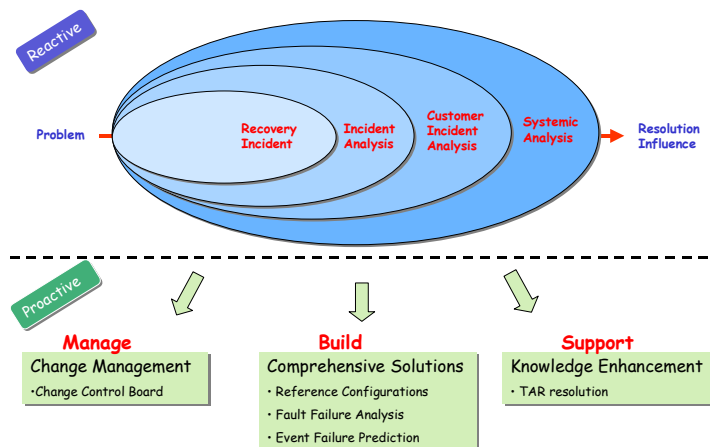
ADAPTING RCA FOR THE ORACLE ON DEMAND BUSINESS

RCA is an analytical methodology to find the root cause of problems. It recursively analyzes the causes of a problem until the root cause is found. RCA is a pragmatic approach. A cause is considered the root cause when there is an actionable solution to the problem that is cost-effective to implement and that has systemic impact to fix one and therefore fix all such problems. As such, RCA is a reactive methodology. However, proactive actions can be derived from such a reactive methodology.

Finding the root cause of any problem takes time. While finding root causes to incidents to “find once, fix many” is needed, incident management cannot be predicated on finding root causes alone. Some problems are incidental rather than systemic. Some problems tend to be customer-specific due to their specific system implementation. To provide on demand solutions, and in particular to provide proper responses to future occurrences, all incidents must be understood at different levels of detail and depth. Recognizing this, Oracle On Demand has divided its RCA methodology into multiple layers of analysis. The following diagram shows how Oracle On Demand implements both Reactive and Proactive Problem Management:

The following sections look at each of these areas in detail.

Root Cause Analysis – A Layered Approach



Reactive Management

This analysis, based on the nature of the incidents and the types of response needed, includes:

- Recovery incident categorization
- Incident analysis
- Customer incident analysis
- Systemic analysis

Recovery Incident Categorization

Recovery incident categorization provides the initial categorization of an incident before service recovery. Categorization is performed for each incident, and is based on a standard problem taxonomy defined by Oracle On Demand. This taxonomy reflects the dynamics of Oracle On Demand systems, with each incident cause being represented by taxonomic data. The intent is to understand the cause of each incident for proper responses during recovery and in some cases, for further actions.

Incident Analysis

When an incident occurs repeatedly, it is an indication of a deeper problem. The initial response of recovery and fix may not be adequate. Such incidents require deeper analysis and trigger the second layer of analysis - incident analysis. The intent is to better understand the cause of these incidents to improve on the responses.

Customer Incident Analysis

Repeated incidents can be due to specific customizations or setups that are unique to a customer. Systemic problems can also cause incidents to repeat across many customers.

Using taxonomic data, Oracle On Demand can track repetitive incidents for all customers to understand trends in the data. Based on this analysis, the next layer of analysis and problem defense – customer incident analysis - is activated to identify core issues within a customer’s system. Customer incident analysis is used to root out the repetitive problems for specific customers. The intent is to pay close attention and solve specific customer issues that may be very important or unique to the customer, but hidden as “noise” in the overall data trends.

Systemic Analysis

Systemic incidents appear as problem taxonomic data from the affected customers. These incidents have the same problem category, sub-category, and occurrence frequency spreading across different customers over time. Once this systemic trend is determined, the last layer of problem defense - systemic analysis - is triggered.

Incident types identified in the System Trends Reports are examined in detail. System runtime logs are often monitored, analyzed, and tested to find the root cause behind systemic problems.

To harness Oracle's expertise, Oracle product development organizations are often involved to assist in identifying root causes. The intent here is to determine *why* certain problems repeat themselves by "peeling the onion" and finding the root cause to fix many systems all at once.

Proactive Management

To ensure that known problems, such as product defects, system configuration and process issues, do not occur in new services, Oracle On Demand has extended RCA to include the following proactive methodologies in various stages of implementation:

- Reference configuration (RC) specification
- Fault Failure Analysis (FFA)
- Event Failure Prediction (EFP)

Reference Configuration

RC ensures that changes based on RCA to find one issue and fix many problems are systematically and automatically included in a base system image for new provisioning. This ensures that new customers do not experience known problems.

Fault Failure Analysis

Root causes in on demand solutions can involve operational errors and process design issues in change management, service delivery, recovery, and other critical service supporting processes.

FFA uses RCA results to analyze and evaluate these process faults by determining, on an a priori basis, how exceptions are handled for continuous process improvement.

