

The Definitive Guide to AIOps

SPONSORED BY...



technologies

A Broadcom Company

Copyright © 2018–2019

All rights reserved.

Second revision, May 2019

Introduction:

The Definitive Guide to AIOps

As hardware and software systems become more efficient, sophisticated and useful, they also tend to grow more complex. For example, when virtual machines replaced bare-metal software environments, virtualization created a new layer of complexity that IT teams had to plan for and manage. The shift in recent years toward microservices and containers similarly increased the number of components that go into a single application, as well as the challenge of orchestrating all of them.

Traditionally, the ability of IT Ops teams to handle ever-increasing complexity has been limited. Hiring more staff is the most obvious response, but that is not a cost-effective solution, or one that can scale well.

Automation tools can also help handle added complexity. However, because traditional automation tools require humans to configure, deploy and manage them, the ability to simplify increasingly complex IT environments is also limited.

AIOps as the Answer to Complexity

In recent years, Artificial Intelligence for IT Operations (AIOps) has emerged as a better solution to the challenge of ever-increasing complexity in IT. AIOps leverages Big Data, data analytics and machine learning to provide insight and enable a higher level of automation (one that does not depend extensively on human operators) for the management tasks that modern infrastructure and software require.

For this reason, AIOps holds tremendous value. Going forward, AIOps will play a key role in enabling new efficiencies for IT teams. It will also make practical the adoption of complex next-generation technologies that cannot be managed successfully using traditional solutions.

In short, businesses of the future won't survive without the assistance of AIOps. If your business has not yet begun adopting AIOps-powered solutions, now is the time for assessing, planning and implementing AIOps tools that can drive business value.

This guide is designed to help you in making the migration toward AIOps. It defines AIOps and assesses the current state of AIOps within the IT industry. It also identifies and explains the core components that drive AIOps, as well as the main use cases for AIOps-powered tools.

Chapter 1:

What Is AIOps?

Defining AIOps

AIOps is the use of machine learning, Big Data and automated decision-making to complete IT tasks. AIOps makes it possible to automate processes that would traditionally require significant manual intervention by humans.

AIOps, which is short for “algorithmic IT operations” or “artificial intelligence for IT operations,” entered the IT lexicon in 2016, when Gartner coined the term as part of an effort to understand how data analytics was enabling new efficiencies for IT Ops teams.

Why is AIOps Innovative?

The use of data analytics and machine learning by businesses has been widespread for years—It did not arise alongside AIOps. IT operations, or IT Ops, also existed as a distinct discipline long before the concept of AIOps appeared.

However, what makes AIOps innovative is that it brings data-driven insights and IT Ops together. Previously, data analytics was used primarily to drive business insights, not to help IT

teams do their jobs. To the extent that data and machine learning played a role in IT Ops, they were limited mostly to basic security and infrastructure monitoring tools. IT Ops teams made use of automated tools to help make their work more efficient, but those tools were not typically capable of making complex automated decisions based on data, and they required significant manual effort to use.

AIOps changes this by providing IT Ops teams with access to tools that can make advanced decisions and perform automated actions by collecting and analyzing data. It represents a much more refined, sophisticated way of integrating data analytics into IT Ops. In addition, it helps traditional IT Ops admins transition into Site Reliability Engineer (SRE) roles and support more scalable workflows that align with business needs.

The State of AIOps

While precise data about current AIOps adoption rates is not available, Gartner [projected](#) in 2017 that “25% of global enterprises will have strategically implemented an AIOps platform supporting two or more major IT operations functions” by 2019. In addition, recent research by TechValidate found that 97% of surveyed IT organizations agreed that AIOps-enabled solutions that deliver actionable insights will help automate and enhance overall IT Operations functions.

AIOps is already seeing early adoption by enterprises, although it will likely take some time before a majority of businesses have deployed AIOps platforms.

The major hurdles currently standing in the way of greater AIOps adoption include a lack of certainty among businesses over whether AIOps reflects true innovation or mere hype. That doubt will likely disappear as more enterprises adopt AIOps and the

value of AIOps becomes clearer. In a 2018 NewVantage Partners [survey](#), 97.2% of executives are investing in building or launching Big Data and AI initiatives.

