dynatrace

Tipping point:

Why IT Operations must change their approach to monitoring

WHITE PAPER

IT Operations has reached a tipping point. In the quest to watch over the application landscape, to ensure operational integrity and intervene as early as possible when problems occur, the complexity of the environment has surpassed the ability of humans to monitor it. In particular, with the introduction of cloud, microservices, and agile development, there are too many "things" and they change too often. The status quo, in terms of techniques and technology, is simply not adequate to enable IT to manage these hyperscale environments. There is too much burden on the human operator to parse too much information to reach a conclusion and remedy. The situation represents a real obstacle to the successful deployment of strategic new digital business initiatives.

It presents an even bigger obstacle to addressing a more fundamental and insidious issue, that is that IT organizations are simply not managing the overwhelming majority of their applications. Analysts estimate that the typical organization monitors no more than 25% of their apps for performance and availability. If only the next 5% represent significant business value, which is ridiculously conservative, a lot of business is at risk. Organizations are not managing the vast majority of their applications because it's perceived to be infeasible — with the current state of APM technology, too many operators would be required.

Recognizing this, Dynatrace has completely redefined monitoring with a fundamentally new approach that based on an Artificial Intelligence (A.I.) engine and automation to eliminate the intractable burden otherwise placed on human operators. It is made possible because Dynatrace efficiently captures the most complete operational data in the industry — for every transaction and every user.

Other vendors claim to embed advanced IT Operations Analytics (ITOA) and even A.I. in their monitoring tools. This paper describes how Dynatrace's approach is different, and significantly advanced even compared with other leading vendors. It is in production today in organizations all over the world. It's not a vision or a planned enhancement — it's live. And it is changing the way organizations monitor their business-critical applications.

Providing answers — not just information

For the purpose of discussion, let's consider IT Operations monitoring in terms of a very simple model encompassing three basic processes.

- **Observe** the operation of the application ecosystem by monitoring key metrics.
- **Analyze** the data to identify patterns, detect problems, and assess the scope of problem impact.
- **Respond** to conclusions derived from the analyses.

The support for these processes provided by most APM solutions today is inadequate.

- They are unable to scale to support the sheer volume of data generated by all the elements in a complex application ecosystem. Some have a narrow scope of domain coverage; most use sampling or capture snaphots in response to some event instead of capturing comprehensive data. Some tools aggregate log data and combine it with simple performance measurements. These techniques produce low value data that inherently produce an incomplete and inaccurate picture of activity. This is particularly damaging to problem identification that relies on anomaly detection.

- They lack the intelligence to automate analysis to a level that sufficiently offloads the burden on human operators to draw salient conclusions. This represents another kind of scalability issue. It is evident in two areas: problem identification, and root cause analysis.

There is clear evidence supporting these assertions:

- An obvious demonstration of poor problem identification is the prevalence of alert storms. Anomalies are not uncommon, and most are not indicative of an actual problem. APM tools with poor problem identification analytics produce a large volume of spurious alerts, relying on the operator to determine which are of genuine concern. These can even be generated by normal events, like maintenance to a system component, or running a back-up.

  Operators often report that the volume of alerts is so high, they simply cannot parse them. Under these all too common circumstances, high impact problems are more likely to be identified through an increase in calls to the service desk than through proactive alerting. The solution with most APM tools is manual tuning of parameters such as anomaly detection sensitivity through trial and error.[1]

- The most common approach to problem root cause analysis is correlation, i.e., finding coincident events that share some common attribute and that are consequently probably related to each other and/or the problem. However, correlation does not indicate the nature of the relationship between events; they could be merely coincidence, or be reacting to a separate unobserved event and have no direct relationship to each other at all. Correlation certainly does not point to cause-and-effect. It merely narrows the scope where an operator would look further for root cause. In effect, "it's probably in here somewhere." It's up to the operator to choose the right metrics and pick out the relevant data in order to determine what has occurred.

