



CMDB: The Key to Jump-Starting ITIL Success

Three experts discuss their Visible Ops methodology, which offers a simple approach to implementing ITIL in four practical steps. In this methodology, the CMDB is a key factor in driving quick wins.

**By KEVIN BEHR, GENE KIM,
AND GEORGE SPAFFORD**

Since 2000, we have met with and observed hundreds of information technology (IT) organizations to find those operating most efficiently and effectively. We identified eight high-performing IT teams with the highest service levels, best security, and greatest operational efficiencies. We documented how these highly efficient and effective organizations work so that others can jumpstart their own IT Infrastructure Library (ITIL®) implementation efforts.

What makes high-performing organizations different from average organizations, both qualitatively and quantitatively? We looked for practices that were universal to all eight high-performing organizations, and we turned to ITIL as a process framework to normalize the terms the different groups used. In the high-performing organizations, common processes occurred in change, configuration, and release management, as well as in incident and problem management.

All of the high performers had repeatable and verifiable processes to release infrastructure components with a known working configuration. They had a culture of change management as a primary way to do work, and they all used causality in their incident and problem resolution processes. Here are the characteristics we identified in each of the high-performing organizations:

- > **Server to system administrator ratios greater than 100 to 1** — In high-performing organizations, each system administrator controls more than 100 servers. In contrast, organizations not using effective processes have ratios of 15 or fewer servers to one system administrator. Figure 1 illustrates the server to system ratio benchmark.
- > **Low ratio of unplanned to planned work** — In high-performing organizations, unplanned work is only 5 percent of operational expenses. Our research showed that average organizations spend 25 to 45 percent of their total operational expenses on unplanned, unscheduled work.

- > **Higher staffing early in the IT lifecycle** — High-performing organizations deploy more resources and staff in the preproduction build phase, where the cost of defect repair is least expensive.
- > **Collaborative working relationships between IT operations and IT security** — In high-performing organizations, IT operations and IT security work together to solve common objectives. IT operations performs most of the work, and IT security acts as coach and consultant. Low performers, in contrast, have a combative relationship between teams whose objectives are not aligned.
- > **Posture of compliance** — In high-performing organizations, IT operations and auditors typically have a trusted working relationship because controls are visible, verifiable, and regularly documented. In low performers, the absence of controls causes auditors to ask questions and make suggestions, which often creates an adversarial relationship with operations.
- > **Culture of change management** — There is a ubiquitous understanding throughout high-performing organizations that changes must be effectively managed to achieve business objectives. Organizations that have large amounts of unplanned work and uncontrolled change have yet to realize the causal relationships between changes and incidents.

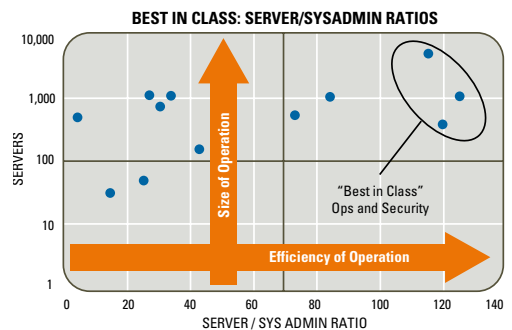


Figure 1. Server to System Administrator Ratio

- > **Culture of causality** — Through the use of controls and metrics, high-performing organizations identify and solve problems by logically studying cause and effect. In contrast, many average organizations have a culture of “let’s see if this works.”
- > **Management by fact** — High-performing organizations value controls and metrics, not only to aid effective problem-solving, but to also aid fact-driven decision-making, as opposed to “management by belief” or “management by the honor system.”

WHERE DO I START WITH ITIL?

After studying these top performers, we set out to develop a methodology to answer the urgent question that everyone was asking: “Where do I start with ITIL?” Although ITIL provides a wealth of best practices, it lacks significant prescriptive guidance about the order in which, as well as how, the processes

should be implemented. Many of the ITIL best practices are seemingly dependent on other practices already in place. Compounding the confusion, much of the third-party information that is publicly available on ITIL isn’t based on what’s been done by top performers, and is often too general and vague to effectively aid organizations.

