




ConversationalGeek®

Conversational Application Management for Federal Government

By Andrey Zhuk (Field CTO for Cybersecurity)



**In this
book, you
will learn:**

- What application performance management really means for Federal agencies
- How observability can help keep you ahead of Federal mandates
- Some key use cases for automatic and intelligent observability platforms in Federal Government

**2nd
Edition**

Sponsored by



Sponsored by Dynatrace

Dynatrace Federal provides federal agencies with deep intelligence about the state of the software running their missions. Our secure platform uses AI to enable observability – contextual awareness into the ‘what’ and the ‘why’ behind application performance and security, the status of the underlying IT infrastructure, and the experience of users wherever they’re located. Our platform continuously monitors data supporting mission-critical operations, freeing IT teams for higher value work. We’re FedRAMP-authorized, FIPS 140-2 authorized and we’re the only full stack, completely automated platform able to provide the depth of software intelligence needed to help simplify cloud complexity and accelerate agencies’ digital transformation.



For more details visit
www.dynatrace.ai/government/

Conversational Application Management in Federal Government (Second Edition)

By Andrey Zhuk

© 2022 Conversational Geek



ConversationalGeek®

Conversational Application Management in Federal Government (Second Edition)

Published by Conversational Geek® Inc.

www.conversationalgeek.com

All rights reserved. No part of this book shall be reproduced, stored in a retrieval system, or transmitted by any means, electronic, mechanical, photocopying, recording, or otherwise, without written permission from the publisher. No patent liability is assumed with respect to the use of the information contained herein. Although every precaution has been taken in the preparation of this book, the publisher and author assume no responsibility for errors or omissions. Nor is any liability assumed for damages resulting from the use of the information contained herein.

Trademarks

Conversational Geek, the Conversational Geek logo, and J. the Geek are trademarks of Conversational Geek®. All terms mentioned in this book that are known to be trademarks or service marks have been appropriately capitalized. We cannot attest to the accuracy of this information. Use of a term in this book should not be regarded as affecting the validity of any trademark or service mark.

Warning and Disclaimer

Every effort has been made to make this book as complete and as accurate as possible, but no warranty or fitness is implied. The information provided is on an “as is” basis. The author and the publisher shall have neither liability nor responsibility to any person or entity with respect to any loss or damages arising from the information contained in this book or programs accompanying it.

Additional Information

For general information on our other products and services, or how to create a custom Conversational Geek book for your business or organization, please visit our website at ConversationalGeek.com

Publisher Acknowledgments

All of the folks responsible for the creation of this guide:

Author:	Andrey Zhuk
Project/Copy Editor:	Pete Roythorne
Content Reviewer(s):	Carolyn Ford Willie Hicks

Note from the Author

At Conversational Geek, we like to break down complex tech into simple terms and do it in a fun way. We will explain why a given piece of technology matters, what problems it solves, and how it can be applied to your own environment whether at work or elsewhere.

This eBook provides a 101-level introduction to the world of application performance monitoring (APM) and observability, and how these technologies apply to United States Federal Government. Software is transforming how federal agencies do business. Mission success of federal agencies is driven by applications. Having insight into the performance, health, and security of these applications and underlying infrastructure is critical, whether this infrastructure is hosted in private, public, or hybrid clouds.

In the past two years, the U.S. Government issued several mandates requiring agencies to implement APM and application security capabilities by the end of fiscal year 2024. In this eBook, we review how APM and observability align with various federal initiatives, including Cybersecurity Executive Order (EO) 14028, and follow-on Office of Management and Budget (OMB) memorandums like M-22-09, M-21-30, M-21-31, M-22-16, and M-22-18. We hope this eBook will equip you with nuggets of knowledge you can immediately use in conversations with colleagues, customers, and even your boss.

Andrey Zhuk
Field CTO for Cybersecurity, CTG Federal



The “Conversational” Method

We have two objectives when we create a “Conversational” book: First, to make sure it’s written in a conversational tone so that it’s fun and easy to read. Second, to make sure you, the reader, can immediately take what you read and include it into your own conversations (personal or business-focused) with confidence.

These books are meant to increase your understanding of the subject. Terminology, conceptual ideas, trends in the market, and even fringe subject matter are brought together to ensure you can engage your customer, team, co-worker, friend, and even the know-it-all Best Buy geek on a level playing field.

“Geek in the Mirror” Boxes


We infuse humor into our books through both cartoons and light banter from the author. When you see one of these boxes it’s the author stepping outside the dialog to speak directly to you. It might be an anecdote, it might be a personal experience or gut reaction and analysis, it might just be a sarcastic quip, but these “geek in the mirror” boxes are not to be skipped.



Greetings. Within these boxes I can share just about anything on the subject at hand.

Read 'em!

