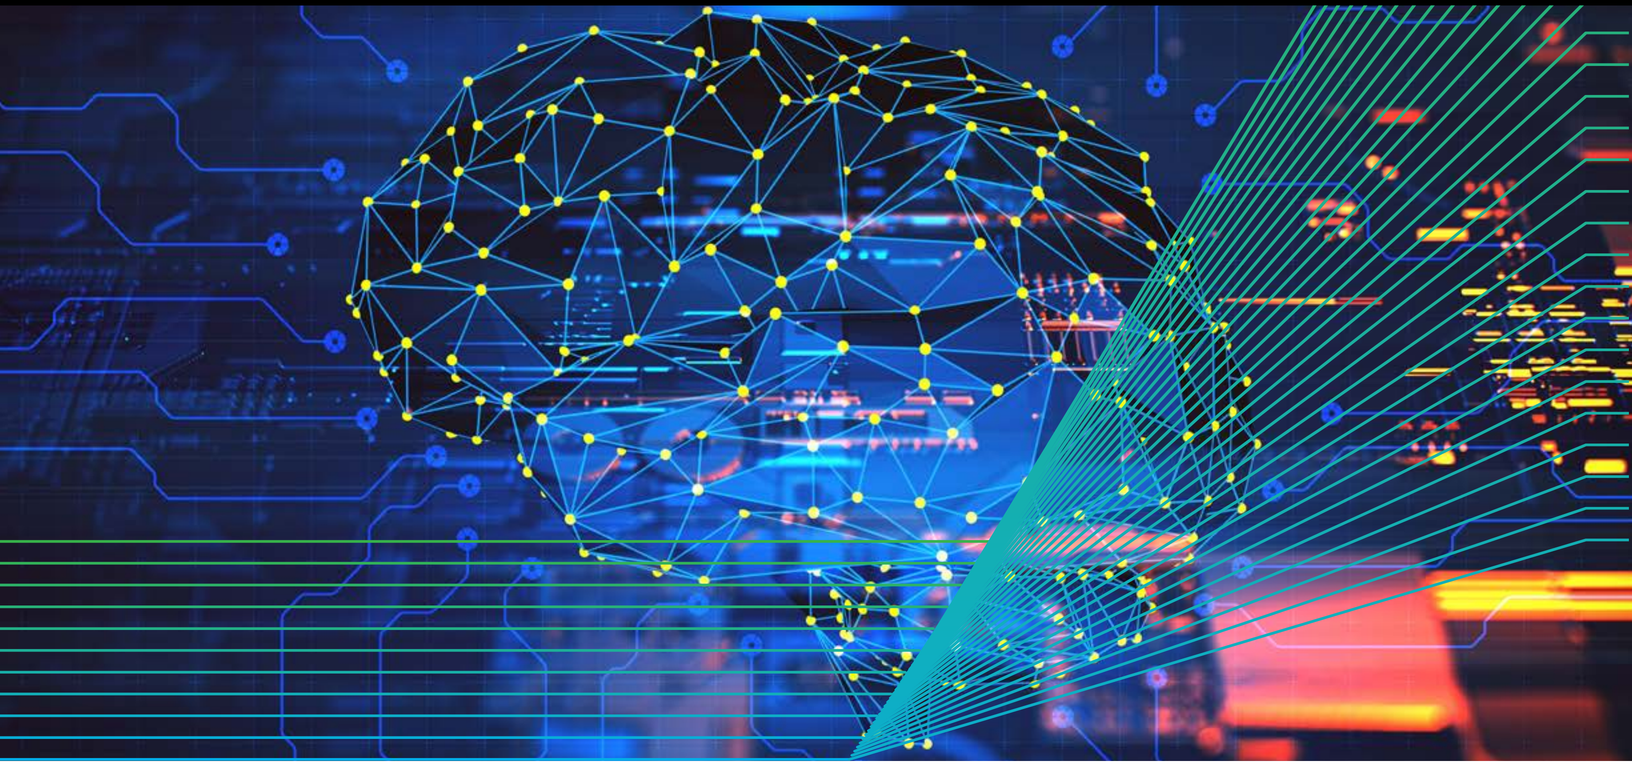

Applying Machine Learning to Boost Digital Business Performance

Assuring performance for multi-layered, architecturally-complex digital applications is a challenge. To address this, we recommend using AI to shift the focus from fixing defects to predicting them.

