ActZero

TechTarget | Custom Media
White Paper

Contextualizing Mean Time Metrics to Improve Evaluation of Cybersecurity Vendors

MTTX speed measures don't speak to the quality of detection from MDR vendors. Buyers need to consider metrics that gauge detection probability and coverage.

Security service providers often tout impressive sub-second performance speeds to prove their efficacy in cybersecurity detection and response. But Mean Time to 'X' (MTTX) metrics—measurements like mean time to detect, alert, respond and so on—lack the operational context necessary to appropriately judge effectiveness of vendors, especially in the era of AI-backed tooling.

Organizations need a better blend of contextual measurements to get a proper sense of how well their security vendors, or even their internal SOCs, are responding to threats. While speed is important, so too is the quality and coverage of detection, as well as the thoroughness of response, and the efficiency of the teams and tools in avoiding time wasted on false positives.

As detection is increasingly being made at machine speed, mean times become a less and less effective means of evaluating providers. When everyone's detection speeds are so fast, it no longer becomes a differentiator. Instead, potential customers should think about metrics that can help them judge detection quality. Due diligence should now consider the principles of signal-to-noise by seeking metrics that offer insight into the thoroughness, consistency, and accuracy of detections over time.

Improving the Signal-to-Noise Ratio in MDR

One of the most fundamental issues in cybersecurity today is that most legacy technology results in so many false positives and uncontextualized alerts that analysts struggle to distinguish the detection signals from noise. When automated alerting is added into the mix, the volume of noise grows, creating an additional burden on IT and security professionals commonly known as "alert fatigue."

Alert fatigue creates an environment where it's all too easy to miss attacks even when systems have accurately alerted on signals of their presence.

The promise of managed detection and response (MDR) services comes from the provider's capability in helping customers boost the fidelity and strength of those legitimate signals compared to the noise being fed to incident responders. The best providers use robust data science, mature threat hunting practices and advanced Artificial Intelligence (AI)/Machine Learning (ML) models to create quality detections that surface context between various indicators of attack and network behavior that would be impossible for a human to do.

Leveraging AI and ML to hyper-automate detection isn't easy. It requires building high-fidelity models that can detect anomalies with amazing accuracy, and automatically alert and mitigate upon detection.

Leveraging AI and ML to hyper-automate detection isn't easy. It requires building high-fidelity models that can detect anomalies with amazing accuracy, and automatically alert and mitigate upon detection. The models need to be constantly tuned based on new learnings. It's not a "set it and forget it" approach.

And this is the obstacle for customers. A provider's ability to sift signal from noise is difficult to measure, and MTTX metrics do not capture it.

When you evaluate speed, you only evaluate to the right of "boom," namely everything after the attack has begun. In an age of hyper-automation, it doesn't make sense to ask how quickly the machine can find an attack. Of course it can run fast, it's a machine. Instead, customers need to to evaluate how well that service is making the detections in the first place. To emphasize this point, let's look closely at MTTX metrics and why they are becoming less relevant today.

The Problem with Mean Time to X (MTTX) Metrics

Security metrics have always been difficult to develop. Mean time statistics are a legacy of the telecommunications and trouble-ticketing worlds. With a long history of success in those arenas, security experts latched onto these many years ago as an easily understandable and measurable way to explain the efficacy of security incident handling. Non-security executives understood how they worked, and they provided a "good-enough" benchmark in the absence of any other quantifiable security metric.

We would argue that the lens that MTTX statistics bring to cybersecurity are not in line with modern considerations of cybersecurity detection and response.

Mean Time to What?

One of the major issues with MTTX is a lack of industry standardization of various mean time acronyms.

For example, take the most common acronym, MTTR. What does it represent to your organization? Mean time to respond? Mean time to remediate? Mean time to recover? Mean time to resolution? It could be any of them. Each has a different meaning with regard to your system and may not even mean the same thing across vendors.