We assume that you may be following some of these practices already, but you may not have a centralized or unified configuration management database (CMDB) in place. A CMDB is critical to creating the linkage between functional groups. This linkage enables the top-performing organizations to provide real business value.

One key to successfully implementing ITIL is effectively utilizing information about the people, processes, and technology that make up

Visible Ops Practice	CMDB Enables you to:
Stabilize the Patient — Reduce access to the production environment	<ul style="list-style-type: none"> > Identify CIs in the production environment quickly and easily > Limit access rights to production CIs > Create maintenance windows for production CIs
Electrify the Fence — Prevent unauthorized changes	<ul style="list-style-type: none"> > Identify relationships between people CIs, infrastructure CIs, and change history > Produce regular change reports about all production CIs > Compare approved change list to detected change list
Modify First Response — Link incident and problem management to recent changes	<ul style="list-style-type: none"> > Accompany incident ticket with history of recent changes to related CIs > If no changes were authorized for the CI, then expand the investigation circle to the next ring of related CIs
Create the Change Team — Manage change resources	<ul style="list-style-type: none"> > Standardize on an approval and communication process, impact assessment, and forward scheduling changes — all based on CMDB CI and relationship data
Utilize a Change Tracking System — Enforce auditable process	<ul style="list-style-type: none"> > Create and store a history of change approval and change history related to each CI

Figure 2. Capabilities Enabled by CMDB for Visible Ops Practices in Phase 1

IT. An enterprise CMDB consolidates this information and enables you to take action on it. We developed the Visible Ops methodology based on four key phases. In each phase, use of configuration item (CI) data contained in a unified CMDB enables the people and processes to effectively implement some of the key practices.

A CMDB is critical to creating the linkage between functional groups.

PHASE 1 STABILIZE THE PATIENT

The goal of the first phase is to reduce the amount of unplanned work as a percentage of total work to less than 25 percent. You curb the number of surprise system outages by freezing change outside of scheduled maintenance windows. To reduce mean time to repair (MTTR), you also ensure that incident and problem managers have all change-related information at hand to determine the cause of an outage.

CI data contained in the CMDB enables the people and processes to effectively implement some of the key practices in this phase. (See Figure 2.)

COMPLIANCE AND THE CMDB

An increasing number of regulations and privacy laws directly affect IT. As a result, key processes must be defined, and IT personnel must follow the defined processes. The processes must include activities to mitigate identified risks. And the process and results must be auditable.

The CMDB provides a central store of data that helps you meet compliance requirements in several important ways. First, the CMDB and IT service model can help you quickly identify what is in scope for various IT audits. Second, CMDB data can help you confirm that the controls function as designed. And third, the CMDB can help you verify that changes and configurations are controlled to the specification of audit requirements.

Without a CMDB, you'll need to audit various data stores and IT processes separately, driving up costs. A unified CMDB can help support common processes, which are more easily documented and controlled, across the entire IT organization. Finally, you can reduce the impact of control and audit requirements on IT functions when you have a central CMDB and data. The result is faster compliance and lower ongoing cost of implementing and maintaining effective controls in IT.

PHASE 2

CATCH AND RELEASE; FIND FRAGILE ARTIFACTS

In this phase, you inventory CIs and services to identify those with the lowest change success rates, highest MTTR, and highest business downtime costs. You identify fragile artifacts — those infrastructure items that are difficult and costly to support and maintain, are tied to an audit requirement, or are critical to the business. You treat these fragile artifacts with extra caution to avert risky changes and massive episodes of unplanned work.

If you haven't yet started a CMDB implementation, the key here is to identify the fragile artifacts and populate your CMDB with information about those CIs first. The CMDB can enable many capabilities, as shown in Figure 3.

PHASE 3

CREATE A REPEATABLE BUILD LIBRARY

In Phase 3, you establish a repeatable build process for the most critical CIs to enable a “cheaper to rebuild than repair” philosophy. By this phase, you typically can reduce unplanned work to less than 15 percent and dramatically reduce the number of different configurations in the production environment.

You'll reap the highest return on investment from implementing effective release management processes, but you'll need to complete the first two phases to successfully get through this phase. You will have to define build mechanisms, create system images, and establish documentation. The result is a repeatable process for building infrastructure from “bare metal.” CI data contained in the CMDB enables the implementation of some of the key practices in this phase. (See Figure 4.)