Event Failure Prediction

Oracle On Demand serves a wide range of customers from large to small organizations. Some customers can put a highly demanding load on their systems at times, causing race conditions in the systems. These conditions are neither product defects nor process errors but they have the same result: service outage or interruption.

Providing spare system capacity and computational power may solve this problem, but not in a cost effective manner for our customers. Instead, Oracle On Demand has extended its RCA methodology using EFP to address this problem. EFP identifies key system status indicators and their specific states that predict the onset of such dynamic problems. EFP then generates automatic alerts for immediate intervention before the system degenerates into a failure state.

EFP is based on structuring and codifying specific RCA results into a set of causal event relationships that probabilistically, but accurately, model this type of dynamic, fatal system behavior. The key system status indicators are monitored

programmatically to find certain indicator states in advance of such failures for immediate intervention.

We are not aware of any vendors in the software industry that are able to offer these comprehensive capabilities to enable on demand solutions. Oracle has found that a superior ownership experience that balances service levels and total cost of ownership cannot be achieved or sustained if any one of the aforementioned service enablers is missing.

ORACLE ON DEMAND RCA IN PRACTICE

The layers of analysis for problem defense and management all contribute to system changes in the form of Request for Changes (RFC) submitted to a Change Control Board (CCB).

The CCB is comprised of several multi-disciplinary teams. These teams review and approve each RFC, creating work orders to manage change. The intent is to provide rigorous oversight to manage change effectively and productively.

Each RFC is reviewed across the different disciplines involved in delivering Oracle On Demand:

- How do the changes compare to imminent Oracle upgrades in cost and effectiveness?
- How extensive are the changes to customer systems?
- How much downtime is required?
- When can the changes be incorporated into a customer's scheduled periodic maintenance plan? Can the customer wait?
- Relative to the other planned changes and customer needs, how important are these changes in priority?
- Are there other similar changes that should be consolidated with the current RFC?
- Are there other changes that may conflict with the current changes?

It is only after these types of questions are analyzed by the CCB that an RFC would be approved unanimously and work orders created for execution.

Proactive Knowledge Transfer

Oracle On Demand hosts a wide range of customers across different industry sectors of various sizes and needs. Oracle On Demand is a fertile ground to learn about additional system behaviors, existing or new problems, and best practices through the RCA process.

Some of these principles are applied directly to Oracle On Demand customers' systems through the change process and some are proactively posted as knowledge articles provided by Oracle's Global Product Support team for off-line customers.

Based on the problem taxonomic data, certain RCA RFCs are algorithmically identified with signatures that uniquely characterize the problem types.

Similarly, customer Technical Assistance Request (TAR) data is assigned signatures and clustered into different problem categories consistent with the RCA problem taxonomy.

Certain RCA RFC signatures can then be mapped to the TAR cluster to determine which customers have experienced similar problems to those defined in the RCA RFCs.

The corresponding RFCs are then proactively sent to these customers to notify and assist them with similar problem prevention.

ORACLE ON DEMAND RCA EXAMPLE

The following customer example illustrates the power of RCA.

A customer's system is running out of file space for applications. The on demand service-supporting infrastructure automatically generates a "running out of space" alert to notify its operational staff that additional file space is needed. The operational staff immediately adds a data file to prevent a service outage incident.

In this case, the type of incident categorization and analysis has been built into the monitoring service-supporting infrastructure to provide not only the automatic alert, but also the immediate diagnosis and recommended corrective action as the first two layers of analysis and defense.

Through the capture and analysis of system and service data trends, if such alerts are observed repeatedly, it seems to indicate a more systemic problem than the customer's one-time usage and load pattern. The third layer of defense and analysis, Customer Incident Analysis is initiated in parallel to the alert to understand what is causing such incident recurrence in these customer systems.

- Are these customers doing a lot more than expected to cause file space shortage? If so, have their system capacities been sized correctly in the first place?
- Or are there hidden problems that cause the file space to be severely fragmented, rendering ineffective use of the file space to cause continued shortage?
- Or are there serious usage problems in how the customers interact with the applications that cause endless file output?

Some of these questions may be specific to only a certain set of customers. In these cases, Customer Incident Analysis can determine the root cause and recommend corrective actions to fundamentally eliminate this problem for these customers.

If these patterns are found across many customer deployments, and if the analysis does not point to customer-specific patterns, the last layer of defense and analysis, Systemic Analysis, is now initiated, to determine the systemic root cause.

For example, if the Customer Incident Analysis finds that the root cause of this problem for a number of customers is due to severe file space fragmentation, and the fragmentation is not due to any customer-specific setup, this is likely a systemic issue that requires systemic analysis to find this issue and fix many systems all at once.

To take this example further, the Systemic Analysis can determine that the file space fragmentation is actually due to a combination of the following two factors:

- Customers sometimes customize and perform certain application functions that require extensive data manipulation and thus enormous temporary data output to use up the allocated file space.
- The application system is smart enough to automatically extend the file space in multiples.

However, automatically extending file space in multiples also creates fragmentation. We can change the application to only extend its file space linearly to minimize fragmentation.