So, to understand the differences, let's define a few of the most commonly used security MTTX terms and offer the challenges posed by each:

Mean Time to Detect: Often referred to as breakout time, mean time to detect is typically the sum of all the detection times for an analyst, a group, or a time period, divided by the number of incidents.

Σ of incident detection times / Number of incidents

Mean Time to Alert: This is the mean amount of time from the detection to the notification, or alert, that goes to the analyst. Some providers also use this to describe the time from detection to informing a customer of the alert, also referred to as Mean time to Inform. This ambiguity can impact your ability to assess providers. Are they reporting the time between the detection and alerting their analyst? Or the time between the detection and the customer getting notified of it? Neither measure really stands to improve much anymore. After all, how much faster can the email go out upon detection of a potential threat?

Mean Time to Respond: Mean time (typically in minutes) for an email or phone call to reply and react to common alerts from staff. Typically, this indicator refers to incidents, not alerts. However, by tracking time to react to an alert, you foster the review/investigative activities that yield the discovery of incidents.

Mean Time to Recover: This is the average time it takes to recover from an incident. Typically calculated from the total elapsed time between the detection of an incident, and the return to "normal" operations, as indicated by removal of all malware on afflicted systems. Note, there is another measure, detailed below, that separates the recovery/resolution from restorative activities that may be required as a result of the incident.

Mean Time to Restore: Mean time (typically in minutes) required to restore active systems so they're ready for use. Understanding how long it takes for restoration can help customers ascertain whether a provider can keep downtime low enough to avoid business impact. Learning from previous restorative efforts and planning steps to achieve them more quickly in the future can help feed your incident response plan.

As AI-based cybersecurity is speeding up the detection of the "typical" or run-of-the-mill attack, most of these mean time metrics offer a diminishing margin of return. If 99% of attacks are automatically detected and blocked in seconds, for example, then that MTTR is going to reflect lightning-fast speed, and those speeds are going to be remarkably similar across vendors.

Often, one of the factors security teams care about is how well and how quickly service providers are able to detect outlier attacks—the 1% that aren't quickly or automatically blocked by normal measures. These are the kinds of targeted and hidden attacks most likely to cause the highest damage and the worst breaches. In other words, the ones that keep executives up at night. It can take days, weeks, or months to detect a handful of the most dangerous attacks, and sometimes the vendor doesn't even detect or respond to these attacks. But that's not reflected in the mean time metrics, because it is all too easy for a provider to inflate a low mean time by only measuring response to signature detections of known threats.

Given that, organizations need to provide some qualifiers for their MTTX metrics to offer greater context around them. Hearing "99.9% of all incidents are responded to within 4 minutes" offers a greater degree of confidence than 'our MTTR is 4 minutes.' In order to get there, consider a comparison of MTTX filtered by or compared with the total percentage of the times that the average falls within that measure (total percentage n) of true detection to get a little closer to the truth.

Choosing the Right Detection Quality and Coverage Metrics

Uncovering great benchmarks for MDR services will take more than simply looking at the total percentage of true detections in a random time period. The correct yardsticks should also measure how well a provider is amplifying signal to noise. Unfortunately, there is no singular signal to noise metric. If there was, everyone would already be using it.

Instead, organizations need to cobble together a few new ways of querying available data to create a body of metrics that can at least hint at vendor capabilities. The question to ask is, if speed isn't the important thing, then what is?

It's more about consumption of security information. Look for metrics that show how much data the provider can ingest and analyze to find an attacker. When you think about it, that's what customers really want from a managed detection response provider. They're paying people who are managing AI to create detections that otherwise wouldn't be available via cookie cutter security controls or services. Otherwise, they'd simply engage a traditional

managed security service provider (MSSP). But we all know that if there is no noise in the woods, then no one is going to react to it. That's why MDR exists in the first place.