Dynatrace automates processes that otherwise must be handled by operators. It's A.I. engine is able to develop conclusions automatically that, with other tools, must be figured out manually. It's exponentially faster than humans, is able to factor in exponentially more data, and unlike human analysis, is not subject to biases or assumptions not based on fact. In short:

- It is more effective in identifying a problem, with self-learning capabilities that automatically discern performance anomalies, leveraging an auto-generated baseline that incorporates a deep understanding of the organization's app and infrastructure behavior and business cycles.

- It is better at determining the specific root cause of a problem, automating much more of the process than solutions that stop with event correlation. An interactive infographic tells you where a problem is and what you can do about it. You can even leverage an instant replay of application problems, and visualize how the various components of your environment were affected over time.

- Problems are automatically prioritized based on their impact on your users.

APM deployment at scale

Just as applications services are frequently deployed and updated in modern IT environments, the associated technology for application monitoring must be able to be as well. This is another area where most APM tools are lacking, requiring manual deployment — and updates — for agents, and different agents for different ecosystem components. Dynatrace instruments your apps from a single installer, and agents are updated automatically. No manual configuration is necessary, and overhead is extremely light. Dynatrace One Agent sees the entire IT stack, and can auto-detect and monitor the broadest set of technologies in the industry — at scale.

1. Dynatrace also allows manual tuning of alert sensitivity. However, rather than doing so to simply reduce alert volume, alerting would be adjusted for unique cases such as setting higher thresholds for applications and services in development or test, or lowering thresholds for critical infrastructure services where default thresholds are considered too tolerant. Auto-baselining is unaffected by such changes.

A PROBLEM SUMMARY

The problem summary description includes its duration and affected resources. Here, an increase in JavaScript error rate is observed.

B IMPACT SUMMARY

The impact is summarized, including affected users and transactions

C ROOT CAUSE

Root cause is determined automatically, and presented succinctly. The steps in the analysis can be explored in detail.