Executive Summary

Applications today mimic the business landscape, where several entities come together as an ecosystem to digitally enable products and services for those who need them. Modern applications are built on a multilayered architecture, comprising a front end (the UI layer), middle level (APIs and microservices) and back-end layer — all supported by a database.

For instance, the ride-sharing app Uber uses third-party APIs for navigation and payment, which together function via a shared user database. However, greater architectural complexity leads to an exponential increase in the number of failure points. This means if an exception occurs in one layer, the entire application could collapse. If an API for navigation responds slowly, the Uber app would be unable to deliver the expected customer experience, even though the app functions as it should. The reason: In a multilayer model, performance is distributed across layers, and must be orchestrated across the technology ecosystem.

To assure performance of applications built on digital ecosystems, businesses need to disengage from traditional quality assurance (QA) approaches that typically focus on automating regression and triaging. Instead, by engineering performance from the start, businesses ensure a consistently high-performing application. This white paper proposes a machine-learning-based AI approach to assuring performance, which analyzes historical data logs and predicts defects well before they may occur. It builds off of our recently-published whitepaper, "From Data to Insights: How IT Operations Data Can Boost Quality."

A machine-learning (ML) algorithm that identifies defects when code is written helps engineer performance earlier in the lifecycle. The algorithm culls performance-related data from performance logs, server logs, business outcomes (the differential between expected and delivered) and customer feedback.

PREDICTING FAILURE TO ENSURE SUCCESS

To meet users' lofty expectations, digital success hinges on consistent and reliable high performance. The British Broadcasting Company's recent loss of 10% of site visitors for every one second delay in loading its website proves the point.[1] To meet the mandate, QA needs to move beyond automation, which only speeds defect fixes. To do this, QA teams must move from a reactive to a proactive process.

A machine-learning (ML) algorithm that identifies defects when code is written helps engineer performance earlier in the lifecycle. The algorithm culls performance-related data from performance logs, server logs, business outcomes (the differential between expected and delivered) and customer feedback. A repository of issues (defects) can be fed to the algorithm, to correlate and predict defects. By triggering alerts, suggesting self-healing paths and tracing defects back to the underlying root causes, machine learning helps businesses orchestrate the performance of complex applications across layers. (To learn more, read our white paper, "Making a Quantum Leap with Continuous Analytics-Based QA.")

ENGINEERING PERFORMANCE THROUGH MACHINE LEARNING

Application performance is measured by response time, data-fetching time between the back-end layer and the database, etc. Monitoring these elements over a prescribed time period creates a rich log of performance parameters, which, when sourced into a machine-learning algorithm, performs predictive analysis that leads to actionable insights.

Time-Series Analysis

Time-series analysis is a linear machine-learning technique which provides an overview of how each layer within the application behaves when a certain number of users access it. Time-series analysis can help QA teams identify the layer that is most susceptible to exceptions or defects during peak load times, as well as establishing interdependencies among the various layers.

Upon analysis, the machine-learning algorithm generates a pattern and creates a heat map of exceptions across each layer (see Figure 1, next page). It would further predict when an exception could occur and trace it back to its root cause.

> The machine-learning algorithm reconciles time-bound data points into a comprehensible pattern. By correlating this data across layers, a performance-modeling function can be created to reflect the cumulative factors impacting performance across layers such as HTML, app server and microservices.

