
IT Service Management Metrics that Matter

Four Key Benchmarks for Improving IT Performance

- page 2** Mean Time to Repair
- page 3** First Fix Rate
- page 4** Change Success Rate
- page 5** Server to System Administration Ratio
- page 6** Reason to Improve: Unintended Consequences of Low Performance
- page 6** About the Author

Do you think great IT performance is achieved through luck or chance? You can bet real money it's not! High performing organizations have figured out which processes and controls really help them achieve their operational effectiveness and efficiency objectives. They have integrated those processes and controls into how they manage almost every aspect of their daily work, helping them achieve their business goals and find variance before it causes a catastrophic outage, failed change, security incident or something that can impact the customer.

High performing organizations, as identified by the landmark IT Process Institute® (ITPI) *IT Controls Performance Study* published in April 2006, are able to achieve these results by implementing and enforcing two controls. These two controls, which every high performer practiced, but none of the medium and low performers did, are:

1. Actively monitor systems for unauthorized change
2. Have defined consequences for intentional, unauthorized changes

These two controls help high performers foster their desired culture of change and causality, and provide a mechanism for key metrics that determine their efficiency and effectiveness. They also make the difference between high performing organizations and those that struggle.

The ITPI study, which was conducted in cooperation with Carnegie Mellon University, Florida State University and University of Oregon, identified the key metrics, which I call "Metrics That Matter," that provide definitive guidance on where to start with IT best practices, and that give the highest rate of performance return for your organization. These metrics have significant impact on your organization's ability to control system availability, compliance, risk and operational performance. These metrics are Mean Time to Repair, First Fix Rate, Change Success Rate, and Server to System Administration Ratio.

Mean Time to Repair

High performers know that 80% of all outages are due to a change, and that 80% of mean time to repair (MTTR) is spent trying to figure out what changed. Therefore, the first question that high performers ask when a system outage occurs is "What changed?" Contrast this behavior to how low performers work. When a system goes down, the first thing they do is reboot the server in question. If that doesn't work, they'll reboot the server next to it. That didn't work? Reboot all the servers! Still not working? Reboot the firewall.

The two extremes of diagnosing and resolving outages have a dramatic impact on how quickly the problem will be found and how long the outage will last. It also serves as an incredibly accurate predictor of the processes, procedures and controls the IT organization will have in place. Analyzing the MTTR of the high, medium and low performers revealed some truly startling insights. The following figures show the MTTR of high, medium and low performers for small, medium and large outages.

For small incidents, all performers experienced similar MTTR rates. These are outages that typically require one to three people to fix, and all incidents are usually resolved in 15 minutes or less.

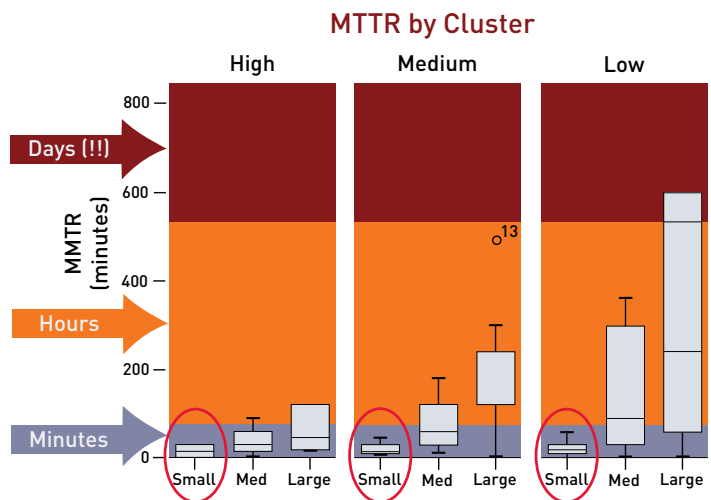


Fig. 1: Small outages and MTTR¹.

For medium severity outages (requiring up to eight people to fix), there starts to appear a growing difference in MTTR between the three groups. High performers are almost always able to resolve the issue in minutes; medium performers' resolution times begin creeping into minutes and hours.