Introduction to APM and Observability in Federal Government



Digital transformation is happening everywhere. The world is now hyperconnected with applications impacting every aspect of our lives: how we work, collaborate, get medical care, file taxes, and even how we get our utilities delivered. All these applications are running on top of compute infrastructures that are no longer just a couple of servers in a rack, but dynamic, hybrid and multicloud environments, containing hundreds of technologies, millions of lines of code, and billions of dependencies. Digital transformation isn't simply about lifting

and shifting apps to the cloud. It is a fundamental shift in how applications are built, deployed, and operated – and, of course, the end user expects all of this to work perfectly, all the time.

Applications have grown in scale and complexity. Traditionally, we managed these through a set of disparate point tools, each requiring human involvement. As applications grew out of the traditional data center into private cloud (and subsequently into public and hybrid clouds), the old way of doing things could no longer scale. Both commercial enterprise and federal agencies were forced to re-think how they maintain the health of their applications, while also monitoring the stability of the underlying infrastructure and ensuring appropriate Quality of Experience (QoE) for the end user. This led to the emergence of APM tools that would ingest multiple telemetry feeds and provide critical insight into application performance. APM provides data on dashboards to troubleshoot application performance in production based on known or expected system failures. Observability collects data across multiple layers of software architecture and analyzes it in real-time. This observability gives users insight into system health and behavior across the entire stack, influencing user experience and business decisions. Over time, these tools have grown into software platforms that enable observation and analysis of application health, performance, security, and user experience.

You may be wondering how this ties back to U.S. Federal Government? While some U.S. public sector agencies are lagging the commercial sector in transitioning their monolithic applications to cloud-native architectures, digital transformation is gaining momentum across all the Fed, with “cloud-first” strategies mandated at the highest levels.



Starting in 2011 with White House publication of “Cloud First” policy¹ and subsequent 2019 “Cloud Smart” practical implementation guidance², public cloud adoption has been at the center of most Information Technology (IT) projects in Fed.

There are several government initiatives and mandates driving digital transformation and implicitly point to the need for APM and observability. The biggest one is Cybersecurity EO 14028 and follow-on OMB memos. Until January 2022, the Cyber EO was mostly applicable to civilian unclassified environments. That changed on January 19, 2022, when the President signed Cybersecurity National Security Memorandum 8 (NSM-8). NSM-8 makes Cyber EO applicable to National Security Systems, expanding EO scope to the Department of Defense (DoD), Intelligence Community (IC), and classified networks. Following is a summary of key milestones:

- **NIST SP 800-207 (August 2020)**³ – National Institute of Standards and Technology (NIST) formalizes the Zero Trust architecture

¹ Vivek Kundra, U.S. Chief Information Officer, Federal Cloud Computing Strategy, February 8, 2011, https://obamawhitehouse.archives.gov/sites/default/files/omb/assets/egov_docs/federal-cloud-computing-strategy.pdf (“Cloud First”)

² Suzette Kent, U.S. Federal Chief Information Officer, Federal Cloud Computing Strategy, June 24, 2019, <https://www.whitehouse.gov/wp-content/uploads/2019/06/Cloud-Strategy.pdf> (“Cloud Smart”)

³ <https://csrc.nist.gov/publications/detail/sp/800-207/final>

- **Cybersecurity EO 14028 (May 2021)**⁴ – The White House calls for cybersecurity modernization across the U.S. Federal Government; the EO is split into eight sections:
 1. EO policy applies to both IT and Operational Technology (OT), like Industrial Control Systems (ICS) and Internet of Things (IoT) devices
 2. Agencies must share cyber information with the private sector and vice versa
 3. Agencies must implement Zero Trust according to NIST SP 800-207, including adopting Multi-Factor Authentication (MFA), encrypting everything, and managing risk using analytics powered by artificial intelligence (AI) (expanded by **OMB M-22-09**⁵)
 4. Agencies must improve the software supply chain with focus on DevSecOps to secure software development (expanded by **OMB M-21-30**⁶ and **M-22-18**⁷)
 5. Establish cyber safety review boards to include federal departments and the private sector
 6. Develop standardized federal incident response playbooks and call for agencies to

⁴ <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/05/12/executive-order-on-improving-the-nations-cybersecurity/>

⁵ <https://www.whitehouse.gov/wp-content/uploads/2022/01/M-22-09.pdf>

⁶ <https://www.whitehouse.gov/wp-content/uploads/2021/08/M-21-30.pdf>

⁷ <https://www.whitehouse.gov/wp-content/uploads/2022/09/M-22-18.pdf>

adopt security orchestration, automation, and response (SOAR) tools

7. Mandate U.S. government-wide endpoint detection and response (EDR) deployment (expanded by **OMB M-22-01**)⁸
8. Logging, log retention and log management; basically, log everything (expanded by **OMB M-21-31**, Maturity Model for Event Log Management)⁹
 - **NSM-8 (January 2022)**¹⁰ – Enforces EO 14028 for DoD, IC and classified networks
 - **OMB M-22-09 (January 2022)** – Makes Zero Trust official U.S. government policy

While the executive order laid out a general vision for federal agencies to “advance toward Zero Trust architectures,” the Federal Zero Trust architecture strategy details specific, directed actions and aggressive implementation timelines as shown below. In short, all agencies must show progress on the various Zero Trust objectives outlined in the guiding policies by September 30th, 2024.

