DATA INSIDER

What Is Mean Time to Repair (MTTR)?

Mean time to repair is an essential failure metric that represents the average time it takes to repair and restore a component or system to functionality. As such, MTTR is a primary measurement of the maintainability of an organization's systems, equipment, applications and infrastructure, as well as its efficiency in fixing that equipment when an IT incident occurs.

MTTR begins the moment a failure is detected and encompasses diagnostic time, repair time, testing and all other activities until service is returned to end users. A low MTTR indicates that a component or service can be repaired quickly and, consequently, that any IT issues associated with it will probably have a reduced impact on the business. A high MTTR signals that a device's failure could result in a significant service interruption and thus more significantly affect the business.

According to ZK Research, **90 percent of MTTR (https://zkresearch.com/blog/2016/10/sd-wan-can-make-network-management-impactful-to-the-business/)** is spent just trying to figure out that there's actually a problem. Incorrect diagnosis or inadequate repairs can also lengthen MTTR. A high MTTR should prompt IT administrators to reevaluate their approach to troubleshooting, taking into account the entire lifecycle, from how they monitor and detect through to how they diagnose and resolve, with the goal of reducing potential downtime.

Most service-level agreements include MTTR in some manner. It's important to remember that MTTR represents a typical repair time, not a guaranteed one. A vendor claiming an MTTR of 24 hours is saying that's how long it usually takes to complete a repair, but individual incidents could take more or less time to resolve.

Depending on the context in which it's used, MTTR may also stand for mean time to recovery, mean time to resolve or mean time to resolution. In all cases, the term denotes the average time required to troubleshoot and repair an issue.

What Is Mean Time to Repair: Contents

In This Article

Additional Resources

- Data Insider home (https://www.splunk.com/en_us/data-insider.html)

MTTR Basics

What are failure metrics?

Failure metrics are performance indicators that allow organizations to track the reliability of their equipment and systems, from common desktop service requests such as basic troubleshooting of a laptop computer or connectivity problems, to server failures and other malfunctioning components that can have a significant impact. The term “failure” doesn’t only refer to non-functioning devices or systems (such as a crashed file server); it can also denote systems that are running but, due to degraded performance, have intentionally been taken offline. Any system that isn’t meeting its objectives can be declared a failure.

Common failure metrics include:

- **Mean time to repair (MTTR):** The average time to repair and restore a failed system. It’s a

measure of the maintainability of a repairable component or service. Depending on the complexity of the device and the associated issue, MTTR can be measured in minutes, hours or days. (May also stand for mean time to recovery, resolve or resolution.)

- **Mean time between failures (MTBF):** The average operational time between one device failure or system breakdown and the next. Organizations use MTBF to predict the reliability and availability of their systems and components. It can be calculated by tracking the elapsed time between system/component failures during normal operations.
- **Mean time to failure (MTTF):** The average time a device or system is expected to function before it fails. Typically, IT teams collect this data by observing system components for several days or weeks. While similar to MTBF, MTTF is normally used to describe replaceable items such as a tape drive in a backup array, whereas MTBF is used with items that can be either repaired or replaced.
- **Mean time to detect (MTTD):** The average time between the onset of a problem and when the organization detects it. MTTD denotes the span of time before IT receives a trouble ticket and when it starts the MTTR clock.
- **Mean time to investigate (MTTI):** The average time between the detection of an IT incident and when the organization begins to investigate its cause and solution. This denotes the time between MTTD and the start of MTTR.
- **Mean time to restore service (MTRS):** The

average elapsed time from the detection of an incident until the affected system or component is again available to users. MTRS differs from MTTR in that MTTR denotes how long it takes to repair an item, while MTRS refers to how long it takes to restore service after that item is repaired.

- **Mean time between system incidents (MTBSI):** The average elapsed time between the detection of two consecutive incidents. MTBSI can be calculated by adding MTBF and MTRS (MTBSI = MTBF + MTRS).
- **Failure rate:** Another reliability metric, which measures the frequency with which a component or system fails. It is expressed as a number of failures over a unit of time.

Failure metrics are essential for managing downtime and its potential to negatively affect the business. They provide IT with the quantitative and qualitative data needed to better plan for and respond to inevitable system failures.

To use failure metrics effectively, you must collect a large amount of specific, accurate data. This would be tedious and time-consuming to do manually, but modern enterprise software can easily collect the necessary data and calculate these metrics, drawing from a variety of sources with just a few clicks.

