Automated Change Management in the Data Center

No Limits Software White Paper #5
By David Cole

Table of Contents

Overview

The only acceptable number of unauthorized changes is zero.

In studying a number of "high-performing" IT organizations, the authors of *The Visible OPS Handbook: Implementing ITIL® in 4 Practical and Auditable Steps* (Behr, Kim, & Spafford, 2004-2005) found these organizations shared a common philosophy which states "the only acceptable number of unauthorized changes is zero". It is this culture of change management which allows these organizations to perform at very high levels as measured by:

- High availability (high MTBF and low MTTR)
- Change success rates over 99%
- Less than 5% time spent on unplanned work
- Server to system administrator ratios greater than 100:1

The authors found the following:

Automated change monitoring can reduce MTTR and increase system availability.

> ***Each of the high performing IT organizations had a common way to resolve service outages and impairments. They all realized that 80% of their outages were due to a change, and that 80% of their MTTR was trying to find out what changed. Consequently, when working on problems, they would first look at changes in the repair cycle. Evidence of this could be seen in the trouble ticketing systems of the high performers – inside the trouble tickets for an outage were all of the scheduled and authorized changes for the affected asset, as well as the actual detected changes on the asset. By just looking at this information, problem managers could recommend a fix to the problem over 80% of the time, with a first fix rate of over 90% (i.e., 90% of the recommended fixes worked the first time).***

For many organizations, modifying this first response represents a complete departure from their "reboot it and see if that fixes it" philosophy. For other organizations, using the change log in problem diagnosis is business as usual, but these organizations understand there are often cases where people have circumvented the change control process and, as a result, all changes may not be approved and logged. This white paper will discuss the importance of automating the change management process to detect changes made to the IT and facilities infrastructure which might have been made outside of the standard procedures. The result will be higher availability, less time spent analyzing outages, and less time spent on unplanned work.

Why Should Change Management be Automated?

If everyone adhered to the change management processes, we would see the following:

- All changes would be submitted to the change advisory board (CAB).
- The CAB would approve the change, which would then be implemented within the approved change window.
- There would be a post-implementation review of the change to determine whether or not the change succeeded.
    - If the change failed, the back out plan would be implemented.
- The change success rate would be updated

Unfortunately, adherence to these processes is not guaranteed. The use of change monitoring software automates the detection of changes, allowing us to determine all changes which were made, regardless of whether the changes were authorized or not. Organizations implementing change monitoring software are often shocked by the number of changes which are being made outside of the standard change management process.

Organizations implementing change monitoring are often shocked by the number of changes made outside of the standard change management process.

In fact, even organizations with strict adherence to change management procedures can benefit from change monitoring software. Change monitoring systems can be used to verify completion of tasks. For example, an authorized change ticket to install a new software application can be verified when the change monitoring software recognizes the new application has been installed. The same would hold true for the addition or removal of hardware, changing a network port, changing set points on a CRAC unit, upgrading the firmware in a UPS and so on. In addition, automatically tracking the changes greatly reduces manual data entry errors, preventing the new operating system installation from being entered as “Microsfot Windws 7”. The change monitoring software would retrieve the new operating system directly from the server and correctly enter the value of “Microsoft Windows 7”.

Requirements for Automated Change Management

There are two primary requirements for automating the change management process:

- Accurate and detailed asset configuration information
- Logged and detected change information

Let’s look at both of these requirements in more detail to explain their importance in the change management process.

Accurate and Detailed Asset Configuration Information

The first requirement for automated change management is accurate and detailed asset configuration information. Why is this important? Let's use a common issue to help explain the need for this information. A user calls into the help desk and says they can't access information provided through a web-based application. Let's put on our troubleshooting hat and go to work. We could start by asking the user basic questions such as

What server does the application run on?

We could then drill down into more detailed information such as

What operating system does the server use?

What service packs are installed?

What version of the database is installed?

What web server is being used?

We must be able to quickly and accurately determine the information we need to trouble-shoot an issue.

Of course, our user may or may not know this information. Worse, the information they provide may be inaccurate. It is therefore important for us to be able to access a source of truth to quickly and accurately determine the information we need. Without this information, we've greatly added to the time required to resolve the issue and, in doing so, decreased our availability.