⁸ <https://www.whitehouse.gov/wp-content/uploads/2021/10/M-22-01.pdf>

⁹ <https://www.whitehouse.gov/wp-content/uploads/2021/08/M-21-31-Improving-the-Federal-Governments-Investigative-and-Remediation-Capabilities-Related-to-Cybersecurity-Incidents.pdf>

¹⁰ <https://www.nsa.gov/Press-Room/News-Highlights/Article/Article/2904637/president-biden-signs-cybersecurity-national-security-memorandum/>



Figure 1 – OMB memos lay out a timeline for both planning and implementing Zero Trust

How does APM and observability play into the security mandates? Zero Trust calls for *continuous verification of the operational picture* using real-time information from multiple sources to determine access and other system responses. In today’s app-centric world, most of this data will come from inside the application and the underlying IT infrastructure. The only way to elegantly pull together all this data and produce intelligent and actionable insights is with an AI-powered observability platform.

It is worth touching on the User Experience EO¹¹ from December 2021 and follow-on President’s Management Agenda (PMA)¹² vision. The User Experience EO outlines government-wide and agency-specific actions to improve QoE as users are interacting with various federal government applications online. Improving QoE is impossible without an observability solution that provides digital experience monitoring (DEM) and business analytics.

¹¹ <https://www.whitehouse.gov/briefing-room/presidential-actions/2021/12/13/executive-order-on-transforming-federal-customer-experience-and-service-delivery-to-rebuild-trust-in-government/>

¹² <https://www.performance.gov/pma/>

APM – The big misconception

As a solutions architect for the U.S. public sector, I talk to many customers, spanning many agencies across defense, civilian, and the sciences. There are two things all these agencies have in common: they run a lot of their own mission-specific applications, and they all claim to do APM. When pressed about what their APM program entails, most agencies admit to some sort of tooling strategy, sometimes using the “best of breed” tools, but typically just using what’s available.



What do we mean by “tooling”?

Tools for server monitoring, Windows monitoring, Linux monitoring, database monitoring (a different tool for different database), network monitoring, packet capture, logs analysis, tools provided by cloud service providers, and the list goes on.

Exhausting just reading that list, isn’t it? Unfortunately, there are many missing pieces to this approach and without those missing pieces you don’t have APM.

Infrastructure and operations teams have been deploying monitoring tools since the dawn of IT to track the performance of networks, storage, compute, and applications. However, as applications have adopted microservices architectures and moved to the cloud, point monitoring tools are showing their limitations. IT teams are spending vast amounts of time learning, deploying, configuring, and managing their own local tools. And despite all these tools, there isn’t a single source of truth. There is just lots of data, dashboards, war rooms, and finger-pointing. Furthermore, these traditional tools fail to provide value in the new world of the cloud, containers,

Kubernetes, and microservices architectures. As one Federal IT leader told me, *“It feels like we monitor everything, but see nothing.”*

The challenges of a typical “tooling” approach to APM are the following:

- Tools are difficult to scale and maintain effectively.
- IT staff invest lots of time and manual effort.
- Tooling requires manual instrumentation and code changes.
- There is no unified view of application and underlying infrastructure.
- There are too many alerts and meaningless metrics with no context or intelligence.
- Agencies end up with hundreds of dashboards.
- The tribal knowledge from managing a DIY, multi-tool approach makes it difficult to maintain the system when key team member leaves.
- The true cost could be significantly higher than the face cost of the tools themselves.

From APM to observability

If implementing multiple point tools to monitor performance of an application is not APM, then what is? A few years back, industry experts defined APM as a solution that provided the following capabilities:

- It monitors and tracks the performance and response time of an application.
- It creates a baseline of performance metrics and alert administrators when performance varies.

- It provides visual data to better understand the performance metrics (aka dashboards).
- It assists in root cause analysis of application performance issues.

But those were the ways of old. Let's look at APM in 2022, as Gartner defines it¹³:

“Application performance monitoring (APM) is a suite of monitoring software comprising digital experience monitoring (DEM), application discovery, tracing and diagnostics, and purpose-built artificial intelligence for IT operations.”

Gartner is onto something here with terms like “digital experience” and “AI” in play. The future of APM lies in AI and automation. What we've seen in the industry start as simple monitoring using a variety of tools has evolved into what we now call “observability”. What is “observability”? The classic systems engineering definition:

“Observability is a measure of how well internal states of a system can be inferred from knowledge of its external outputs.”

In the world of IT performance monitoring, *external outputs* are the raw data types we observe being output by the IT system supporting an application. These external outputs consist of three principal data types: *metrics, traces, and logs*. These are also referred to as the three pillars of observability.

