ROI GUIDE

Runbook Automation for Incident Management

RUNDECK

We all instinctively know that incidents are costly. However, do you have a model for calculating the full cost of incidents? And what can you do to reduce those costs?

This guide will cover:

- Identifying the full cost of an incident
- The ROI of adding Runbook Automation to your Incident Management capabilities

Start with the Basics: Outage Time

> ***How much would you save if you could shorten Outage Time by 60%?***

Outage Time is how long a particular service is down or has degraded performance. This is the classic view of an incident that measures the value lost from missed transactions, reputation loss, or service level agreement (SLA) payouts.

Outage Time x Outage Cost Per Hour = Value Lost

Anything you can do to reduce the Outage Time is going to have a direct return on investment. Most companies have already gone through a risk management exercise to determine outage cost per hour. Sometimes the cost per hour is packed into one number. Other times it is broken into component numbers. Ask around your company and you should find either a full-baked analysis or enough ingredients to put together your own.

So how do we shorten Outage Time?

You can separate Outage Time into four distinct phases.

Detect - How quickly can we identify the outage? If the customer spots the outage first, that is the worst case. If you detect the outage before any customer impact, that is the best case.

The length of this phase largely depends on your organization's ability to leverage monitoring, observability, and alerting technologies. Like all tools, how you use them is more important than the tool you pick (and there are a lot of options from which to choose).

Mobilize - This is how long it takes to get the appropriate person(s) alerted and into a position where they can take action. Be sure not to mistake the level of commotion — or the number of people showing up —for being any closer to being able to have the right person(s) with the necessary skills in a position to work on resolving the incident.

Well-rehearsed incident management procedures supported by ITSM tools (e.g., Service Now, Remedy/Helix, Cherwell) and Incident Management tools (e.g., PagerDuty, VictorOps, xMatters, OpsGenie) can speed up an organization's ability to mobilize. However, remember that other work, headcount constraints, context switching costs, and life outside of work will always lengthen mobilization times.

Diagnose - This is how long it takes for the responder(s) to orient themself to the error, determine what the underlying cause may be, and decide on a course of action.

Note: Keep in mind that the path a particular incident takes through these phases isn't always linear. For example, during the Diagnose phase your responders might retreat to the mobilize phase if it is determined that different people are required to weigh-in on the incident.

Use Runbook Automation to Diagnose Incidents Quicker.

The more complicated the environment, the more you have to rely on a subject matter experts to figure out what has gone wrong. And, by definition, those subject matter experts are going to be limited resources.

What would happen if we could capture — as automated procedures —the steps that those subject matter experts take when diagnosing issues?

We could then put those automated procedures into the hands of anyone who responds to incidents or even have our monitoring/alerting tools kick-off the automated procedures.

Rather than waiting for the subject matter expert to be in a position to investigate, others could be running the expert diagnostics almost immediately.

For those incidents where subject matter experts still have to weigh in during the diagnostic process, Runbook Automation saves time by making available standard operating procedures (designed collaboratively with colleagues during calmer times).

Additionally, the execution history of the Runbook Automation enables the subject matter expert to get up to speed more quickly once pulled into the incident.

Rundeck Tip: What can you expect? Analysis of the Rundeck user community shows us that using Runbook Automation can shorten the time it takes to diagnose known problems by 95% and unknown/new problems by 20%.

Repair - This is how long it takes the responder(s) to undertake the necessary steps to repair or restore the service. Repair is measured from when the responders decide on a course of action to when the service is restored.

Use Runbook Automation to Repair Services Quicker.

Similar to diagnosing incidents, putting the ability to take expert action into the hands of your initial responders will shorten the length of incidents.

First, capture (as automated standard operating procedures) the actions that your experts routinely perform when responding to incidents. Then use Runbook Automation to safely distribute the ability for anyone responding to an incident to safely execute those actions. By adding Runbook Automation, repair actions happen as quickly as possible (e.g., restarts, manual failover, reconfiguration, rollback, etc.).

These automated standard operating procedures can also be helpful for the subject matter experts who still need to be called into "unknown" or significantly complicated incidents. Having these standard operating procedures allows them to work quickly and reliably. A Runbook Automation tool also allows for easier logging of operations activity and improved transparency for all involved.

Rundeck Tip: What can you expect? Analysis of the Rundeck user community shows us that using Runbook Automation can shorten the time it takes to repair known problems by 75% and unknown/new problems by 40%.

Warning

When making any calculations regarding Outage Time, beware of the limits of the usefulness of averages. This warning includes the popular Mean Time To Repair (MTTR) metric. While there is value looking at MTTR from the broadest of management perspectives, you have to remember that no two incidents are the same. You generally aren't averaging like things, even if you classify them as similar. "Known" incidents will still have some form of "unknown" in them. The more specific you can be ("this is the outage times for these specific incidents"), and the more you can analyze each incident as its own thing, the more accurate you will be.