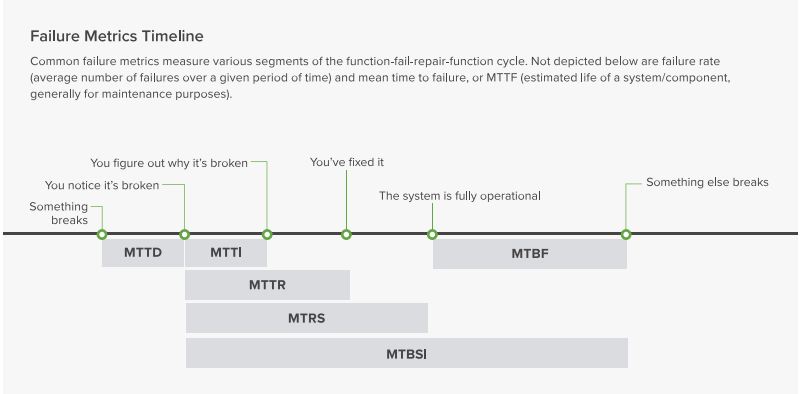

What are reliability, availability and maintainability?

Often abbreviated as RAM, reliability, availability and maintainability are system design attributes that influence the lifecycle costs of a system and its ability to meet its mission goals. As such, RAM can be a measure of an organization's confidence in its hardware, software and networks.

Each of these attributes can illuminate the strengths and weaknesses of a system and their respective impact on productivity, customer satisfaction and the organization's bottom line.

- **Reliability** refers to the probability that a system will consistently perform its intended function without failure for a given period of time. Because reliability is to some extent a function of the quality of a product, it's an inherent characteristic of a system's various components (as well as its design). However, all hardware and software is subject to failure, so failure metrics like MTBF, MTTF and failure rate are often used to measure and predict component and system reliability.
- **Availability** is the probability that a system is

operating as designed when it needs to be used. Therefore, it is a function of both reliability and maintainability, and can be calculated by dividing MTBF by the sum of MTBF and MTTR (A = MTBF / (MTBF + MTTR).)

- **Maintainability** describes the ease and speed with which a system and its component parts can be repaired or replaced, then restored to full operation after a failure. A system's maintainability depends on a host of factors, including the quality of the equipment and its installation, the skill and availability of IT personnel, and the adequacy and efficiency of maintenance and repair procedures. MTTR is one of the benchmark metrics used to determine the maintainability of a component or system, with a low MTTR figure indicating high maintainability.

Taken together, RAM can be used to determine a system's uptime (reliability) and downtime (maintainability) patterns, as well as its percentage of uptime over a particular span of time (availability).

Why is MTTR important?

Because MTTR ostensibly measures how long business-critical systems are out of service, it's a powerful predictor of the impact an IT incident will have on the organization's bottom line. The higher an IT team's MTTR, the greater the risk that the organization will experience significant downtime when IT incidents occur, potentially leading to

business disruptions, customer dissatisfaction and loss of revenue.

Technological failures are inevitable. Understanding MTTR gives organizations an idea of how quickly and efficiently they can expect to respond to these failures and return business operations to normal. On the whole, lower MTTR ratings are a sign of a healthy computing environment and a positive IT function.

What is the difference between MTTR and MTBF?

Essentially, MTBF tells an organization how often its equipment breaks down, while MTTR tells it how quickly it can get things running again. These metrics can be used together, however, to calculate a system's uptime, or availability. An organization's goal should be to both reduce MTTR and increase MTBF to minimize or avoid unplanned downtime.

How is MTTR calculated?

MTTR is calculated by dividing the total downtime caused by failures by the total number of failures. If, for example, a system fails three times in a month, and the failures resulted in a total of six hours of downtime, the MTTR would be two hours.

MTTR = 6 hours / 3 failures = 2 hours

While repairs can take minutes or days to complete, depending on the severity of the failure, MTTR of IT systems is typically measured in hours.

Applying MTTR

What is MTTR in ITIL?

MTTR is a key metric included in an IT infrastructure library (ITIL). ITIL is a series of written volumes that detail best practices for better aligning IT service management (ITSM) with business needs. It currently includes five core publications that map the ITIL service lifecycle, from "identification of customer needs and drivers of IT requirements, through to the design and implementation of the service and, finally, the monitoring and improvement phase of the service," according to **Axelos (https://www.axelos.com/best-practice-solutions/itil)**, the current owner of the library's license.

ITIL breaks down IT functions into several measurable processes, including service catalog management, service level management, risk

management, capacity management, availability management, IT service continuity management, compliance management, IT architecture management and supplier management.

MTTR is included as a part of the availability management process, whose goal is "ensuring that all IT infrastructure, processes, tools, roles, etc., are appropriate for the agreed availability targets." MTTR is noted along with MTBF, MTBSI, and MTRS, as a measurement for incident and problem management that may be included in a service level agreement (SLA).

What is MTTR in DevOps?

In **DevOps (https://www.splunk.com/en_us/data-insider/what-is-devops-and-why-is-it-important.html)** — where MTTR is normally referred to as mean time to recovery — MTTR is used to measure how long it takes for the DevOps team to recover from a production failure. Here it's typically calculated as the average production downtime over the last 10 downtime incidents.