¹³ <https://www.gartner.com/en/information-technology/glossary/application-performance-monitoring-apm>

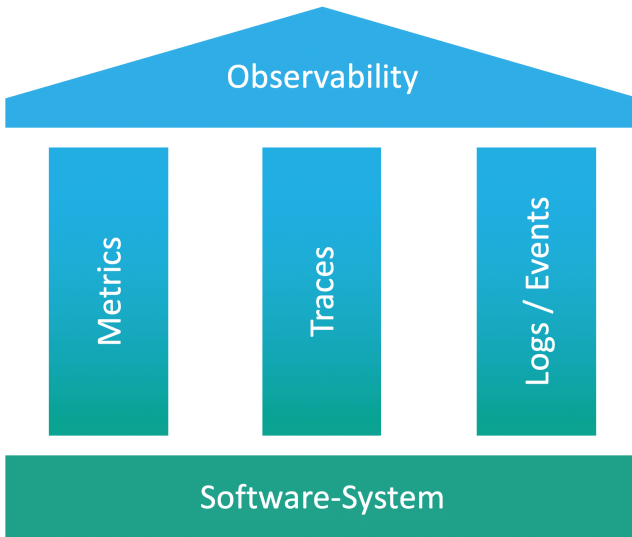


Figure 2 – The Three Pillars of Observability

The problem in modern environments is the sheer volume of data aggregated by monitoring and observability tools. By applying AI and machine learning (ML), modern APM and observability systems can analyze massive amounts of network and machine data to find patterns that are impossible for humans to identify. As Gartner defines it, the distinction between “APM vs. observability” is that “observability-centric solutions support an exploratory, analytics-driven workflow that may bear more resemblance to business intelligence (BI) than IT operations.”¹⁴ The term “software intelligence” refers to this new AI-driven approach: delivering APM, observability, infrastructure monitoring, DEM, digital analytics, and application security with AI at the core and in an all-in-one

¹⁴ <https://www.dynatrace.com/gartner-magic-quadrant-for-application-performance-monitoring-observability/>

platform. Gartner lists the following critical capabilities¹⁵ for a modern observability solution:

- Application debugging and distributed profiling (ADDP)
- Root cause analysis
- IT service and infrastructure monitoring
- Behavior analysis
- Business Analysis
- Runtime application self-protection (RASP)

These capabilities are critical for supporting the digital transformation and cybersecurity modernization efforts taking place across the entire U.S. federal space.

¹⁵ <https://www.dynatrace.com/gartner-critical-capabilities-for-application-performance-monitoring-observability/>

Use Cases for Observability in the U.S. Public Sector

Use Case 1 – Cloud migration

Starting with 2011 “Cloud First,” then 2019 “Cloud Smart,” and the wild success of the FedRAMP program, public cloud adoption is at the forefront of U.S. Federal IT. Cybersecurity EO 14028, M-22-09 memo on Zero Trust and M-22-16¹⁶ memo on budget priorities for FY 2024 further emphasize cloud adoption and securing cloud applications. Unfortunately, moving to the cloud is hard. You cannot just forklift workloads to AWS, Azure, or GCP. Well, technically, you can. But then you’ll be looking at steep cloud usage bills. Instead, migrating to the cloud involves rethinking, re-architecting, and re-building many of your agency’s legacy monolithic applications.

This is where an observability solution can help produce an inventory of your current environment and capture a performance baseline for the applications you are planning to migrate. It can help you understand the current system utilization, CPU, memory, storage requirements, etc. This way, when you start spinning up resources in the public cloud, you are not overprovisioning and are deploying only what’s necessary.

We have also seen situations where an agency is thinking of moving an app to the cloud, but after adopting an observability platform, they realized that no one was using the application. From here, the agency would decide whether they would refactor, adjust, or simply decommission the app. In making this decision, the agency would once again be relying on dependency mapping provided by the platform AI.

¹⁶ <https://www.whitehouse.gov/wp-content/uploads/2022/07/M-22-16.pdf>

Finally, an agency can use an AI-powered observability platform to confirm application performance after migration. There are two aspects to this use case. On one hand, it makes sense to verify that the migrated version performs similarly or better than the original. On the other hand, there are also situations where users still complain about performance after an application migration. Is it just perception or is there really performance degradation? The only way to know for sure is by using an observability solution to compare performance stats before and after migration.

Use Case 2 – Software supply chain risk management for DevSecOps

DevSecOps simply means security built into application development and operations, or DevOps. DevSecOps aims to make software development security-focused and places additional responsibility on developers to ensure their code doesn't have vulnerabilities, like the Log4Shell vulnerability¹⁷ that targeted the ubiquitous Log4J Java logging utility in 2021. In today's fast-paced environment of continuous integration and continuous delivery (CI/CD), agencies need an automatic and intelligent observability platform that makes it easy to integrate application security into the DevOps pipeline.