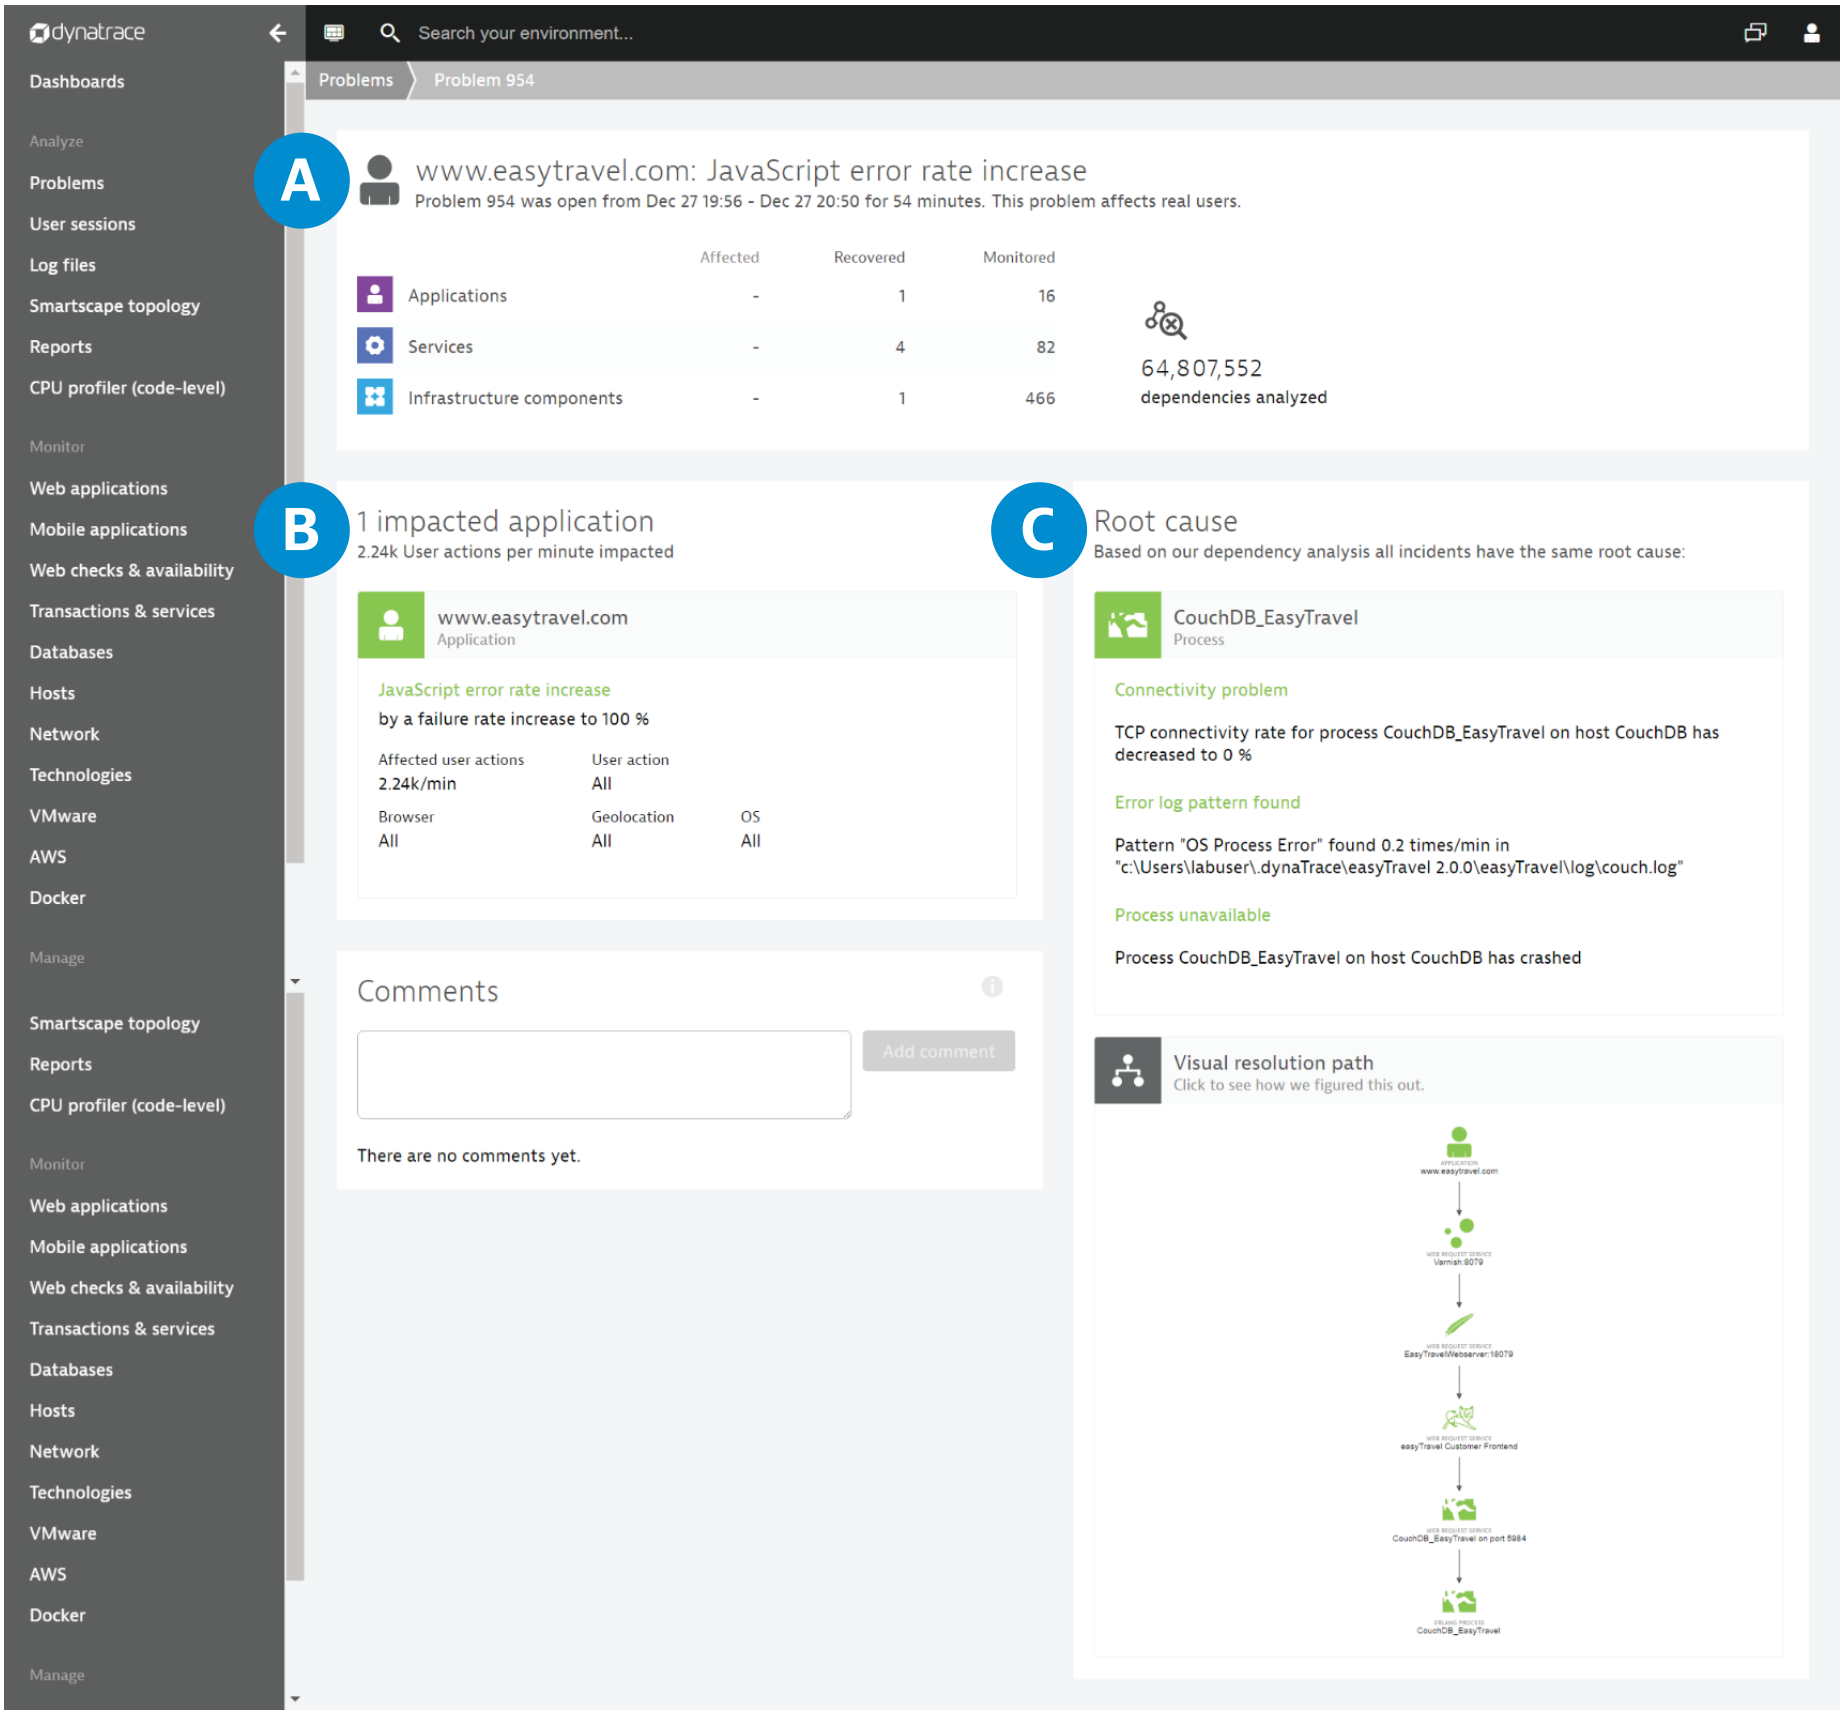

Dynatrace AI automatically determined that a spike in the rate of JavaScript errors is caused by the failure of a backend "CouchDB" process, affecting all the users of the "Easy Travel" web application.

AI-powered automation

Dynatrace is able to do all this because the A.I. engine in Dynatrace understands causation. Causation is the capacity of one variable to influence another. This is often confused with correlation. However, unlike correlation, causation conveys an understanding of the specific relationship between variables, not just that there probably is one.