Metrics are always essential to ensure and quantify DevOps success. Though MTTR can be skewed by the volume of new features being added to an app, code complexity and other production variables, it generally provides an accurate measure of a team's capabilities. Ideally, MTTR will shrink as an organization's DevOps implementation matures.

MTTR can also be helpful in communicating the positive business impact of DevOps to executives and other business leaders if, for example, it can

be translated into dollars saved by increasing productivity or decreasing downtime.

What is MTTR in continuous development?

MTTR is used as one measure of the stability of an organization's continuous development process.

Speed of software development and delivery is a vital driver to the success of most organizations. A robust continuous delivery process employs a "build, measure, learn" feedback loop to ensure that it's always improving and meeting business goals.

Because speed and stability are the foundation of continuous development, metrics that help evaluate and improve these issues are essential. There are no standardized metrics to monitor for continuous development, and ultimately each organization must decide which metrics are right for its goals. However, MTTR is commonly used to evaluate how quickly teams can address failures in the continuous delivery pipeline, and MTTR can serve as a guide to improving its stability.

How do you lower MTTR?

While many of the issues that contribute to a high MTTR will be unique to each organization (requiring specific evaluation of its particular IT processes and procedures), there are six general steps to lower MTTR that are likely to benefit any business.

Understand your incidents: To start reducing MTTR, you need to better understand your incidents and failures. Modern enterprise software can help you automatically unite your siloed data to produce a reliable MTTR metric and glean valuable insights about causes and contributions to this critical metric.

2 **Make sure to monitor:** Before you can fix a problem, first you need to identify it — and the sooner, the better. A good monitoring solution will provide you with a continuous stream of real-time data about your system's performance — usually in a single, easy-to-digest dashboard interface — and alert you to any issues as they develop.

3 **Have an action plan:** While ad hoc responses are often necessary for smaller, resource-strapped organizations, large enterprises should follow more rigid procedures and protocols. For many organizations, this will require a conventional ITSM approach with clearly delineated roles and responses. Companies that have successfully undergone full digital transformation may take a more flexible approach, employing cross-functional collaboration tools and constructing specific responses to each incident. Whatever your plan, make sure it clearly outlines whom to notify when an incident occurs, how to document the incident, and what steps to take as your team starts working to solve it.

4 **Automate your incident-management system:** A quick

response starts with making sure the right people receive accurate information about a problem quickly. Alerting a team member with a phone call may be fine for low-priority incidents during business hours. But what happens if a failed server takes your website offline at 8 p.m. on a Friday? An automated incident-management system can send multi-channel alerts — phone call, text message, email — to all designated responders at once, significantly saving time that would otherwise be wasted attempting to locate and manually contact each person individually.

5 **Designate response teams and roles:** Clearly defined roles and responsibilities are crucial for effectively managing incident response and lowering MTTR. Although the structure of a support organization will be shaped by your business's needs, ITIL offers a framework consisting of the following roles.

- **Incident manager:** This role drives the incident-management process, adapting and improving it as required. According to ITIL, this role is usually assigned to the service desk manager in small and midsize organizations, or it may be a separately defined role in the enterprise. The incident manager is primarily responsible for managing the incident management system and enforcing the incident management process. The IM also leads the response team, reports

key performance indicators (KPIs) back to management, and manages first- and second-line support.

- **First-line (level 1) support:** This role is the single point of contact for end users reporting service disruptions. It's responsible for classifying incidents and attempting to restore a failed service as quickly as possible. If first-line support is unable to resolve the incident, this group must route it to appropriate second-line support personnel, monitor repair activities and update users on the incident status.
- **Second-line (level 2) support:** Second-line support technicians usually have more advanced knowledge than first-line support. Therefore, they may be enlisted to address incidents that first-line support can't resolve. Second-line support responders may also be responsible for interacting with third-party software or hardware vendors to help quickly restore normal service. In very large, complex or sensitive environments, a third-line (level 3) support group with even more advanced skills and knowledge may also be required.

Your incident response team doesn't have to look exactly like this. But however you construct it, be sure it has a designated

leader who oversees incident response and ensures strong communication with stakeholders within and outside the team, and that all team members are clear on responsibilities.

6 **Cross-train team members for different roles:** Having focused knowledge specialists on your incident-response team is invaluable. However, if you rely solely on these specialists for relatively menial issues, you risk overtaxing them, which can diminish the performance of their regular responsibilities and eventually burn them out. It also handcuffs your response team if that specialist simply isn't around when an incident occurs.

You can mitigate these issues — and ultimately lower your MTTR — by making sure all team members have a deep understanding of your system and are trained across multiple functions and incident-response roles. Your team will be positioned to respond more effectively no matter who is on call when a problem emerges.