The White House in Section 4 of Cyber EO 14028¹⁸ calls out improving Software Supply Chain security, with M-21-30 calling agencies to protect critical software through enhanced security measures and M-22-18 specifically requiring software vendors to provide a Software Bill of Materials (SBOM). The EO emphasizes automation and employing automated tools for software development security. NIST formalized Cyber EO

¹⁷ <https://www.dynatrace.com/news/blog/what-is-log4shell/>

¹⁸ <https://www.nist.gov/itl/executive-order-14028-improving-nations-cybersecurity/software-security-supply-chains>

recommendations in May 2022 with their Special Publication (SP) 800-161 “Cybersecurity Supply Chain Risk Management Practices for Systems and Organizations.”¹⁹ Cybersecurity and Infrastructure Security Agency (CISA) even set up a web portal to serve as a “nexus for the broader for the broader set of SBOM resources across the digital ecosystem and around the world.”²⁰

This is why AI-powered observability is now more important than ever for Federal IT. Not only does it enable development teams to accelerate DevSecOps processes through automation and elimination of mundane work, but it also allows them to verify composition of software libraries used in mission-critical applications. An observability platform with AI at its core can also provide RASP to automatically detect vulnerabilities and block attacks by continuously analyzing applications, libraries, and code at runtime, in production and pre-production. An intelligent observability platform can automatically detect CI/CD changes, including multi-version deployments, runtime container updates, rollback, and elastic scaling with real-time detection, alerting and re-validation. Further, such a platform can uncover information on potentially compromised code through integration with Application Security Testing (AST) tools like Snyk, an open-source platform to find and fix security issues in code.

Use Case 3 – CI/CD quality gates

The CI/CD pipeline is a familiar concept for most in DevOps, but what exactly are “quality gates”? A quality gate is a benchmark that defines specific, measurable, and achievable success criteria a service must meet before it advances to the next phase of the software delivery pipeline. Of course, this needs

¹⁹ <https://csrc.nist.gov/publications/detail/sp/800-161/rev-1/final>

²⁰ <https://www.cisa.gov/sbom>

to happen automatically and intelligently. This is where the observability platform can automate quality checks earlier in the development lifecycle before software is released into production.



CI/CD Quality Gates should be part of any Agency's "shift left" strategy. Ultimately, "shift left" is all about using more production data earlier in the development lifecycle and, ultimately, answering the question: "Is this a good or a bad change that we are trying to push into production?"

Some examples of DevOps quality gates include the following:

- **Unit testing** – ensure that units of codes are without issues
- **Unit integration testing** – test how units of code work with one another
- **Deployment environment validation** – ensure the deployment environment is without issues; this may take the form of configuration checks for Infrastructure as Code (IaC) templates
- **Code build verification** – check if the build is meeting performance criteria; if not, stop the build
- **Static code analysis** – useful for checking code for security vulnerabilities, which can take the form of integrating with Snyk code for automatic static application security testing (SAST)
- **Post deployment testing** – conduct a functional test of the application; for example, can users log in and do required tasks?

An AI-powered Observability solution enables unbreakable continuous delivery of application updates and allows DevOps teams to automate the entire CI/CD pipeline:

- **Automate quality**
Performance-gates-as-code – perform quality checks in development before bad code gets to production
- **Automate deployments**
Validation-as-code – release higher quality applications more frequently
- **Automate operations**
Auto-remediation-as-code – self-heal bad deployments in production

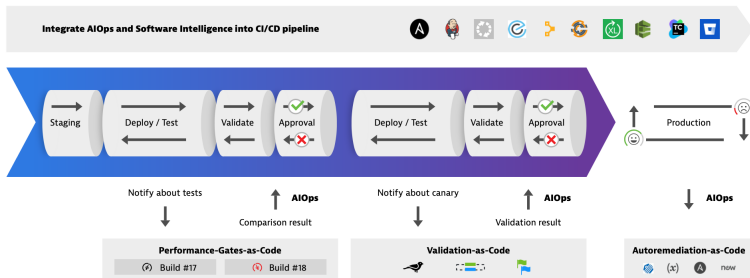


Figure 3 - Integrating AI-powered observability with the CI/CD pipeline

Use Case 4 – Mission continuity between contracts

In government, agencies often have a situation where a particular project was supported by one contractor and after the contract is re-competed, another bidder takes over. It often happens that contractor A built an application, but contractor B must operate and support it. These custodian transitions are rarely smooth. Documentation is scant and key personnel have moved on. Imagine you were awarded such a contract. How do you go about figuring out how the application works, the

process-to-process relationships, and what's considered normal performance?

An AI-powered observability solution can simply do it for you. It will map out dependencies both horizontally (between components of the same type) and vertically (between components of different types). You will get a real-time map of your entire application stack, end-to-end, from end-users' web browsers to your applications, down to containers, infrastructure, and cloud platforms. You'll know exactly which services are supported by what process and which processes are running on what hosts. Moreover, you'll know exactly how your environment is being used so you can plan a cloud migration strategy for greater efficiency.