Expand Your Focus: Escalations & Total Incident Response Hours

> ***How much would you save if you could reduce total Incident Response Hours 50%?***

Incident Response Hours are the total hours spent responding to incidents. This measure provides insight into the labor costs involved with your incidents.

Total Response Hours x Hourly Labor Cost = Value Lost

Incident Response Hours shows the impact of escalations injecting delay and disruption into your organization. Each escalation interrupts someone, disrupting either their other work or their life. That disruption and the resulting ripple effects are expensive.

The higher the Total Response Hours, the more waste and friction there is diminishing the capacity of your most valuable assets — your people.

Which hours to count?

The most popular version of this metric takes into account anyone whose primary job function isn't Level 1 incident management. Using this definition, anytime there is an escalation beyond the helpdesk or NOC (automated or human manned), you count the total time of the person(s) responding.

So How Do We Reduce Incident Response Hours?

There are three ways to reduce Incident Response Hours: 1) Avoid escalations wherever possible 2) Pinpoint escalations 3) Improve the effectiveness of those who respond to escalations.

1. Avoid Escalations Wherever possible

Failure is inevitable, but we don't have to escalate every time!

The trick is empowering those closest to the problem to respond with the same capabilities as your subject matter experts.

Using Runbook Automation to Avoid Escalations

First, use Runbook automation to automate the typical tasks and procedures performed by your subject matter experts. Include both diagnostic and repair actions.

Of course, you shouldn't expect to replicate the full expertise (that is impossible). However, there is a base set of capabilities that

should help to resolve a common set of issues (or determine who to escalate to). These capabilities would include diagnostics, performance checks, restarts, failover, rollbacks, configuration updates, and more.

Next, use the access control and "guardrails" features of your Runbook Automation to delegate those expert actions to your first responders.

Now, those closest to the problem can take swift action to diagnose and resolve issues. If they can't resolve the issue, the diagnostics they ran should be able to tell them to whom, specifically, they need to escalate.

Security and compliance often create roadblocks to allowing a broader team to respond to incidents in sensitive environments. Runbook Automation can open things up because you aren't giving out direct access to an environment. Instead, you are giving access to run a specific, pre-vetted automated procedure, and the outcome is logged.

Rundeck Tip: What can you expect? Analysis of the Rundeck user community shows us that using Runbook Automation can reduce escalations by 50%.

2. Pinpoint Escalations

Despite our best efforts to arm our first responders with the automation to execute expert diagnostic and repair actions, there

will always be a level of new or complicated incidents that require escalation.

We want to avoid the classic "all hands on deck" chain escalations. These mass mobilizations are often driven by an incident commander whose job is to bring onto a bridge call anyone who could have something to contribute.

Unfortunately, in practice, these escalations cast a wide net as those who first encountered the incident lack the expertise to diagnose the problem. As the incident continues, there is a high likelihood that even more people are paged to join the call to either rule out their area of expertise or be there "just in case."

The "blast radius" of this style of escalation drives a considerable amount of costly disruption throughout your organization. Either their other work or their life is being interrupted. And, how many of those people were necessary for the resolution of the incident?

Using Runbook Automation to Avoid Broad Escalations

By using Runbook Automation, you can empower your first responders with the ability to execute expert automated diagnostic procedures.

Based on the response, the first responder running the diagnostic automation should be able to either know what corrective action to take (e.g., run the automated restart job) or know precisely to whom to escalate.

With these quick diagnostics and pinpoint escalations, the need for the expensive big-bang, all-hands-on-deck escalations go away. And, on top of the cost savings, your resolution time goes down!

Rundeck Tip: What can you expect? Analysis of the Rundeck user community shows us that using Runbook Automation can significantly reduce the number of people who need to be paged/alerted during an incident because of the increased quality of diagnostics early in the incident.

3. Improve the Effectiveness of Those Who Respond to Escalations

When it comes time to take action, even the most seasoned subject matter experts benefit from Runbook Automation.

Using Runbook Automation to Improve the Effectiveness of Responders

Rather than relying on written documentation or tribal knowledge, Runbook Automation puts standard operating procedures at their fingertips 24/7.

Runbook Automation also improves collaboration throughout the lifecycle.

Before an incident, teams can collaborate to define their automated

diagnostic and repair procedures. This preparation gives everyone on the team the same basic capabilities as their most senior experts.

During an incident, teams have standard operating procedures at their fingertips. The centralized logging from the Runbook Automation platform improves team members' ability to follow along as actions are taken or to get caught up when coming in late to the incident.

After an incident, review, and learning is supported by the Runbook Automation platform having a clear record of what actions were taken, who took them, and what the outcome was.

Rundeck Tip: What can you expect? Analysis of the Rundeck user community shows us that using Runbook Automation can improve the overall efficiency of Ops teams by 15% and reduce the amount of self-inflicted/prolonged incidents by 60%.