Some important types of metrics to consider that would offer greater insight into an MDR provider's ability to detect signal from the noise include:

Ingestion Metrics: Measurables like number of data sources and volume of logs analyzed

Coverage Metrics: Number of data sources per data type analyzed, e.g., analyzing three different Cloud SaaS data sources—Office 365, G Suite and Salesforce

False Positives: Measuring false positives generated by detections over a specific time period or by number of "things"—be it endpoints, accounts and so on—can offer a view into noise levels

Recall: Number of known attacks detected by the service

By themselves, mean times don't necessarily connote security effectiveness, or even offer a good yardstick for differentiating vendors.

The last metric—recall—can be particularly useful. But it is also problematic as it can be easily cherry-picked by adjusting the source of truth from which the baseline number of "known" attacks is pulled from. So, we caution customers to consider it but also take care in how it is analyzed and used.

Mean Times Still Serve a Purpose

Speed matters when it comes to judging MDR effectiveness, but as a secondary means of validating a provider's capabilities. By themselves, mean times don't necessarily connote security effectiveness, or even offer a good yardstick for differentiating vendors. However, in context with or filtered by metrics that measure probability of detection, they can offer value —especially if an organization maximizes which mean time metrics they use, and clearly defines them, relative to the relevant segment of the attack lifecycle for which they'd like to improve response.

The process of identifying which specific MTTX measurables you need is all about thinking about use cases and the variables you require.

Businesses are likely to start rejecting the idea that MTTD is the most important metric, because many vendors don't detect the worst problems in the first place. The smarter organizations, on the other hand, might find that mean time to recover—the best MTTR—delivers more value. In that case, it does a good job of assessing the mitigation capabilities of a provider—how well they can drive down the impact of actual incidents over time. And, to throw yet another competitor into the mix, mean time to remediate has important implications for the proactive hardening efforts/recommendations that providers advocate on behalf of IT teams to ensure there are fewer opportunities for hackers to take advantage of in the first place, as quickly as possible.

Final recommendations

Ultimately, a prospective MDR customer must remember that there's no magic formula for picking the right provider, especially not one involving the mean time of anything. As you chart your path, here are some recommendations:

- Consider alternative metrics, both for acquiring new services and setting up service-level agreements (SLAs)
- If you must use a mean time, make it mean time to recover
- No metrics are perfect, so don't forget the power of due diligence

- Interview technical team and ask them piercing "left of boom" questions, whether they have the metrics or not to back up the answers
  - How many kinds of log sources can your ML operate on?
  - What proportion of your detections leverage machine learning?
  - How would you describe your data science team? Does it include those with expertise in security engineering?
- Ask for references and really talk to these customers about their pain points and why they like the vendor

ActZero challenges cybersecurity coverage for SMB and mid-market companies.

Whether shoring up an existing security strategy or serving as the primary line of defense, ActZero enables business growth by empowering customers to cover more ground.

Learn More

To learn more about how ActZero uses machine learning to ensure higher quality and breadth of detections, with so few false positives that responses can be automated to respond at machine speed, check out our other white paper, "The Hyperscale SOC and the Minds Behind It: A Machine-learning Foundation for Effective Cybersecurity." Or, to see the results in action, request a demo of our intelligent managed detection and response (MDR) service.

ActZero

About ActZero

ActZero challenges cybersecurity coverage for SMB and mid-market companies. Intelligent MDR provides 24/7 monitoring, protection and response support that goes well beyond other third-party software solutions. Our teams of data scientists leverage cutting-edge technologies like AI and ML to scale resources, identify vulnerabilities and eliminate more threats in less time. We actively partner with customers to drive security engineering, increase internal efficiencies and effectiveness and, ultimately, build a mature cybersecurity posture.

Whether shoring up an existing security strategy or serving as the primary line of defense, ActZero enables business growth by empowering customers to cover more ground.

ActZero | ActZeroAI | ActZero.ai

---

www.ActZero.ai/contact

TORONTO
207 Queens Quay, Suite 820
Toronto, Ontario M5J 1A7

MENLO PARK
2882 Sand Hill Road, Suite 115
Menlo Park, California 94025

SEATTLE
925 4th Ave., 20th Floor
Seattle, Washington 98104

This content was commissioned by ActZero and produced by TechTarget Inc.