Use Case 5 – Help mitigate IT staff shortages

Shortage of qualified IT staff is a real problem that we all know about first-hand. This is also a perfect use case for an AI-powered observability platform, allowing your agency to make the most of valuable human resources. By eliminating manual configuration, monitoring, and troubleshooting, your personnel can focus on critical thinking, strategic problem solving, and innovation. This kind of solution is capable of self-discovery and automatically adjusts to infrastructure changes. You can now rely on the observability platform to proactively detect and remediate issues, all without human intervention.

One civilian agency integrated their observability solution with Ansible to proactively detect and remediate memory leaks in a large enterprise application. The observability platform receives telemetry from a running process. It then uses AI to determine if the telemetry indicates a memory leak. If yes, the tool triggers Ansible to restart the faulty process. No human intervention required.

Use Case 6 – Mean time to recovery

Mean time to recovery, or MTTR, refers to the ability to recover a system back to operational state after a failure or degradation in service. Imagine an all-too familiar situation: an anomaly in your agency's application generates a flurry of alerts across various monitoring products in your environment. You are presented with a tangled mess of clues. How do you go about isolating the root cause of the issue?

The following graphic comes from a U.S. federal agency that recently implemented an observability solution²¹. Before implementing the observability solution, the time to restore exceeded the service level agreement, involved multiple phone calls, and required triage time. After implementing the observability solution, which integrated with their messaging services xMatters, PagerDuty, VictorOps, and Atlassian Opsgenie, no manual communications were necessary, and the service was restored well within the SLA time. The agency was able to triage the problem 40% faster, resolve issues 90% faster, and reduce help desk calls by 90%. AI-powered observability can help pinpoint the root cause of the issues. In many situations, remediation steps may not even have to be done by humans. We are moving to a world where AI-powered intelligent observability solutions can automatically recover the failed service. With an AI-powered observability solution, agencies can reduce MTTR from hours to minutes.

²¹ <https://www.dynatrace.com/resources/ebooks/aiops-done-right/>




Figure 4 – Implementing AI-driven observability tools reduces system MTTR from hours to minutes

Use Case 7 – Helpdesk automation

Helpdesk automation is another example of how AI-driven observability augment IT staff and simplify IT processes. A solution that can integrate with incident resolution systems, such as ServiceNow, Atlassian Jira, and Zendesk can automatically detect a problem, determine the root cause, open a service ticket, and trigger the system to take an appropriate action.

For example, the observability solution may know that a monitored application worked prior to a new code release, but now users are experiencing issues. The solution can automatically roll back code to the last working version. It will take care of all interdependencies and do it correctly every time. Imagine doing this manually. The room for human error is immense!

Use Case 8 – “Out-of-band” security through performance heuristics

Typically, when we think of application security, we envision tools and techniques like micro-segmentation between containers, Web Application Firewalls (WAFs) for the application front end, Network Detection and Response (NDR) for flow analytics, malware and vulnerability scanning for the workloads, and the list goes on. With the help of observability, we can add another technique to the list – detecting out-of-band (OOB) security attacks through baselining performance metrics (heuristics).

Let’s start with a simple use case. Imagine you are running a large enterprise application with thousands of microservices. An opportunistic attacker was able to get access to your environment (malicious insider) and spin up a Bitcoin miner. Why not? There are thousands of services running. No one is going to notice. And they are right! Without appropriate monitoring in place, it would be virtually impossible to detect.

Another use case is more subtle. Imagine a nation state crafting a Zero Day exploit of your agency. There are no signatures to scan for. The attacker is using low and slow methods of communication, so your network security has not caught on yet. How do you go about detecting such an attacker?

For both cases, an AI-powered observability solution can provide you with information about the normal baseline for application performance and the new post-breach baseline with suspicious processes running. The observability solution can pin-point the exact process that is out of the norm. It can uncover resource overutilization, creation of unrecognized new processes, or additional containers spinning up – anything outside the normal baseline. In this way you can use infrastructure performance heuristics to provide OOB security

protection. An additional security layer for your agency's defense-in-depth strategy.

The Big Takeaways

In today's world of ever-expanding applications, federal agencies are looking to make sense of the data produced by their hybrid, multicloud environments, and the "alert noise" created by multiple performance monitoring tools. These tools require a tremendous amount of human effort to set up, instrument, and manage. Today's level of complexity has surpassed what can be managed by humans manually. Automation and intelligence are required. As you rethink your agency's monitoring strategy going forward, you want to make sure you choose an automatic and intelligent observability solution that provides the following capabilities:

Advanced observability

- Automatic visibility at scale across hybrid and multicloud infrastructures, including metrics, logs, traces, code, user experience, and behavior data, all in context
- One single source of truth, not an alert storm
- Goes beyond dashboards, filters out the noise to give you precise answers

Continuous automation

- Continuous self-discovering automatic deployment for multicloud and dynamic containers, adjusting in real time as workloads change
- Auto-discover infrastructure and full-stack dependencies
- Eliminate manual configuration to save resources and reduce strain on IT staff