Time-Series Representation of Exceptions Across Layers

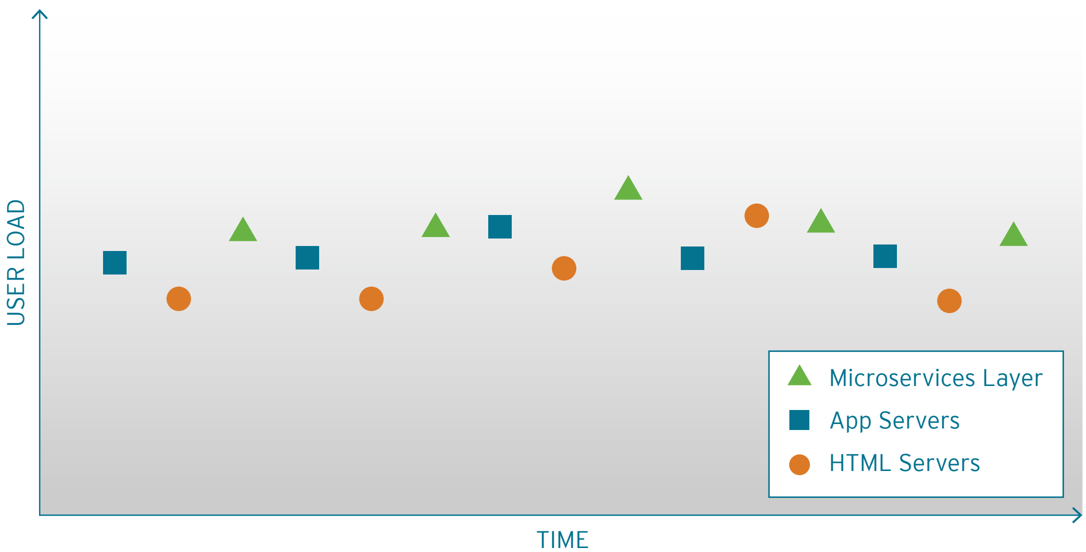

Figure 1

Key questions an ML-based algorithm addresses include:

- Where will an exception occur?
- In how many minutes or seconds will the exception occur?
- In which layer will the exception occur?

Arriving at Performance Modeling Function

As an application is under development, several layers are typically added to enrich the customer experience. However, with each new layer, performance of the overall application can be altered, often making it difficult for QA teams to identify the exact impacted layer. Performance modeling solves the issue by helping QA teams with a root-cause analysis of an exception, which leads them to the layer impacting the performance (see Figure 2, next page).

The machine-learning algorithm reconciles time-bound data points into a comprehensible pattern. By correlating this data across layers, a performance-modeling function can be created

Performance Monitoring by the Numbers

App Response Time (F)	# of Concurrent Users (G)	HTML5 Server (H)	Cookie Server (I)	App Server (J)		Microservice Response Time (K)		Transaction Timestamp
		Latency Value	CPU Utilization %	# of Active Sessions	CPU Utilization %	Response Time	CPU Utilization %	
5 sec	1200	40 ms	60%	200	74%	3 sec	73%	2017-11-22 16:47:08
7 sec	2200	70 ms	75%	350	89%	5.5 sec	88%	2017-11-21 16:47:08
6 sec	1700	65 ms	78%	280	80%	5 sec	80%	2017-11-20 16:47:08
–	–	–	–	–	–	–	–	2017-11-19 16:47:08
12 sec	2400	90 ms	95%	410	95%	11 sec	92%	2017-06-22 16:47:08

Performance Modeling Function ***F(x) = G(x) + H(x) + I(x) + J(x) + K(x)***

Figure 2

to reflect the cumulative factors impacting performance across layers such as HTML, app server and microservices.

Numerous technological attributes such as network bandwidth and central processing unit utilization could impact a particular layer's performance. Modeling would help identify the performance threshold required at various layers to achieve the intended performance and identify the bottlenecks (see Figure 3).

Preventing Performance Issues Through Log Analysis

Building on the performance modeling function, the machine-learning algorithm prompts the developer to consider the impact each line of code has on the overall performance as the application is placed into production. The developer can then take remedial measures to ensure that the code is of shippable quality (see Figure 4, next page).

This approach can be used to establish a pattern between an application's performance degradation and respective file changes.

How Performance Management Is Calculated

Performance Modeling Function ***F(x) = G(x) + H(x) + I(x) + J(x) + K(x)***

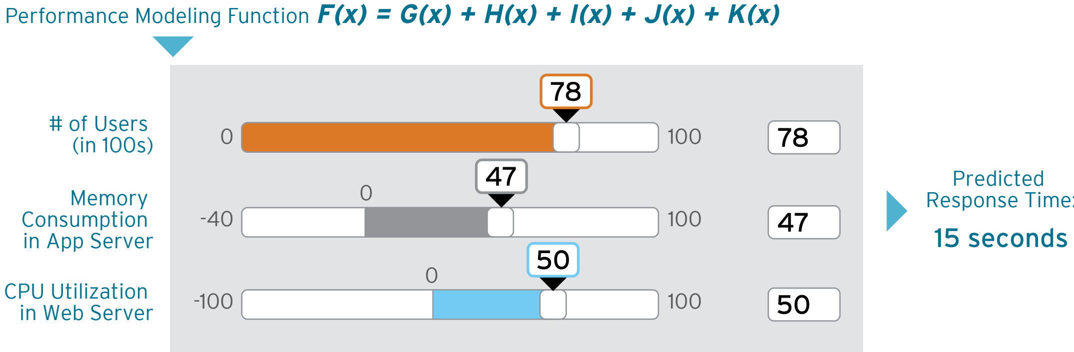

Figure 3

How Pattern Identifications Come into Play

Identify pattern between files and transactional performance from historical Sprint data.