Low confidence among businesses in their ability to collect high-quality data (combined with uncertainty over how best to implement AIOps-enabled solutions that will deliver broad, long-term value) are also holding back AIOps adoption in some organizations. These challenges, however, can be overcome with sufficient research and planning.

AIOps Components

Such planning begins by identifying the core components that make AIOps possible and assessing your business's ability to implement them effectively.

The chief components of AIOps include:

- *Data collection.* Collecting data is the first step in enabling AIOps. Successful AIOps and Big Data technologies are used to collect data from disparate sources, to transform and aggregate the data as needed, and backup and retain data effectively and maintain data quality sufficient for powering data analytics and machine learning.
- *Data analytics.* Once data has been appropriately collected and transformed, statistical analytics are performed to draw out insights from the data.
- *Machine learning.* Machine learning is the process of using the insights gleaned from data analytics to make automated decisions. Machine learning is implemented through algorithms that allow software to react automatically to information revealed by data.

- *Artificial intelligence (AI)*. AI refers to the broader category of automated decision-making, of which machine learning is one component.

We will explore each of these components below within the context of discussing AIOps use cases and practices.

AIOps Use Cases

By combining data collection, data analytics and machine learning to form a complete AIOps solution, IT Ops teams can support several key use cases:

- *Anomaly detection*. Perhaps the most basic use case for AIOps is detecting anomalies within data, then reacting to them as needed.
- *Causal analysis*. AIOps also helps IT Ops teams automate root cause analysis so that issues can be resolved quickly.
- *Prediction*. AIOps allows tools to make automated predictions about the future, such as how user traffic is likely to change at a given point in time, then react accordingly.
- *Alarm management*. AIOps plays an increasingly important role in helping IT Ops teams to contend with the deluge of alerts that they must handle in order to support operations.
- *Intelligent remediation*. AIOps drives closed loop remediation through automation tools without relying on human operators.

The following chapters dive deeper into each of these cases by explaining in detail what they involve and how to support them successfully.

Chapter 2:

Data Collection and Normalization

Data forms the foundation for AIOps. For this reason, implementing effective processes for collecting and normalizing data is an essential first step in creating an AIOps-enabled solution.

Data collection refers to the task of moving data from the sources where data originates, which are usually of a diverse nature, to a location where it can be processed and analyzed.

Data normalization is the process of preparing data for analysis. Normalization involves converting data from one format to another so that it is compatible with data analytics tools. It also often entails integrating diverse datasets so that they can be analyzed efficiently from a single location.

To perform data collection and normalization effectively, organizations should keep the following challenges and best practices in mind.

Disparate and Diverse Data Sources

In most cases, the data that powers AIOps platforms is not “born” in a single location or a single format. It instead originates in many different formats and is spread across multiple locations.

For example, some of the data that your business collects for AIOps might originate from web server logs that are available in plain text. At the same time, you might collect other data from operating system logs that are stored as compressed files and need to be unpacked before they can be analyzed. Similar challenges arise when collecting data from different types of databases. For instance, data inside MySQL databases is typically formatted differently from data within so-called NoSQL databases. Time-series datasets also pose data collection challenges because they require data that was collected at different times to be normalized before it can be analyzed.

The diverse nature of data sources and formats creates two distinct challenges:

- **Organizations must be able to collect and aggregate data from multiple locations.** The amount of effort and complexity required to perform this task will vary depending on how many data sources you have and how widely distributed they are. In most cases, data collection will require running agents on your various systems that can collect the data they generate and send it to a central location for storage and processing.
- **Normalizing data by translating it into formats that are compatible with analytics tools.** Normalization doesn't necessarily require transforming all of your data into a single

format. It does, however, typically involve performing at least some dataset transformations, as well as taking advantage of Big Data tools like Apache Hive, HBase and Elasticsearch, which provide an interface for integrating data inside conventional databases with Big Data analytics tools such as Hadoop.

Real-Time Data Operations

When planning and implementing an AIOps solution, it is important to strive for real-time data collection and normalization. Real-time data operations mean that you can collect and analyze data as quickly as it is generated and gain instant or near-instant insights as a result.