AI-assisted analysis

- Continuous analysis of dependencies for application and underlying private/hybrid cloud infrastructure
- Auto baselining and continuous anomaly monitoring
- Ability to provide the precise root cause of failures prioritized by business impact


Cross-team collaboration

- Single view across infrastructure and apps with precise answers, prioritized by business impact, for all stakeholders
- Capability to interconnect your agency teams, whereby technical staff and the agency administration are working from the same pane of glass

User experience and business analytics

- Real-time, outside-in perspective on user experience and how it impacts the agency mission
- Ensure applications are available and high performing by combining performance data, real user behavior, synthetic monitoring, and session replays
- Measure mission results with real-time visibility into key performance indicators (KPIs) with contextual analysis

Dynatrace – Bringing Intelligence to Your Applications



Dynatrace is a market leader in observability, providing APM, application security, digital experience, business analytics, DevOps cloud automation, and infrastructure monitoring. The Dynatrace Software Intelligence Platform delivers AI-powered observability and automation to help government agencies accelerate digital transformation. Dynatrace automatically discovers and captures high-fidelity data from applications, containers, services, processes, and infrastructure. It then automatically maps the billions of dependencies and interconnections in these complex environments. The Dynatrace AI engine, Davis®, analyzes this data and its dependencies in real-time to instantly provide precise answers, prioritized by business impact.

This level of automation and intelligence enables agencies to overcome the challenges presented by modern hybrid and multicloud environments. With an automatic and intelligent observability platform, teams can develop better software faster, automate operations, and deliver better business results for their agencies and their constituents. This is why many of the world's largest enterprises, including 72 of the Fortune 100, trust Dynatrace.

Dynatrace is a Leader in the 2022 Gartner® Magic Quadrant for APM and Observability²² and ranked #1 out of 19 vendors in the 2022 Gartner® Critical Capabilities for APM and Observability report in four of six use cases:

- IT Operations
- Digital Experience Monitoring
- DevOps/AppDev
- SRE/Platform Operations

For our customers in the U.S. federal government, Dynatrace achieved FedRAMP Moderate authorization in July 2020²³ and is in progress to achieve FedRAMP High.

²² <https://www.dynatrace.com/gartner-magic-quadrant-for-application-performance-monitoring-observability/>

²³ <https://www.dynatrace.com/news/press-release/dynatrace-secures-fedramp-moderate-impact-level-authorization/>



Figure 5 – Components of Dynatrace Software Intelligence Platform

How Dynatrace aligns with federal Zero Trust initiatives

Dynatrace works closely with U.S. Government customers to help them implement capabilities mandated by policies like cybersecurity and user experience EOs. Dynatrace empowers agencies with automatic and intelligent observability so they can do the following:

- Continuously monitor and capture all data from logs, metrics and end-to-end traces in context with user behavior and
- Apply AI to establish baselines and root-cause analysis
- Automatically identify anomalous and potentially threatening activity to enforce least privilege

Following is how Dynatrace extends observability to every device, browser, and application supporting every user:

- **Dynatrace OneAgent** automatically and continuously collects all relevant metrics and tracks all dependencies along your full application-delivery chain.
- **Dynatrace SmartScope** maps in real-time all the components and dependencies in an agency’s ecosystem, including what is predictable – hosts, networks, and other infrastructure – and unpredictable website, application, services, and process activity.

- **The Dynatrace Apdex** rating system is calculated for each discrete user action and application.
- **Dynatrace AI engine**, Davis, scours for risk vulnerabilities within applications running in the cloud, containers, virtual machines and traditional servers, while analyzing logs, metrics and traces for additional context to immediately pinpoint the root cause of issues.
- **Dynatrace integrates** with tools like ServiceNow and Ansible so agencies can rapidly resolve problems through auto-remediation, ensuring that applications are always on and fully optimized.

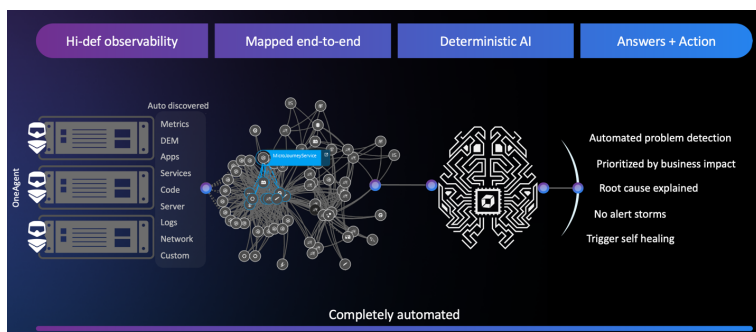


Figure 6 – How Davis works to deliver answers in real time

Many Dynatrace public sector customers have security restrictions that require them to operate in a secure or air-gapped infrastructure. To accommodate these requirements, Dynatrace offers the Dynatrace Managed mode of deployment, on premises. However, to provide SaaS-like experience, the solution can optionally connect to Dynatrace Mission Control for automatic product updates. For air-gapped networks, like those found in classified environments, Dynatrace can operate in a completely disconnected fashion, with product updates

delivered separately, as required by the agency's security policy.