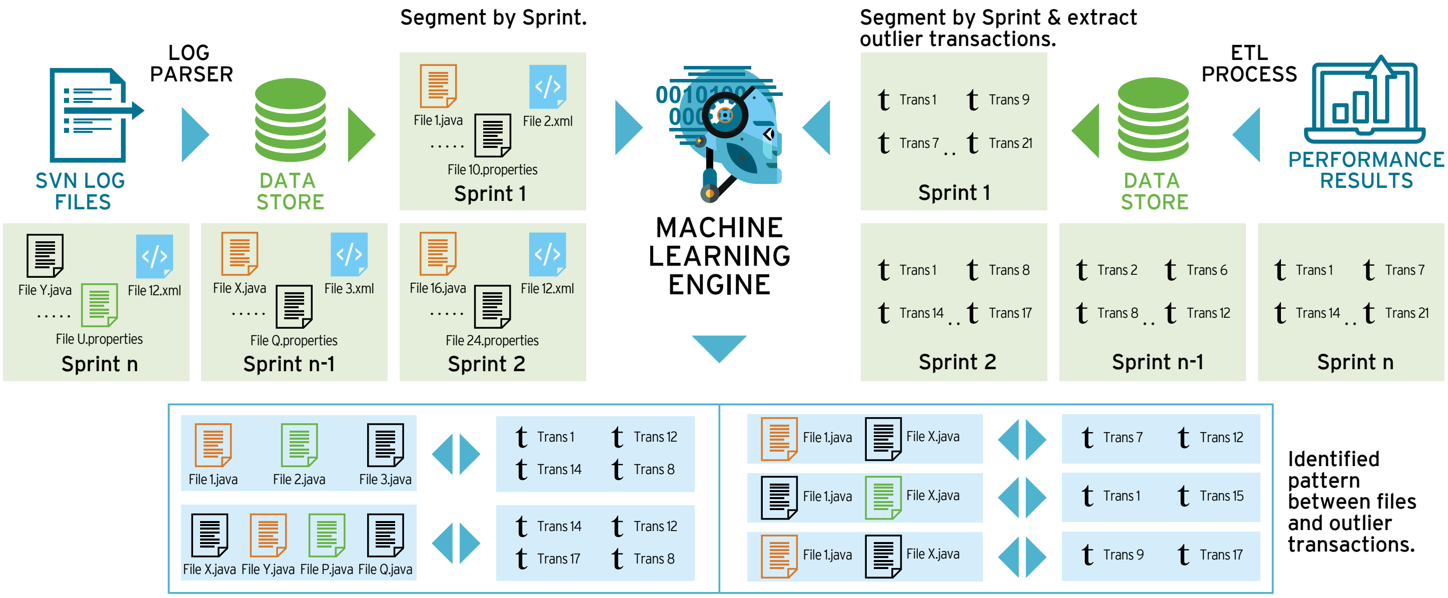

Figure 4

THE WAY FORWARD

These are linear machine learning techniques that businesses can deploy to ensure architecturally-complex applications perform up to customer expectations. However, an advanced genre of machine learning — known as deep neural learning — could induce self-healing patterns that fix defects in code without involving a developer or a tester.

The use cases for deep neural learning are under consideration by machine-learning experts. However, as a more inclusive and well-rounded approach to machine learning, deep neural learning involves domain SMEs who bring added value in their understanding of the business, along with technical experts who train the machine-learning mechanism to act independently.

With deep neural learning, enterprises can assure end-to-end customer journeys without human intervention. This would allow for better orchestration of performance through a community of QA professionals, domain SMEs and industry experts.

With deep neural learning, enterprises can assure end-to-end customer journeys without human intervention. This would allow for better orchestration of performance through a community of QA professionals, domain SMEs and industry experts.