In large outages, the differences are significant. High performers again resolve issues in minutes or hours, but medium performers resolve issues in a low number of hours. Low performers resolve issues in a much higher number of hours, sometimes taking even days to fix a problem.

Considering that large incidents are mobilizing somewhere between 25 to 50 people, low performers are sustaining the mobilization of an "all hands on deck" situation for a considerable portion of the workday. The level of disruption that this causes for an IT organization is difficult to overstate. While outages will occur, the frequency of incidents can be reduced, the problems quickly fixed, and the duration of the outage shortened, if change control is avidly practiced and enforced.

First Fix Rate

First Fix Rate measures the percentage of incidents that are successfully restored on the first fix attempt. It is a leading indicator of system availability and MTTR; that is, how well an IT organization manages First Fix Rate will also result in radically improved MTTR. First Fix Rate is often used in the connotation of the service desk, where it measures how often the incident is resolved at the first point of contact between a customer and the service provider. However, in this context, it is not the service desk that is measured, but the IT operations staff who are actually working on the incident to restore service.

This metric is one of my favorites because, almost more than any other metric, it indicates the degree to which a culture of causality exists in an IT organization. The ITPI wasn't the only organization to uncover this phenomenon. The Microsoft Operations Framework (MOF) study showed that their high-performing customers reboot servers 20 times less often than average and have five times fewer "blue screens of death."

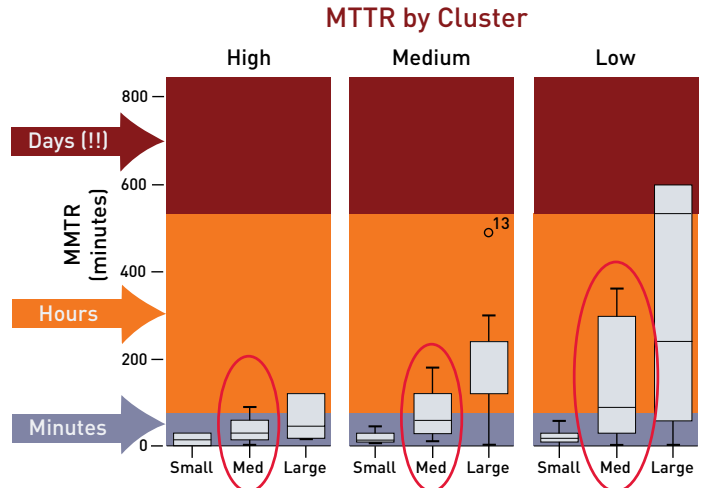


Fig. 2: Medium outages and MTTR.




Fig. 3: Large outages and MTTR.




Fig. 4: First Fix Rates for high, medium and low performers.

In the First Fix Rate chart, note how the high performers have the highest First Fix Rates at 90%, with some variance. Medium performers have lower First Fix Rates of about 80%, and low performers have the lowest First Fix Rates at 50%, both with considerably higher variance. The numbers show very dramatically how much more effectively and efficiently high performers are at figuring out the root cause of an outage (which is usually because of change). But, variance tells another story.

Think of variance in terms of target practice. High performers (those with greatest control) will shoot a rifle five times, hitting the target five times in nearly the same place. Low performers will also shoot five times, but can't hit even remotely in the same place twice. Which organization do you think is easier to improve, the consistent or the haphazard shooters?

Why is the First Fix Rate so low in the low performers? It is because outages are not worked in a systematic fashion, appropriately ruling out change early in the repair cycle. Instead, IT staff will often work by intuition. Rather than a culture of causality, they have a culture of "let's see if this works," resulting in significantly longer outages and lower First Fix Rates. These are the organizations that say, "Let's reboot this server and see if that works. It didn't work? Okay, let's try rebooting this other servers." This is exactly what the MOF study showed. And this behavior, more than any other, is responsible for pulling down First Fix Rates.

High performers, on the other hand, always rule out change first, and always use a list of authorized and detected changes to guide their efforts. Because they have most probable cause at hand, they enjoy better First Fix Rates.