Automation is at the heart of the Dynatrace platform, which frees up teams to focus their time and energy on the agency's mission rather than troubleshooting applications and underlying infrastructure. Dynatrace is a solution that is self-service in nature and doesn't require a degree in computer science to use and an army of engineers to maintain. Dynatrace is designed to be used by administrative teams, operations teams, and development teams alike, all working off the proverbial single sheet of music, seamlessly sharing data between different departments and cost centers.

Following are just a few of the many Dynatrace customer success stories in U.S. federal space:



Department of Homeland Security (DHS) Headquarters (HQ) recognized the need for application performance monitoring that spanned all DHS components across the OneNet network. They required visibility into the true root cause of application performance degradation and

the assurance that mission-critical applications were meeting requirements.



Army & Air Force Exchange Services (the military base Exchange stores) use Dynatrace on their eCommerce platform to break down silos between application and development teams and give management enterprise-wide visibility

through a single pane of glass. Dynatrace's automated root cause detection has empowered teams to take ownership of their code, reduced troubleshooting time, and enabled recovery of technical debt.



Social Security Administration uses Dynatrace's Software Intelligence platform to manage both legacy applications and new applications that leverage the latest Amazon Web Services (AWS) capabilities, microservices and containers, as they continue their Agile development journey. Disability Claims Processing System 2 (DCPS2) is one of many applications benefitting from the application performance monitoring and management and AIOps realized with Dynatrace.



Department of Veterans Affairs (VA) deployed Dynatrace to provide IT teams with ability to assess performance of new and existing applications, down to the user's desktop. This helps to ensure that veterans receive the best possible digital experience when accessing their health care applications.



Air Force 45th Test Squadron uses the Dynatrace Software Intelligence platform in their testing of integrated information systems and their individual components. These systems supply real-time data needed to plan and execute missions making application performance a critical function. With millions of lines of code and highly complex applications, Dynatrace can detect, solve and optimize automatically, ensuring mission success.



National Geospatial Intelligence Agency (NGA) leverages Dynatrace to ensure the availability of their most mission-critical imagery datasets. NGA depends on Dynatrace daily to meet their mission-critical goals.

Dynatrace provides an AI-driven observability platform that simplifies federal cloud complexity and accelerates digital transformation. With AI and complete automation, the all-in-one Dynatrace Software Intelligence Platform provides answers, not just data, about the performance of applications, the underlying infrastructure, and the experience of all users. That's why many U.S. Federal Government organizations trust Dynatrace to modernize and automate enterprise cloud operations, release better software faster, and deliver unrivaled digital experiences for the end users.

Acronym Key

ADDP	Application Debugging and Distributed Analysis
AI	Artificial Intelligence
AIOps	Artificial Intelligence for IT Operations
APM	Application Performance Monitoring
AWS	Amazon Web Services
BI	Business Intelligence
CI/CD	Continuous Integration / Continuous Delivery
CISA	Cybersecurity and Infrastructure Agency
DEM	Digital Experience Monitoring
DHS	Department of Homeland Security
DOD	Department of Defense
DOE	Department of Energy
EO	Executive Order
FedRAMP	Federal Risk and Authorization Management Program
HQ	Headquarters
IaC	Infrastructure as Code
IC	Intelligence Community
ICS	Industrial Control Systems
IoT	Internet of Things
IT	Information Technology
KPI	Key Performance Indicator
MFA	Multi Factor Authentication
ML	Machine Learning
MTTR	Mean Time to Recovery

NDR	Network Detection and Response
NIST	National Institute of Standards and Technology
NSM	National Security Memorandum
OMB	Office of Management and Budget (The White House)
OOB	Out of Band
QoE	Quality of Experience
RASP	Runtime Application Self-Protection
SAST	Static Application Security Testing
SBOM	Software Bill of Materials
SLA	Service level agreement
SOAR	Security Orchestration, Automation and Response
SSA	Social Security Administration
VA	Veterans Administration
WAF	Web Application Firewall

Notes

Quickly become conversational about managing application complexity in Federal Government

As the U.S. Public Sector continues to negotiate the challenges of “cloud-first” mandates, it is seeing the applications it uses grow in scale and complexity. As a result, federal agencies have been forced to re-think how they manage the applications they use. This ebook will help you understand how you can simplify operations and standardize processes, while saving time and money, and to understand the key role software intelligence plays in maintaining the health of your critical applications.



About Andrey Zhuk

Andrey Zhuk is field CTO for Cybersecurity at CTG Federal, where he helps US Government Agencies adopt new cloud services and secure agency assets in the cloud. Andrey is an experienced cloud, cyber and network architect with over 13 years of experience in US Federal Government space.



ConversationalGeek®

For more books on topics geeks love visit

conversationalgeek.com