zenoss
Own IT.

FORRESTER®

Powering Digital Transformation with

Intelligent Monitoring & Analytics

How Using the Right Data in Context Improves Speed, Stability, Incident Command and MTTR

During our recent webinar, Using the Right Data In Context to Prevent IT Outages, guest speaker, Forrester Principal Analyst Charles Betz shared his insights on the importance of leveraging the different types of machine data, collected through intelligent monitoring and analytics solutions, to enable observability and act on critical insights that allow you to speed IT and accelerate digital transformation. Here are some of our takeaways from this webinar.

1. Digital Transformation Means Need for Speed

The overriding theme for much of digital transformation is this need for speed. As we can see, it's no longer about having large, well-established firms that operate at their own pace and at their own leisure. This is an age where speed matters. Speed is essential — speed to market, speed to experiment. And this is driving so much of what we are seeing with IT and digital operations today. We know that the customer is king. Forrester calls this the age of the customer. Customers are driving vendors and suppliers with their insistence on fast, delightful experiences. All of this comes together into this overall drive for speed and how the industry has been responding to that. What we see is that first, overall, the customers and the business technology professionals are agreeing that digital technology enhancements are driving the business. And when we ask these same people, "What do you mean by that? How do you want to improve the customer experience?" Many answer that improving online customer experience is essential.

2. Experimentation and the Speed-Quality Paradox Cannot Be Ignored

Another way that speed plays out is in the ability to experiment more quickly — and this has been one of the great paradoxes. **As we move into these more agile and DevOps-influenced forms of development and deployment and delivery, how do we maintain quality?** So, you have an idea, and you need to test that idea. You need to bring it in contact with the market. You need to build, and the term for this is to create a minimum viable product, or MVP. You then need to test this in various ways against reality — against customer demand, for example — and measure the response. That is what gives you the learning that you need that then informs your decisions, e.g., building along the same lines with some further improvements, or going back to the drawing board and generating some new ideas, or whether the whole endeavor is something that should be abandoned. And this is how we ultimately understand our business outcomes.

66% of professionals agree that digital technology enhancements are the driving force behind their business strategy

3. For DevOps, Change Can Actually Improve Stability

DevOps is expanding across the industry — it has passed the experimenters, the innovators. It is now firmly in the early majority. And, in fact, I think we're going to be seeing late majority in the next 18 months. We already have nearly half of the industry already aware of it, experimenting with it. The success stories are just too compelling. But DevOps is still overcoming the big debate

of stability versus change. But one of the fundamental realizations that the DevOps pioneers started to work with is - that **making smaller, more frequent changes actually might lead to better system stability. Imagine that.**

More frequent changes in and of themselves do not equal better results and stability. However, when you make changes more frequently, by nature, you will also be making smaller changes. And this has been shown in company after company to improve your time to market and retain stability.

74% of developers say their organizations are using DevOps to some degree, with 52% increasing

4. Traditional Incident Management is Becoming Incident Command

We are so dependent on digital systems as a society, and this is resulting in some interesting responses. **Amazon realized some years back that traditional IT incident management was no longer sufficient** for the level of dependency other companies were beginning to have on the Amazon cloud. And so, they started looking more broadly and found some very interesting material coming out of the U.S. emergency services, including a protocol and a method known as the National Incident Management System, which is currently maintained by the U.S. Department of Homeland Security. There are certain sets of defined protocols and expectations that people must follow when in these very critical situations. And increasingly, we are seeing more and more interest in this framing, in these perspectives as incidents in digital systems become more socially impacting. In fact, this was a major theme at the DevOps Enterprise Summit in San Francisco, where a number of safety professionals, non-IT, were gathered on stage and challenged the DevOps Enterprise Summit attendees to increase the level of professionalism because digital systems were becoming so socially and economically critical.

5. Automation Drives Need for Monitoring and Observability

One of the themes that has presented itself recently is this problem of what happens when we automate all of the easy things. We have amazing resiliency in our digital systems. We have the ability to automatically heal containers. But the common sense concern is this: If I've automated all the easy stuff, then what is left? All the hard stuff.