Real-time data processing is essential for most AIOps use cases. If your goal is to use AIOps to detect anomalies that could indicate a security breach, for example, being able to gain that insight as soon as the breach occurs will drive a much greater deal of business value than discovering the problem after attackers are already exploiting your data and infrastructure. Similarly, when you use AIOps for root-cause analysis of a software or infrastructure problem, you want to be able to get to the root of the issue as quickly as possible so that you can resolve it before it impacts end users. In both of these examples, delays of even just a few minutes in collecting and normalizing the data that you depend on for your AIOps processes could undercut your ability to achieve your business goals.

Achieving data collection and normalization in real time requires full automation of these processes. The agents that help you collect data, and the tools that help you transform or access it for analytics purposes, must be able to operate without the assistance of humans. Otherwise, if you rely on admins to help

collect data or perform data transformations manually, you won't be able to achieve real-time insights.

Data Retention and Backup

Although real-time data processing is an important part of AIOps, keeping data available after AIOps processes are complete is also valuable. You may be required to retain data for a certain period for compliance reasons, and even if you are not, being able to perform retrospective analysis of data can be useful.

This is why your AIOps planning should include assessment of how long your business will retain data after it has been collected and normalized, as well as how data will be backed up in order to protect it against unexpected disruptions. While data retention and backup policies vary widely depending on business needs, a common means of deciding which policies are the best fit for your organization involves analyzing two factors:

- *Recovery Point Objective, or RPO.* RPO is the amount of data that your business can afford to lose permanently without serious consequences. If you have high RPO needs, it is essential to back up data on a constant, routine basis.
- *Recovery Time Objective, or RTO.* RTO refers to the amount of time that your business can wait for data to be made available again following a disruption. High RTO requirements necessitate backup processes that allow you to restore data very quickly, as well as period tests of recovery time to ensure that you are able to meet recovery goals.

In order to reduce data storage and backup costs, businesses can take advantage of discounted data storage services available from public cloud providers. These services, which are typically referred to as “cold storage,” provide low-cost data storage, with

the caveat that accessing the data often entails a delay. For data that is no longer in active use for AIOps, that delay is typically acceptable.

Openness

A final key factor to consider when preparing data collection and normalization solutions for AIOps is the issue of closed versus open source-based solutions.

In general, choosing open solutions is better than adopting proprietary, closed-source tools. The latter can lead to lock-in and restrict your business's ability to modify its AIOps toolset and processes in the future.

For this reason, it is a best practice to adopt open source data collection and normalization tools. Examples include:

- Apache Kafka
- Apache Hive
- Apache HBase
- CloverETL
- KETL
- Rsyslog
- Logstash
- Elasticsearch

Keep in mind that many proprietary data collection and normalization tools are built on top of these open source solutions. Some such platforms are more “open” and compatible with third-party tools than others. If you consider commercial tools for data collection and normalization, assess how suitable they are for integration with third-party options in order to avoid lock-in regrets. Similarly, when considering an open source

solution, be sure to assess whether the effort required to set up and maintain the tool offsets the cost savings of using open source.

Chapter 3:

Detection

Detecting anomalies in order to locate problems and understand trends within infrastructure and applications is a key use case for AIOps. Detection allows tools to both recognize behavior that is out of the ordinary (such as a server that is responding more slowly than usual, or uncommon network activity generated by a breach) and react accordingly.

What is an Anomaly?

An anomaly is a data point or event that is consistent with normal operating conditions. In other words, it is an outlier.

Dynamic Baselineing

While understanding the concept of an anomaly is easy enough, what makes anomaly detection particularly challenging for AIOps in modern software environments is that, in many cases, there is no consistent means of defining “normal” operating conditions. The amount of network traffic, memory and storage space that a given environment consumes might fluctuate widely throughout the day, for example. So could the number of active users or application instances.

Effective detection under these circumstances requires AIOps tools that are intelligent enough to set dynamic baselines.

Dynamic baselines allow the tools to determine what constitutes normal activity under given circumstances (such as the time of day and the number of registered users for an application), then detect data or events that do not align with the dynamic baseline.

Univariate vs. Multivariate Anomalies

Alongside dynamic baselining, another important factor to bear in mind for AIOps detection use cases is the difference between univariate and multivariate anomalies.