Change Success Rate

Change Success Rate is determined by the number of changes successfully deployed without creating an incident, service impairment or disruption of work. As an example, suppose an organization makes 100 changes in a given week, and they find that 20 of those changes caused issues. Their Change Success Rate is 80%, which puts them right in the middle of the pack of medium and low performers. However, suppose another organization also makes 100 changes per week, but they only have five incidents. That's a Change Success Rate of 95%, which is typical for high performers.

What is interesting, though, is how high performers define a "failed change." To them, a change doesn't necessarily have to cause an incident to be defined as unsuccessful. For instance, when a change does not go accordingly to plan would also count as a process exception. This could be such as when a change exceeds the scheduled implementation time or varied from the deployment plan, requiring undocumented steps. High performers would count all these examples as failed changes. (High performers often use the term "change process exception" to define this condition in order to distinguish from a failed change.)

The Change Success Rate graphic shows high performers with 95% Change Success Rate whereas medium and low performers had 80% or less Change Success Rate. This means high performers had failed changes 25% less often than other organizations.

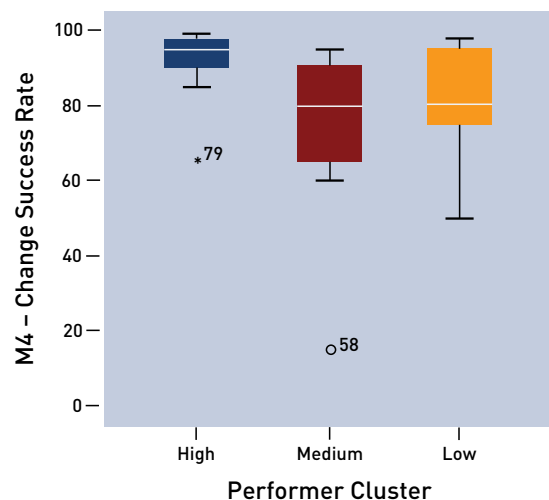


Fig. 5. Change Success Rates for high, medium and low performers.

Like First Fix Rate, the average measures of the three clusters only tell only a fraction of story. Observe how the variance in the high performers is much smaller than the measures in the medium or low performers. High performers are not only better at implementing changes, but they are more consistently hitting the target. As we look at the variance of Change Success Rate scores, we see that medium performers have a much higher variance than high performers, but the low performers have the highest variance of them all.

The ITPI also found that high performers make more than 10x more changes than medium or low performers. My interpretation of this finding is that high performers have more confidence in making changes. In other words, because they have a good track record in making successful changes that deliver value to the business without causing undesired side effects, they can sustain a higher throughput of changes. This results in the organization’s perception of nimbleness and ability to meet business needs.

Why are Change Success Rates so low in medium and low performers? This is a question that I’ve been thinking about since 1999. I believe the answer is simple: medium and low performers do not have a culture of change management, meaning that changes are not authorized, and scheduled and tested in a manner that leads to successful changes.

Server to System Administration Ratio

In 2000, when Kevin Behr and I first started studying high performing IT organizations, we noticed there might be something very special about the high performers. This “special” thing is the Server to System Administration Ratio. What we noticed back then was that high performing IT organizations were not only the most effective, but they were also the most efficient—those with the best Mean Time to Repair, First Fix Rate, and Change Success Rate also had the highest Server to System Administration Ratio. Where typical IT organizations had one system administrator managing 15-25 systems, many system administrators in the high performers were managing more than 100 systems! That’s a 4X difference in span of control! The 2006 ITPI study, besides providing new insights on the performance of high performers, also validated this six-year-old observation about Server/SysAdmin Ratio.

As you can see in the graph, the high performing cluster median Server/SysAdmin Ratio was around 125:1, while both the medium and low performing clusters was around 25:1.

In virtually every vocation, there are simple organizational indicators that are used to benchmark effectiveness and efficiency. A sales organization may want to benchmark itself against competitors by calculating its revenue per quota-bearing salesperson. If company management desires to increase the revenue-to-salesperson metric, they cannot achieve it by merely firing salespeople! Instead, they must do myriad things to systematically increase their sales productivity and sustain it over time. The outcome of pulling this off successfully is that the organization may very well end up hiring even more sales staff to continue their growth.