There is some evidence about mean time to resolution (MTTR) going up in organizations that, on the surface, appear pretty mature. But it seems to make sense to me that as we, with DevOps, have faster release cycles, it means we are solving problems more quickly. If we solve problems more quickly, we lose a longstanding concept in IT service management called the "known error." In the old days when someone found an issue, you created a known error record so when somebody ran into the issue again, the help desk could do x, y and z. Now we are in a day and age where things like that get fixed right away. They get fixed, they get patched, you roll the patch out, you have increasing release frequency.

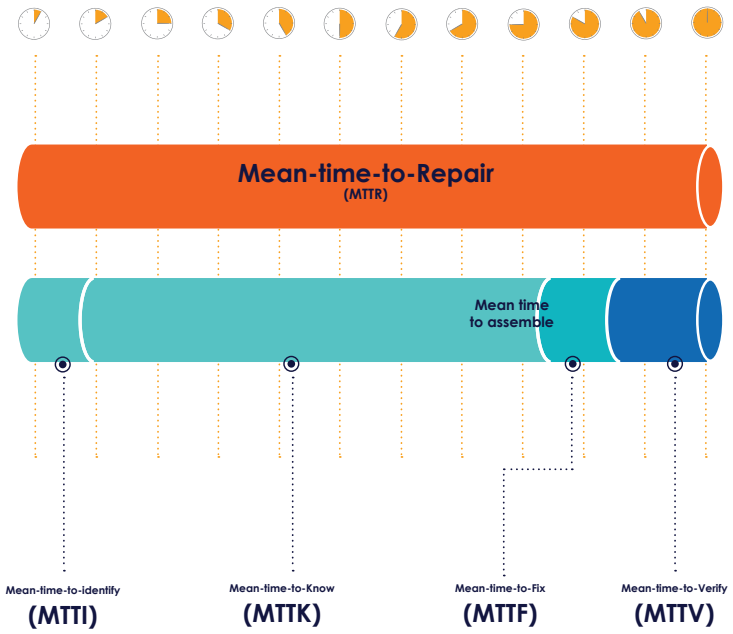

And so, your known error life cycles start to decrease — which means your understanding of the likely operational outage scenarios, issue scenarios — that stock of knowledge — is going down as you increase your release cycles. That has been essential for the resolution of IT and digital incidents and issues. Broadly speaking this is all a form of knowledge work, and we know that knowledge workers spend up to 20 percent of their time looking for this kind of information. **And all of this, in the context of digital systems, leads to a critical need for monitoring and observability.**

6. Using Intelligent Analytics to Understand Critical Data Types Enables Fast IT

Monitoring and observability are key to the success of digital transformation, experimentation, speed and quality, stability in DevOps, incident command, and improving MTTR — key to the success of "fast IT." Monitoring and observability, in turn, are tremendous producers and consumers of various kinds of information that enable efforts to modernize IT. It's important to note that there are two sets of equally important capabilities: collecting this data and analyzing this data. **Companies should seek out the platforms that have the most robust analytics capabilities for the broadest sets of data.** And we assert that there are five critical data types for this endeavor.

Log Data

Computers and the software running on them all produce logs. They show, in great detail, the state changes, the events, the various functions, users logging in, attempting to log in, hackers trying to log in. We also then use various techniques to analyze them. And it's, of course, the analytics that provide the greatest value when you are looking at log data. At the lowest level, you can see what is being logged as an event, and then as you analyze the logs to determine what happened. When we see these 3 or 5 or 10 things in a log, we know that this may mean there is a situation, so we generate a higher-level notification, i.e., an event.

Metrics

Metrics can tell us about operational concerns in the infrastructure like performance degradation or capacity utilization. There's a wide variety of infrastructure metrics that give us understanding into operational concerns that might even require the dispatch of a crisis team or, at the very least, require proactive remediation. There are also longer-cycle business-as-usual concerns, like, maybe we need to invest in some more capacity. We can collect these on interval. With proper analytics, metrics provide rich data for diagnosing issues that have occurred and for preventing issues that may occur. These analytics can produce higher-level alerts in the form of events.