FOOTNOTE

[1] www.creativebloq.com/features/how-the-bbc-builds-websites-that-scale.

ABOUT THE AUTHORS

Vikul Gupta

Market Head for Digital Assurance, Quality Engineering & Assurance Practice, Cognizant

Vikul Gupta is Market Head for Digital Assurance within Cognizant's Quality Engineering & Assurance business unit. He has more than 18 years of experience in strategy, management, delivery and development on multiple technologies, and deep domain experience in analytics, cloud, DevOps, mobile and the Internet of Things (IoT), with roles ranging from key strategist to solution architect. Currently, Vikul is leading a team that is responsible for creating differentiated service offerings and intellectual property that defines the go-to-market strategy for digital assurance for analytics, cloud, DevOps, mobile and the IoT. He is a graduate of the National Institute of Technology, Surat. Vikul can be reached at Vikul.Gupta@cognizant.com.

Vasanthkumar Velayudham

Analytics, Solution Architect, Quality Engineering & Assurance Practice, Cognizant

Vasanthkumar Velayudham is the Analytics, Solution Architect leading the effort for quality intelligence and predictive analytics/machine learning within Cognizant's Quality Engineering & Assurance business unit. Quality insight is a key initiative that he drives as part of his focus on digital assurance initiatives. Vasanth has 12-plus years of experience in testing, analytics, technical solution architecture, experimental design, product management, technical product planning, product innovation and market research on enterprise software products. Moreover, he is responsible for creating new machine-learning analytics solutions to improve applications' quality. Vasanth is a graduate of SSN College of Engineering, Anna University. He can be reached at Vasanthkumar.Velayudham@cognizant.com.

Saravanan Palanivelu

Analytics, Data Scientist, Quality Engineering & Assurance Practice, Cognizant

Saravanan Palanivelu is the Analytics, Data Scientist for Quality Intelligence and Predictive Analytics/Machine Learning, within Cognizant's Quality Engineering & Assurance business unit. Quality insight is a key initiative that he drives as part of his focus on digital assurance initiatives. Saravanan has more than 13 years of experience in testing, analytics, machine learning solution architecture, product management, technical product planning, product innovation and market research on enterprise software products. Moreover, he is responsible for creating new machine-learning analytics solutions to improve applications' quality. Saravanan graduated from National Engineering College, Manonmaniam Sundaranar University. He can be reached at Saravanan.Palanivelu@cognizant.com.

ABOUT QE&A

Cognizant's Quality Engineering & Assurance (QE&A) Practice is an independent business unit within Cognizant, offering a comprehensive suite of assurance services spanning QA consulting, quality engineering, environment management and digital QA. With the establishment of this business unit in 2001, Cognizant pioneered the concept of a practice dedicated to independent verification and validation. Since then, the practice has grown organically to more than 31,000 professionals, making it one of the world's largest providers of QA services. Cognizant's QE&A Practice is organized by areas of industry specialization to leverage Cognizant's domain expertise and provide domain-aligned QA services. Today, more than 550 global, regional and local enterprises trust Cognizant QE&A with their QA needs. For more information, please visit www.cognizant.com/quality-engineering-and-assurance.

ABOUT COGNIZANT

Cognizant (Nasdaq-100: CTSH) is one of the world's leading professional services companies, transforming clients' business, operating and technology models for the digital era. Our unique industry-based, consultative approach helps clients envision, build and run more innovative and efficient businesses. Headquartered in the U.S., Cognizant is ranked 195 on the Fortune 500 and is consistently listed among the most admired companies in the world. Learn how Cognizant helps clients lead with digital at www.cognizant.com or follow us @Cognizant.

Cognizant

World Headquarters
500 Frank W. Burr Blvd.
Teaneck, NJ 07666 USA
Phone: +1 201 801 0233
Fax: +1 201 801 0243
Toll Free: +1 888 937 3277

European Headquarters
1 Kingdom Street
Paddington Central
London W2 6BD England
Phone: +44 (0) 20 7297 7600
Fax: +44 (0) 20 7121 0102

India Operations Headquarters
#5/535 Old Mahabalipuram Road
Okkiyam Pettai, Thoraipakkam
Chennai, 600 096 India
Phone: +91 (0) 44 4209 6000
Fax: +91 (0) 44 4209 6060

TL Codex 3672