In the same way, the Server/SysAdmin Ratio serves as an interesting, easy-to-calculate indicator of IT organization effectiveness and efficiency. Perhaps a far more tactical metric is percentage of time spent on unplanned work, yet benchmarking suggests a linear relationship between unplanned work and the Server/SysAdmin Ratio. All the hard work and process improvement initiatives pay off not only in less time spent on unplanned work and better service levels, but better efficiencies as well.

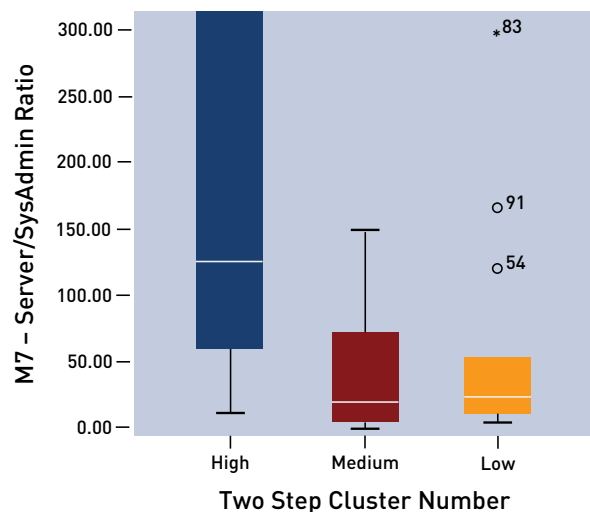


Fig. 6: Server to system administration ratio.

Reason to Improve: Unintended Consequences of Low Performance

When IT organizations perform poorly on basic metrics, one of the inevitable outcomes is a lower quality of life for the IT employees. This can have a great impact on the ability of an organization to attract and keep a trained staff. When unplanned work eats up resources and time, the organization begins a downward trajectory. Projects don't get completed, service is slowed, functional delivery is delayed, compliance problems occur, more outages happen and the availability of the network is reduced. In the view of the company, this is seen as poor organizational performance. This can quickly result in lower morale and higher turnover. If turnover rates begin climbing, management is faced with the need to constantly train new people to work in what still remains a poorly performing organization.

Conversely, when unplanned work is low, staff is given the opportunity to work on strategic business projects. This creates a virtuous cycle where the company becomes a "best places to work" organization and easily recruits and retains its staff.

Regardless of how you measure up to high performers, there are steps you can take to improve your processes. Start by controlling change to manage unplanned work—it is an amazingly simple and significant way to improve performance and processes, especially when compared to wading through a sea of best practices literature that may not be relevant to your cause. It's also helpful to know that in any improvement endeavor, the 80/20 rule applies: 20% of the set of IT controls result in 80% of the realized benefit. And let metrics that matter help guide your success.

About Tripwire

Tripwire, Inc. is the recognized leader of configuration audit and control solutions, serving over 6,000 enterprises. As the first in the industry to combine configuration assessment with configuration change auditing, Tripwire helps IT organizations automate compliance across the data center, reducing risk and increasing operational efficiency. Tripwire ensures the organization achieves continuous operational, regulatory and security compliance, helping IT achieve and maintain a known, trusted and compliant system state. Tripwire is headquartered in Portland, Oregon, with offices worldwide.

About the Author

Gene Kim

Co-Founder and CTO, Tripwire, Inc.

Co-Founder, IT Process Institute

Gene Kim is co-founder and chief technology officer of Tripwire, Inc. He is also co-founder of the IT Process Institute and co-author of *The Visible Ops Handbook: Implementing ITIL® in 4 Practical and Auditable Steps*, (IT Process Institute, 2003.) He recently worked with the Institute of Internal Auditors (IIA) on the Guidance for Auditing IT General Controls (GAIT) project, intended to help management appropriately scope the IT portions of SOX-404, published in January 2007. He was recently named by *Computerworld* magazine as one of the "40 Innovative IT People Under the

¹ All figures ©2006, 2007, The IT Process Institute, used by permission.



www.tripwire.com

US TOLL FREE: 1.800.TRIPWIRE MAIN: 503.276.7500 FAX: 503.223.0182
326 SW Broadway, 3rd Floor Portland, OR 97205 USA