To provide the same benefits to other customers, the change can be incorporated into the next release of the system image in the form of a standard RC for all existing and new customers.

This certainly works to the extent that some customers are not unknowingly creating more “rogue” customized applications that require even more extensive data manipulation and thus more file space than before.

Like any service offering, resource utilization is commensurate with a specific service level agreement (SLA). Customers either purchase the next level of SLA for more resources if there is a real need for it, or consider how to customize some of the application actions carefully to use less space.

How do customers know how much space is reasonable?

With EFP, using the knowledge we learn from deep systemic analysis, we can begin to identify the key status indicators and their critical state values to predict the causality between certain “rogue” actions and the onset of file space full.

We can then provide additional space in advance of any problems as a courtesy to our customers while working with them to address the “rogue” application effects to stop this from happening again.

CONCLUSION

Software recovery and redress is a very critical component of on demand solutions. Along with its other supporting services and processes, Oracle On Demand's extended RCA methodology takes on demand to another level of service quality.

This paper contains detailed descriptions of the factors differentiating Oracle On Demand from competitive offerings and highlights, by example, one of the many best practices customers will be able to adopt by engaging with Oracle On Demand.

TERMINOLOGY

Change Control Board

The Change Control Board (CCB) consists of several multi-disciplinary teams. These teams review and approve each RFC, creating work orders to manage change.

Event Failure Prediction

Event failure prediction (EFP) is an Oracle extension of the RCA methodology. EFP identifies key system status indicators and their specific states to predict the onset of runtime dynamic problems. EFP then generates automatic alerts for immediate intervention before the system degenerates into a failure state.

Fault Failure Analysis

Fault failure analysis (FFA) is an Oracle extension of the RCA methodology. FFA uses RCA results to analyze and evaluate process faults by determining, on an a priori basis, how the processes handle process exceptions for continuous process improvements.

Incident

An incident is any event that causes or could cause an outage or disruption in service quality.

IT Infrastructure Library

IT Infrastructure Library (ITIL) is a de facto industry standard for IT service management.

Known Error

A Known Error is a problem for which root cause analysis and a workaround or a solution has been issued.

Proactive Management

ITIL defines Proactive Management as avoiding occurrence of incident by detecting potential issues or errors and triggering a change control and a change management process.

Problem

A problem is a set of incidents with common symptoms, or a unique incident for which the root cause is unknown.

Reactive Problem Management

ITIL defines Reactive Problem Management as detecting the root cause of past incidents and proposes workarounds or a solution as a resolution path.

Reference Configuration

Reference Configuration (RC) is an extension of RCA methodology by Oracle. RC ensures that changes based on RCA to find one issue and fix many problems are systematically and automatically included in a base system image for new provisioning.

Request for Changes

A Request for Changes (RFC), in this whitepaper, is a request made by the RCA teams based on RCA analyses to submit change requests to the CCB for acceptance to effect the requested changes.

Root Cause Analysis

Root cause analysis (RCA) is a component of the ITIL Problem Management process.

APPENDIX

ITIL & ITSM

ITIL provides a great blueprint to solve common problems of IT Service Management (ITSM) as shown in the following table.

IT Service Management	
<i>Service Support</i>	<i>Service Delivery</i>
Service Desk	
Incident Management	Service Level Management
Problem Management	Financial Management for IT
Change Management	Capacity Management
Release Management	Service Continuity
Configuration Management	Availability Management

ITIL Problem Management

ITIL defines the Problem Management process goals as:

- Reducing the disruption of business caused by incidents and problems resulting from errors in the IT infrastructure
- Avoiding the recurrence of incidents related to those errors.

Proactive management detects potential issues or errors in the infrastructure and triggers a change control and management process. Due to lack of standardization and automation, very few IT organizations can achieve proactive management. However, Oracle On Demand not only provides proactive management, it also extends the process to include:

- Reference Configurations (RC)
- Fault Failure Analysis (FFA)
- Event Failure Prediction (EFP)

Reactive problem management detects the root cause of past incidents and proposes workarounds or solution as resolution path. ITIL classifies the reactive problem management process in two areas: Problem Control and Error Control.

- Problem Control detects, analyzes, classifies, and records incidents from the other ITIL domains such as configuration management, incident management or capacity management.
- Error Control is the systematic monitoring process to provide workarounds or solutions to known errors until the errors are closed.

Oracle On Demand is working with Oracle service engineers and Oracle software development teams in a unique initiative to implement closed-loop multi-level RCA.



Oracle On Demand Best Practice: Root Cause Analysis

October 2005

Authors:

Charles Sum

Mukul Goyal

Eric Tranle

Oracle Corporation

World Headquarters

500 Oracle Parkway

Redwood Shores, CA 94065

U.S.A.

Worldwide Inquiries:

Phone: +1.650.506.7000

Fax: +1.650.506.7200

oracle.com

Copyright © 2005, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice.

This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, JD Edwards, and PeopleSoft, are registered trademarks of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.