Visible Ops Practice	CMDB Enables you to:
Catch and Release — Audit and tag all infrastructure components	<ul style="list-style-type: none">> Serve as a central location to store information about each identified CI> Enable identification of dependency information about related CIs> Enable centralized storage of key CI attribute data
Find Fragile Artifacts — Identify and label critical or hard-to-repair components	<ul style="list-style-type: none">> Correlate and analyze change history and change success rate of each class of CI> Enable a change risk assessment of planned changes to historically fragile CIs
Prevent Configuration Mutation — Prevent out-of-process changes	<ul style="list-style-type: none">> Enforce the change process in Phase 1 as a way to enforce the update of the CI inventory in the CMDB

Figure 3. Capabilities Enabled by CMDB for Visible Ops Practices in Phase 2

Visible Ops Practice	CMDB Enables you to:
Enable a Release Team — Deploy reliable configurations into production	<ul style="list-style-type: none"> > Identify production CIs > Identify functional state of CIs > Provide clear picture of target release environment to ensure successful release
Create a Repeatable Build Process — Streamline production release	<ul style="list-style-type: none"> > Clearly identify CIs in the preproduction and production environments > Identify CIs in a build catalog > Create a documented build process CI associated with each item in the build catalog
Maintain a Definitive Software Library (DSL) — Standardize storage of approved builds	<ul style="list-style-type: none"> > Identify production CIs not in the DSL and those that don't have a related build process CI > Inventory CIs found in the DSL
Create an Acceptance Process Contract — Get approval from preproduction and post-production resources	<ul style="list-style-type: none"> > Store the contract as a revision-controlled CI
Move from Production Acceptance to Deployment for all approved builds	<ul style="list-style-type: none"> > Process request for change (RFC) for system rollout — including risk analysis, forward scheduling of changes, etc.

Figure 4. Capabilities Enabled by CMDB for Visible Ops Practices in Phase 3

PHASE 4 ENABLE CONTINUOUS IMPROVEMENT

In the previous phases, you progressively built a closed loop between the release, control, and resolution processes. This final phase implements metrics to enable the continuous improvement of all of these processes to best meet business objectives.

Philipp M. Nattermann wrote in *The McKinsey Quarterly* that if all you are doing is adopting best practices, then eventually, all you are going to get is competitive parity.¹ To really excel, you need to optimally apply all your resources to achieve the real business goals.

The CMDB is IT's knowledge repository and, as such, is a critical enabler of cross-functional performance measurement. When information is federated between a centralized CMDB and various other applications that help you manage functional processes, the consumers

and users of vast quantities of data are brought together.

We recommend a core set of performance metrics based on what we learned from studying top performers. Much of this data can be found in a CMDB or integrated from other data stores:

RELEASE MEASURES

- > **Time to provision known good builds** — How long does it take to build and provision the infrastructure from bare metal? Shorter is better, and should be shorter than any MTTR requirement.
- > **Number of turns to establish a known good build** — How many times must the build be modified before it is acceptable for deployment? Lower is better. A high number indicates the need for a more automated process.

1. Philipp M. Nattermann, "Best practice does not equal best strategy," *The McKinsey Quarterly*, 2000 Number 2. www.mckinseyquarterly.com/article_page.aspx?ar=809&L2=21&L3=35

- > **Shelf life of builds** — How long will each build be in production until it is replaced? Longer is typically better, because it enables release management teams to stay out of reactive mode.
- > **Percentage of systems that match known good builds** — According to the detective controls, how many production systems actually match their corresponding golden builds? Higher is better, because it indicates the absence of uncontrolled production configuration drift.
- > **Percentage of builds that have security signoff** — How many configurations are approved by security? Higher is better, because it indicates that security is involved in the standard “blessing” process.
- > **Number of fast-tracked builds** — How many builds are rushed into production through the emergency change process? Lower is better, because each fast-tracked build represents a deviation from the intended process.
- > **Ratio of release engineers to system administrators** — What percentage of staff is deployed on preproduction processes? Higher is typically better because the cost of defect repair is much lower in preproduction.
- > **Change success rate** — How many changes are successfully implemented, without causing an outage or episode of unplanned work? Higher is better: Best-in-class organizations achieve better than 99 percent.
- > **Number of service-affecting outages** — How many changes result in service impairment or an outage? Lower is better.
- > **Number of emergency changes** — How many changes require using the CAB emergency change process? Lower is typically better, since it indicates a higher percentage of planned work.
- > **Number of “special” changes** — How many changes, for whatever reason, are made outside of the change process? Lower is better, because these indicate that a change process is not fully functional and that management is allowing certain categories of changes to bypass change management entirely.