Dynatrace's A.I. engine works because of two key ingredients: full-stack connected data, and smart algorithms that tune themselves over time through learning. The next sections explore each of these in more detail.

Connected data

Greater intelligence relies on high value data — more complete and connected. Only Dynatrace has the ability to capture every swipe, every click, for every transaction, every user, 24x7. Dynatrace tracks applications and transactions from end-to-end with patented PurePath Technology®. As a result, the metrics that are captured are connected; the transaction provides the context that reveals the relationship among data points. PurePath captures timing and code-level context for every transaction, method, and service across all tiers — no matter what your application environment.

Infrastructure data is also connected, as it relates to the application and transaction, with Smartscape® Technology. It detects billions of causal dependencies between websites, applications, services, processes, containers, hosts, networks, and cloud infrastructure. The entire application topology is visualized in an interactive infographic.

Understanding how data is connected means dependencies can be inferred with significantly more accuracy than is otherwise possible.

Algorithms that learn

Dynatrace uses a combination of proven A.I. technologies found in advanced applications such as Google search, Facebook, and modern fraud detection programs. These algorithms become more accurate over time as they automatically learn about the environment. A brief discussion of some of the more significant algorithms follows.

1. Learn what's important

Dynatrace's A.I. engine uses a modified and highly scalable Random Surfer model to determine an object's relevance. The result of the ranking algorithm ultimately determines the probability of root cause candidates. Google's famously effective PageRank algorithm is based on the same Random Surfer model.

2. Learn how are things related and interconnected

A graph database is used for analyzing complex relationships and interconnections. For example, real time event processing and causation detection employ a weighted and directed graph model. A graph model is also used to track topology discovery and mapping.