Univariate anomaly detection focuses on identifying outliers based on single metrics or data points. For example, a univariate anomaly might generate an alert when disk storage space surpasses a normal threshold.

Multivariate anomalies, in contrast, detect outliers based on a series of different metrics. In a basic multivariate detection scenario, an AIOps-enabled tool might analyze disk usage, memory usage and network traffic at the same time and assess whether overall behavior is out of the ordinary. In this case, excessive disk usage alone may not trigger an alert, but if high disk usage coincides with unusual memory consumption and network traffic, the tool would be likely to determine that an anomaly exists.

This is a simple multivariate detection example. In more complex situations, multivariate methods rely on neural networks to model interactions between various metrics and make decisions based on them.

Multivariate anomaly detection thus provides deeper, more comprehensive insight. However, multivariate detection is also more difficult to implement effectively because it requires identifying multiple metrics that can be analyzed in common and creating algorithms that can interpret them accurately. Multivariate detection is also more difficult to scale because complexity grows as more metrics are introduced.

In most cases, an AIOps solution that combines univariate and multivariate detection methods will deliver the best results. Univariate detection can be useful for basic alerting and monitoring, while multivariate methods can power more complex automated decision-making.

Another critical capability of an AIOps solution is to hide the algorithmic complexity. It should automatically pick the right algorithm based on the type of data being analyzed.

Detection Model Extensibility

The detection processes that power AIOps should be extensible and future-proof, rather than designed only to meet a finite set of needs.

Extensible detection strategies are characterized by the following features:

- The ability to add new metrics, or modify the weights afforded to various metrics within detection models.
- The ability to add new data sources and technologies into detection models.
- Support for continuing to adapt dynamic baselining techniques as behavior grows more nuanced and complex.

What this means in practice is that detection models often start small, but grow in scale and complexity over time. At first, your AIOps strategy may be driven primarily by univariate detection models, coupled with some basic multivariate methods. And they may focus on simple metrics such as memory and disk usage. Over time, however, you will likely want to make your techniques more complex by adopting more sophisticated multivariate models that collect metrics from advanced technologies—such as the startup time of containers or the execution time of serverless functions.

Chapter 4:

Causal Analysis

Another key use case for AIOps is causal analysis. This refers to the task of tracing a problem to its source or sources in order to help resolve it.

The Challenge of Causal Analysis

AIOps-driven causal analysis is increasingly important as software environments grow more complex, and the dependencies between different components become increasingly difficult to map on the surface level.

Consider, for example, a web application that consists of a frontend component as well as a backend database, and that is deployed as a set of microservices hosted in containers. In the event that the IT Ops team notices that the web server has started responding slowly, tracing the problem to its root cause could be quite difficult without the assistance of automated, data-based tools. The issue could be caused by network bottlenecks. It could be the result of failing disks, or a database configuration problem. The container orchestrator could be failing to balance application load properly across multiple container instances. The application code itself might be the source of the problem.

Rather than investigating each potential cause of the problem manually, an IT Ops team could deploy AIOps tools that automatically analyze data in order to determine the likely cause or causes of the issue. By parsing information such as network traffic patterns, container statistics, application profilers and database logs, an AIOps tool could provide quick visibility into the issue. AIOps tools also provide end-to-end visibility, enabling IT Ops teams to identify problems that they might not recognize on their own.

Causal Analysis Data Collection and Contextualization

Your causal analysis efforts are only as effective as the data you collect. You must determine which types of causal analysis you intend to perform, then ensure that you are collecting and normalizing the right data to support those analyses.

It is also important to collect contextual information, or data that is not directly related to the problem whose causes you are analyzing. This includes information such as how often similar problems have occurred in the past and what their causes were, or whether other systems are experiencing similar issues. Contextual information such as this can help you interpret the scope and significance of a problem, and prioritize it accordingly.

Handling Multiple Causes

It is sometimes the case that there are multiple causes of an issue. For instance, in the web application example above, it is possible that network bandwidth limitations and disk I/O problems are both causing a slow application response. Your causal analysis strategy and reaction should therefore be designed to handle situations in which multiple causes must be addressed in order to resolve a problem.

Causes can come in multiple layers, too. To go back again to the web application example, slow response times could be caused by improper load balancing, which is in turn caused by a lack of memory resources for the container orchestrator. In this case, resolving the first cause (the load balancing problem) won't solve the underlying issue.