With the asset information in hand and using our first response method of examining changes to determine what might have caused the issue, we take the followings steps:

1. Determine the server on which the web application runs

   Data Source – Detailed configuration information from the CMDB

2. Examine the changes on the server to see if any of these changes might have caused the issue

   Data Source – Event log

This methodology will typically allow us to successfully diagnose problems over 50% of the time. While this is a great first step, it only works if the change was on the asset itself. If there was a hardware, software, or network configuration change on the server, we might be able to successfully diagnose the problem. But what if the change occurred upstream from the server? What if the router into which we are plugged had a firmware upgrade or was moved to a different port on its upstream router?

It becomes clear that we need more information at this point. We need to know the next level of infrastructure – the network and power devices to which the server is connected. We might also need to know the rack in which our server resides. Where do we get this information? We go back to the detailed configuration data for the server to get its location, its power and network connections and any other pertinent information. Again, we want to be able to retrieve this information quickly and accurately.

If we can't find a change on the server which may have caused the issue, we now expand our search to changes made at the next level of infrastructure. We'll look at changes on the network and power devices to see if they've had changes which may have caused our issue. By expanding our search to the next level of infrastructure, we can increase our successful diagnosis of the issue to over 70% of the time!

The change log for the asset and the next level of infrastructure provides successful issue diagnosis over 70% of the time!

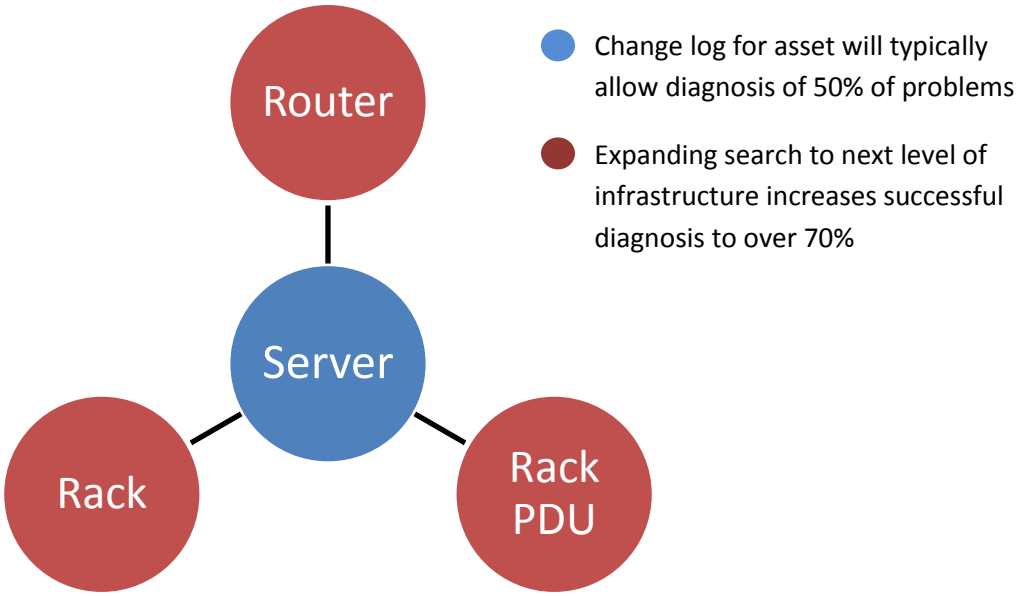

Remember the importance of having quick access to accurate data. But how does the data get into the system? In most cases, the entry of data into the configuration management database (CMDB) is done manually. This manual entry presents several problems.

The first is the time and cost to collect (and to later audit) the asset data. Data centers can contain thousands of servers, power, cooling, storage and network devices. Each of these devices has a relationship to other devices. It is a very daunting task to collect data about each piece of IT equipment, particularly if you are starting from scratch. While "visible" data (name, manufacturer, model, serial number, location) can be gathered reasonably quickly, retrieving the detailed configuration data we need (processor, storage, memory, network and power

connections, software, and so on) may involve logging into the server and using various tools to collect the information. This information, in turn, must then be manually entered into the CMDB.