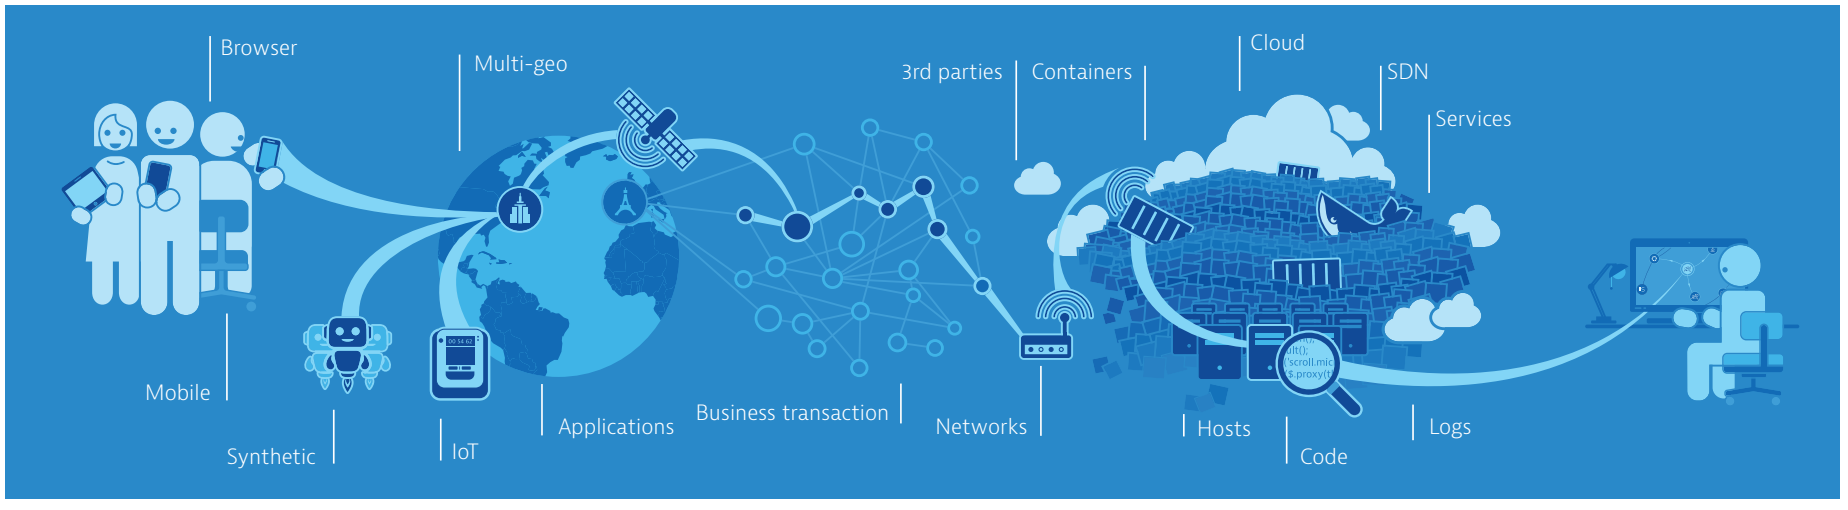

Dynatrace's A.I. relies on complete and connected data

3. Learn what's normal

Multi-dimensional baselining and advanced algorithms power Dynatrace's prediction-based anomaly and frequent issue detection.

- **Prediction-based anomaly detection**

  Anomaly detection is an effective means of identifying unexpected events and measurements. The term "unexpected" can also be read as "statistically improbable," which is why anomaly detection depends heavily on deep knowledge of a system's baseline performance and behavior. This is why Dynatrace monitors entire technology stacks end-to-end, capturing the baseline performance and behavior of applications, services, infrastructure components, and more, capturing metrics across millions of entities related to availability, error rates, response times, service load, user traffic, and resource dependencies.

  While baselining is used to automatically detect anomalies in the response times and error rates of applications and services, a prediction-based methodology is used to detect abnormalities in application traffic and service load. This is because traffic and load are entirely dependent on daily, seasonal, and business-cycle related patterns, such as weekends/workweeks, workday/evening hours, and holiday-driven customer activity. Black Friday is a great example of an extraordinary seasonal event that occurs on an annual cycle.

- **Frequent issue detection**

  There are situations where application performance is impacted as result of normal occurrences such as regular backup jobs that cause a CPU spike, or an unimportant disk that has been full for weeks. The A.I. engine in Dynatrace is able to differentiate unhealthy situations that are caused by normal versus abnormal events.

Expert knowledge built-in

Expert knowledge from Dynatrace is applied to these algorithms to make them smarter. It contains hardcoded information about how components can be affected by different topological connections. Some are obvious, for example, that it's likely that a process can affect another process on the same host or hypervisor; they do not affect each other if they are not on the same host.

Automation is key to scalability

As has been shown, Dynatrace delivers a level of automation that greatly exceeds what is possible with other APM solutions. It's important to recognize that, while it automates the processes that are typically the most time consuming for IT operations professionals, the primary benefit is not merely a productivity gain, however large it may be. Without automation, managing emerging hyperscale together with complex hybrid environments is virtually impossible. With Dynatrace, organizations can support emerging environments while maintaining support for existing applications, and even expanding the scope to include apps that are presently unmanaged.

Beyond dashboards: Extending access to insights with A.I.

Dynatrace is extending its already unprecedented level of automation with davis — an A.I.-based virtual assistant that changes how operators work, interfacing with them as an extension of the team. Davis's conversational interface uses advanced natural language processing and data-driven analysis to provide answers and insights instead of raw data. It learns who you are and customizes responses. Interact with your Dynatrace environment via Amazon Alexa and Slack, or integrate it into your own fabric via open APIs. Davis is even revolutionizing the way IT delivers application intelligence insights into the hands of non-technical business leaders.

Dynatrace Digital Performance Platform — it's digital business...transformed.

Successfully improve your user experiences, launch new initiatives with confidence, reduce operational complexity and go to market faster than your competition. With the world's most complete, powerful and flexible digital performance platform for today's digital enterprises, Dynatrace has you covered.

Learn more at dynatrace.com

Dynatrace has redefined how you monitor today's digital ecosystems. AI-powered, full stack and completely automated, it's the only solution that provides answers, not just data, based on deep insight into every user, every transaction, across every application. More than 8,000 customers use Dynatrace to optimize customer experiences, innovate faster and modernize IT operations with absolute confidence.

03.06.17 2066_WP_jw_US