In situations like these, where multiple causes are at play, graphical modeling can be helpful for separating intermediate causes from root causes when resolving a problem.

Drilling Down

In addition to helping you identify the cause of a problem, AIOps tools should provide the ability to drill down into a problem in order to investigate it at a deep level. For example, if a web application is failing and you determine that the cause is network bandwidth limitations, you might want to be able to drill down and determine whether a certain type of network traffic— such as traffic from a specific region—was associated with the bottleneck that caused your application problem.

Insight such as this can help your IT Ops team improve systems so that they are more resilient to the recurrence of problems. In this way, causal analysis with the assistance of AIOps not only helps to resolve problems in real time, but also helps to achieve continuous improvement by preventing problems from happening again.

Chapter 5:

Prediction and Trend Identification

AIOps can also facilitate continuous improvement by helping IT Ops teams to predict future developments and identify trends.

To understand the value of prediction and trend identification, consider the following examples.

Predictive Capacity Analytics

Right-sizing infrastructure is a constant challenge for most IT Ops teams. If an organization fails to provide enough compute, storage and other resources to its applications, it risks performance problems. On the other hand, providing excessive resources leads to cost-inefficiency, because the organization pays to set up and maintain more infrastructure than it needs.

Predictive capacity analytics do much to address this challenge by helping IT Ops teams to predict how their infrastructure needs will grow over time. They could even enable cyclical resource allocation adjustments. For example, if an online retailer experiences significant traffic peaks on certain days of

the year, predictive capacity analytics could enable the retailer to allocate extra infrastructure resources on those days, while scaling back on other days in order to save money.

Application Performance

Alongside software testing, AIOps can play a role in helping to optimize application performance prior to application deployment. For example, prediction could help the IT Ops team determine how an application will respond to a certain condition, such as a sudden increase in network traffic that results from a DDoS attack. Obtaining such insights before the event occurs in production positions the IT Ops team to prepare for it more effectively.

IT Ops Performance

Prediction and trend identification can also help the IT Ops team itself to optimize its own performance. How has the team's ability to resolve incidents within a certain window of time changed? Which types of problems are causing the greatest numbers of issues? By analyzing data to identify trends, AIOps can answer questions such as these, so that the IT Ops team knows where to concentrate its efforts going forward.

Chapter 6:

Intelligent Remediation and Automation

AI Ops not only helps IT Ops teams to identify problems and areas of concern, it also enables quick resolution of issues once they have been identified.

The Need for Fast Resolution

The importance of quick resolution of problems is easy enough to understand. In a world where more than half of users will [abandon a site](#) that takes longer than three seconds to load, and in which a one-second delay in page load time [results](#) in 7 percent fewer sales, organizations cannot afford to leave software availability or performance problems unaddressed for any length of time.

While collecting and analyzing data in real time in order to find issues quickly is one component in achieving fast resolution, so is the ability to interpret the problem quickly using AIOps-driven insights. As noted above, AIOps-enabled solutions can help engineers trace the cause of an issue and suggest remediation approaches so that problems can be resolved as quickly as they appear.

Leveraging Historical Data for Remediation

AIOps can also enable faster incident resolution by helping IT Ops teams to interpret historical data associated with past issues in order to suggest solutions for similar incidents as they occur. Without AIOps, parsing through reams of logs and other data in order to identify the similarities between two incidents, and determine whether the resolution that worked for the first will also effectively address the second is not feasible. AIOps-enabled solutions, however, can provide rapid insight based on historical data to help respond to this challenge.

Automated Resolution

AIOps tools can even take automatic action to resolve problems after they have identified them. They could block a host or close a port automatically in response to a security threat, for example, or spin up additional instances of an application if they determine that the existing instances are insufficient to meet demand.

Automated resolution is not practical in all situations; sometimes, the ultimate resolution to an incident will have to be implemented manually, even though AIOps can provide insights that help lead to the resolution. Yet as machine-learning algorithms grow increasingly sophisticated, the problems that AIOps tools can resolve automatically will increase in number, enabling even faster and more seamless incident resolution.