Application Data

Application data is a form of specialized metrics. But your application metrics are very different, and that's where you start seeing business metrics, like, for example, shopping cart abandonment. Well, perhaps the shopping cart abandonment on your e-commerce site is due to the fact that your performance is degraded to the point where people are no longer going to do business with you that day. Like infrastructure metrics, proper analytics application metrics can provide higher-level alerts about application health and performance in the form of events.

Event Data

An event is a significant marker that helps us understand and contextualize and turn the logs and metrics from mere data into actual information that is then actionable. For example, we may see an attempt to log in a dozen times on a server. That's just 12 data points, but then we are also reasonably confident that this probably means that there is a brute force attack or other potential security exploit in progress. And that can then be understood to be an event, and that is actually much more useful information. We've taken a set of information — from the logs, from the infrastructure, from the applications — created higher-order events, then correlated that information to parse out actionable events.

Model Data

While events can be correlated to find patterns, event engines have no knowledge of systems, their interconnectedness, nor their dependencies on each other. This is where model data comes into play. The overall model data may be the hardest form of data to collect and understand. It is data, the metadata, that actually gives you the overall context for how to understand all of the previous data, the log data, the event data, the infrastructure and application metrics. There is an immense level of complexity as we start to model these systems and we look at the dependencies between datastores and midtier processing and edge processing. We need to understand this model once it is actually operationalized and we start to understand and look at the dependencies and the telemetry we are getting off of the runtime operational systems. But actually, compiling and maintaining this dependency data is still a nontrivial problem.

SUMMARY

As we venture into the brave new world that awaits us over the next couple of years, we will bring all the data together and start using data warehousing techniques. Because understanding all of these data types, in context, is critical, there will be significant focus on platforms that bridge data silos. These platforms will be key to enabling "fast IT"and will be a key ingredient in businesses that survive digital transformation and remain relevant for years to come.

ZENOSS CLOUD

Zenoss works with the world's largest organizations to ensure their IT services and applications are always on. As the leader in software-defined IT operations, Zenoss uniquely collects all types of machine data to build real-time IT service models that train machine learning algorithms to predict and eliminate outages in hybrid IT environments, dramatically reducing downtime and IT spend.

Zenoss Cloud is the first SaaS-based intelligent IT operations management platform that streams and normalizes all machine data, uniquely enabling the emergence of context for preventing service disruptions in complex, modern IT environments. Zenoss Cloud builds the most granular and intelligent infrastructure relationship models possible at any scale and proactively provides unparalleled holistic health and deep performance insights to optimize any IT environment.

Technology vendors have taken many different approaches over the years to help prevent IT service outages and improve overall IT performance. These approaches include infrastructure monitoring, artificial intelligence operations, log analytics and more. Some approaches collect performance data from systems directly, some rely on logs, some rely on events, while others rely on data sent from agents. Zenoss Cloud is the unique platform that combines all of these approaches.

ZENOSS CLOUD HELPS CUSTOMERS:

Increase Operational Agility

- Automate processes and streamline collaboration to enable faster service delivery
- Support new business models at the speed of demand
- Deliver management as a service for DevOps teams

Accelerate Technology Adoption

- Simplify cloud migrations and adoption of software-defined and converged technologies
- Eliminate risk associated with digital transformation
- Apply consistent monitoring policies across all cloud and on-premises systems

Ensure Service Reliability

- Identify issues, isolate root cause and accelerate resolution before disruptions impact users or business
- Evolve from availability and performance to capacity and optimization
- Transition IT to event-driven outcomes

Consolidate Monitoring Tools

- Increase IT visibility and eliminate silos while reducing overhead and spend
- Streamline across teams with collaboration workflows (ChatOps)
- Drive new efficiencies with Smart View, the machine learning–powered dynamic user interface

For more information, or to request a Zenoss Cloud trial, please visit **https://www.zenoss.com.**

www.zenoss.com

1-512-687-6854 (direct)
1-888-936-6770 (toll free)

twitter.com/zenoss

www.linkedin.com/
company/zenoss-inc-

zenoss
Own IT.