When information is federated between a centralized CMDB and various other applications that help you manage functional processes, the consumers and users of vast quantities of data are brought together.

CHANGE AND CONFIGURATION MEASURES

- > **Number of changes authorized per week** — How many changes, as measured by the change management process, are implemented? In general, higher is better, as long as the change success rate remains high as well.
- > **Number of actual changes made per week** — How many changes, as measured by detective controls, are implemented? In general, higher is better, but should not be higher than the changes authorized by the Change Advisory Board (CAB).
- > **Number of unauthorized changes** — How many changes circumvent the change process? This is typically measured by using the detective controls, or worse, through unplanned outages. Lower is better.
- > **Number of “business as usual” changes** — How many low-impact changes occur? This metric reflects the number of changes that have been identified as “standard” and do not require review. This metric also reflects the number of changes that can pass through without requiring change management scrutiny, but are still logged. In general, higher is better.
- > **Change management overhead** — How much effort (in hours or staffing) is the change management function consuming? In general, this number should be low. A high number may indicate a bureaucratic process, rather than one that enables productivity.
- > **Changes submitted versus changes reviewed** — What is the ratio of evaluated change requests against the total change requests turned in? A lower number is better.

INCIDENT AND PROBLEM MEASURES

- > **Mean time to repair (MTTR)** — The average time to restore service after an interruption.
- > **Mean time between failures (MTBF)** — The average time between service incidents.

Top performers have repeatable and verifiable processes for release management, a culture of change management, and incident and problem resolution processes that rely on causality.

CMDB ENABLES TOP PERFORMERS

The practices we learned from top-performing IT organizations illuminate an approach to solving IT operational process issues. Top performers have repeatable and verifiable processes for release management, a culture of change management, and incident and problem resolution processes that rely on causality. The data contained in a unified CMDB enables people and processes to effectively implement these practices, and the four phases of Visible Ops can provide a practical framework for more effectively implementing ITIL. We are confident that if you follow these steps, you will be able to replicate the transformation that other IT practitioners have achieved with their organizations. •

ABOUT THE AUTHORS



Kevin Behr is the CTO and chief operational strategist for IP Services at the IT Process Institute. He

cofounded the ITPI with Gene Kim. Kevin is an active member of the Information Systems Audit and Control Association. Kevin is a frequently invited speaker called on to address a broad range of technology and management framework topics. Kevin is co-author of The Visible Ops Handbook.



Gene Kim is cofounder of the IT Process Institute and is also the CTO and cofounder of Tripwire. Gene co-chaired the Best in Class Security and Operations Roundtable (BIC-SORT) with the Software Engineering Institute. He is co-author of The Visible Ops Handbook and is a primary researcher for the IT Controls Benchmarking Survey.



George Spafford is the managing director of Spafford Global Consulting. He is a recognized expert in IT process and audit. He is a prolific author, contributing articles to a wide range of IT publications. His Daily News e-mail list subscribers include high-level executives from Fortune 500 and international companies. George is co-author of The Visible Ops Handbook.

5 BENEFITS OF USING A CMDB TO SUPPORT ITIL PROCESSES

A CMDB consolidates and enables you to act on the information about the people, processes, and technology that make up IT.

A CMDB enables key practices that help reduce the amount of unplanned work.

A CMDB helps you to better correlate and analyze change history, assess the risk of planned changes, and enforce the change process.

A CMDB helps you identify CIs and their status so you can establish a repeatable process for building your infrastructure.

As the knowledge repository for IT, a CMDB is a critical enabler of cross-functional performance measurement.

This article is based on "The Visible Ops Handbook, Starting ITIL in 4 Practical Steps" written by the three authors and published by the Information Technology Process Institute (ITPI).