This visibility into your infrastructure can help you diagnose problems more quickly and more accurately. For example, having real-time data on the volume of a server's incoming queries and how quickly the server is responding to them will better prepare you to troubleshoot an issue when that server fails. Data also allows you to see how specific actions to repair system components are impacting system performance, so you can craft an appropriate solution more quickly.

The Bottom Line

Mean time to repair can have a profound impact on your business

With enterprise IT being pressured to increase service levels while reducing costs, incident-response times are becoming critically important. While MTTR isn't a magic number, it's a strong indicator of an organization's ability to quickly respond to and repair potentially costly problems. Given the direct impact of system downtime on productivity, profitability and customer confidence, a firm understanding of MTTR and its functions is essential for any tech-centric orrganization.

Splunk Insights for AWS Cloud Monitoring (/en_us /software /splunk-enterprise/aws-cloud-monitoring-insights.html)

Splunk App for Infrastructure (/en_us /software /splunk-enterprise /server-and-infrastructure-monitoring-and-troubleshooting.html)

VictorOps (/en_us /software /victorops.html)

Splunk Enterprise Security (/en_us /products /premium-solutions /splunk-enterprise-security.html)

Splunk Phantom (/en_us /software /splunk-security-orchestration-and-automation.html)

Splunk User Behavior Analysis

Aerospace and Defense (/en_us /solutions /industries /aerospace-defense.html)

Communications (/en_us /solutions /industries /communications.html)

Energy and Utilities (/en_us /solutions /industries /energy-and-utilities.html)

Financial Services (/en_us /solutions /industries /financial-services.html)

Healthcare (/en_us /solutions /industries /healthcare.html)

Higher Education (/en_us /solutions /industries /higher-education.html)

Manufacturing (/en_us /solutions /industries /manufacturing.html)

Nonprofits (/en_us /solutions

White Papers (https://www.splunk.com/en_us /resources.html#filter/filter3 /WhitePaper)

More... (/en_us /resources.html)

STRATEGY AND BUSINESS INSIGHTS

AI Ops (/en_us /it-operations /artificial-intelligence-aiops.html)

Machine Learning (/en_us /explore /machine-learning-artificial-intelligence.html)

Data Insider (/en_us/data-insider.html)

More... (/en_us /about-us/business-strategy-and-vision.html)

PARTNERS

Become a Partner (/en_us /partners /become-a-partner.html)

Partner Login (https://login.splunk.com

(/en_us /support-and-services /support-programs.html)

TRAINING (/EN_US /TRAINING.HTML)

ABOUT SPLUNK

Careers (/en_us /careers.html)

Events (/en_us /about-us/events.html)

Leadership Team (/en_us /about-splunk /leadership.html)

Locations (/en_us/about-splunk/contact-us.html#tabs /tab_centercontent_tabs_1)

Newsroom (/en_us /newsroom.html)

Splunk for Good (/en_us /about-us/splunk4good.html)

Splunk Protects (/en_us/legal /splunk-data-security-and-privacy.html)

More... (/en_us /about-

Answers (https://answers

Blogs (https://www.sp /blog)

Community (/en_us /community.htm

.conf (https://conf.spl

Developers (http://dev.splur

Documentation (http://docs.splu /Documentatior

Splunkbase (https://splunkb

SplunkLive! (https://splunkli

T-shirt Store (https://www.my /splunk)

User Groups (https://usergro

(/en_us /software/user-behavior-analytics.html)

Splunk Insights for Ransomware (/en_us /products /splunk-insights /splunk-insights-for-ransomware.html)

Splunk for Industrial IoT (/en_us /software /splunk-industrial-asset-management-and-intelligence.html)

Splunk Business Flow (/en_us /software /business-analytics-and-process-mining.html)

/industries /nonprofit.html)

Online Services (/en_us /solutions /industries /online-services.html)

Public Sector (/en_us /solutions /industries /public-sector.html)

Retail (/en_us /solutions /industries /retail.html)

/en_us splunk.html) /index.php?module=roles& func=showloginform& redirecturl=https %3A%2F %2Flogin.splunk.com%2Fpage%2Fsso_redirect%3Ftype artid=partner)

More... (/en_us /partners.html)

FREE TRIALS AND DOWNLOADS (/EN_US /DOWNLOAD.HTML)

PRICING (/EN_US /SOFTWARE /PRICING.HTML)

CALCULATORS

Splunk Value Calculator (/en_us/value-calculator.html)

Critical IT Incident Calculator (/en_us/form/critical-it-incident-calculator.html)

> (https://www.splunk.com)

(https://twitter.com/splunk) (https://www.facebook.com/splunk) (https://www.linkedin.com/company/splunk) (https://www.youtube.com/user/splunkvideos) (https://www.instagram.com/splunk/)