The second problem with manual data entry is the accuracy of the information. In the Computer Associates technology brief *Striving to Achieve 100% Data Accuracy: The Challenge for Next Generation Asset Management* (Watson & Fulton, 2009), the authors point out the difficulty in maintaining the accuracy of manually entered information. The authors point out that

The error rate for manually entered asset configuration information can be 10% or higher.

> ***Manual tracking with pen and clipboard, or even spreadsheets is time consuming and highly error-prone. Organizations can typically expect a 10% error rate in manual data entry due to typing and transcribing errors.***

In a survey of the International Association of IT Asset Managers (IAITAM) members, respondents said an 85% accuracy rate for tracking IT assets was above average and that a 90-95% rate was exceptional. Consider for a moment the impact of a 10% error rate in a data center with 1,000 servers. As many as 100 of the servers will have inaccurate data recorded! Remember that this is the crucial information we need to be able to diagnose and resolve issues in the data center. If it is inaccurate, we've severely reduced our capabilities to quickly resolve issues.

Logged and Detected Change Information

The second requirement for automated change management is access to both logged and detected change information. Why is this important? If change is the cause of as much as 80% of outages, the ability to examine the change logs for an asset and its supporting infrastructure is a crucial tool in resolving problems.

The list of changes must be complete in order to properly diagnose issues. This means the change log must contain all changes which occurred – whether the changes were authorized or not. There are many instances where a manual change log will be incomplete or contain inaccurate information:

- Someone forgot to log the change
- The change has made but not yet entered into the change log
- Someone circumvented the process and made a change without authorization
- The change information was incorrectly logged
- Information from a vendor isn't entered into the system (a firmware upgrade as part of a PM, for example)

Automated change monitoring will recognize and automatically log ALL detected changes and will resolve issues with information being forgotten or entered incorrectly as well as detecting changes made outside of the standard change management procedures.

Summary

Change is a major cause of outages in the data center. Knowing this, many data center managers have changed their first response to problem solving. Instead of using intuition to attempt to determine the correct response or taking a more drastic step such as rebooting a server, the first step is to examine the change log for the asset and its supporting infrastructure. This typically allows successful diagnosis of the problem over 70% of the time.

In order to modify this first response, there are two key requirements:

- Accurate and detailed asset configuration information

  Manually entered data is prone to errors of 15% or more. Automated gathering of device configuration will greatly improve the accuracy of the information.

- Logged and detected change information

  It is important to have a complete record of all changes made – whether authorized or not. It is recommended that the systems are scanned for changes at least once a day. The ability to track both authorized changes and detected changes – changes made but not necessarily authorized – is key functionality needed to reduce MTTR and increase overall system availability.

Recommendations for Automated Change Management

1) Automate gathering of device configuration information
2) Employ change monitoring software to track both authorized and detected changes

About No Limits Software

No Limits Software is a leading provider of data center solutions, including asset management, capacity planning, and power and environmental monitoring. No Limits Software provides a unique solution by taking asset management to the rack unit. The RaMP (Rack Management Platform) solution improves system availability by automating the asset management and change management processes. It also eliminates the need for physical audits, dramatically reduces the time to find and repair equipment, and improves data center energy efficiency by providing accurate real-time power and cooling monitoring to improve capacity planning.

No Limits Software solutions are fully scalable and all data is fully accessible via published web services to allow you to easily integrate the solutions with your existing management platforms.

No Limits Software was founded in 2009 by industry experts in data center monitoring and management solutions. To learn more, visit www.nolimitssoftware.com or email info@nolimitssoftware.com.

Bibliography

Behr, K., Kim, G., & Spafford, G. (2004-2005). *The Visible Ops Handbook: Implementing ITIL(R) in 4 Practical and Auditable Steps.* IT Process Institute (ITPI).

Watson, T., & Fulton, J. (2009, March). *Striving to Achieve 100% Data Accuracy: The Challenge for Next Generation Asset Management.* Retrieved from http://www.ca.com/files/TechnologyBriefs/33603-itam-tech-brief-qxd-final_201662.pdf