Chapter 7:

Breaking Through the Noise: Managing Alerts

While AIOps enables broader, faster visibility into infrastructure and software than IT Ops teams can achieve using manual tools, it also creates a risk of information overload. If AIOps tools are deployed or managed improperly, they generate so many alerts that IT Ops engineers become overwhelmed and begin ignoring notifications. This issue, commonly described as alert fatigue, undercuts the value of AIOps-based monitoring and analysis.

Avoiding alert fatigue and managing alerts successfully requires several best practices:

- *Avoiding manual alerting thresholds.* Alerts that are configured manually to fire based on fixed thresholds do not work well in today's dynamic environments. Not only do manual alerts require considerable time to configure, but they can also lead to false positives because what constitutes acceptable disk, network or other resource consumption at one moment may change in the next moment, along with the environment. Instead of configuring manual alerting thresholds, AIOps

tools can set thresholds automatically. They can also leverage dynamic baselining (discussed above) to configure when an alert should fire.

- *Actionable alerting.* Rather than merely indicating that a problem has occurred, alerts should be accompanied by information that will help IT Ops teams to respond to the problem. This means providing contextual information that will help engineers to understand a problem more thoroughly. Actionable alerting can also include data-based resolution recommendations for engineers to consider.
- *Avoiding redundant alerts.* It is often the case that a single root-cause problem triggers multiple alerts. A database failure could impact multiple applications and cause an alert for each one, for example. In this case, multiple alerts will distract the IT Ops team rather than help it to resolve the incident quickly. Instead of generating redundant alerts, AIOps tools should intelligently map and/or cluster multiple issues to a single root cause, and generate one alert that will help engineers resolve that cause quickly.

Chapter 8:

The Future of AIOps and Data Analytics

Beyond increasing enterprise adoption of AIOps-enabled solutions, what does the future hold in store for AIOps? The following trends are likely to form an important part of the field of AIOps as it continues to evolve in coming years.

Support for Increasingly Dynamic Environments

As noted above, part of the appeal of AIOps is that it enables IT Ops teams to handle highly dynamic infrastructure and software environments, such as IoT devices, containers, and serverless platforms more effectively.

Going forward, it is likely that newer technologies will appear that introduce even more dynamism to the deployment models that IT Ops teams have to support. While the exact nature of these technologies remains to be seen, it is a safe bet that AIOps will be a key enabler for managing them.

Graphical Pattern Recognition

We have noted above how graphical modeling can assist in efforts to understand particularly complex causal relationships.

While some AIOps tools are already making use of graphical pattern recognition to a certain extent, expect the role of graphical modeling to increase in importance in the future. Graphical models will help AIOps tools to deliver new levels of insight as the data they process grows in volume and complexity.

Genetic Algorithms

The use of genetic algorithms within AIOps applications is likely to increase significantly as AIOps evolves. Genetic algorithms refer to software logic that improves over time by using data and machine learning to refine itself automatically. For example, as an AIOps algorithm deals with a certain type of problem repeatedly, it will learn automatically which solutions work best, and train itself to pursue those in the future.

Genetic algorithms form an important part of the ability of AIOps tools to achieve continuous improvement. They not only enable faster resolution of problems with less manual effort on the part of IT Ops engineers, but also make it possible for AIOps tools themselves to become faster, more accurate and more effective over time, without having to be updated manually.

CA's AIOps-Enabled Solutions

CA Technologies has been at the forefront of the AIOps ecosystem since its introduction several years ago, and will continue to lead the way as AIOps evolves.

AIOps functionality is delivered by CA through its [Digital Experience Insights](#) platform and comprises an essential part of the features of CA's [APM](#) and [API monitoring](#) solutions. CA also recently introduced [Digital Operational Intelligence](#), an AIOps-enabled solution that provides cross-domain contextual intelligence to help IT Ops teams to make smarter, faster decisions for enhancing user experience and improving IT service quality and capacity. Built on an open, powerful engine, it provides users with comprehensive insights by ingesting and analyzing a diverse dataset including metrics, topology, text, and log data. The machine learning-driven analytics, along with out-of-the-box visualization and correlation, help drive a superior user experience and deliver significant operational efficiencies.

[Learn more](#) about how CA can guide you in your AIOps journey.

About CA Technologies

CA Technologies helps customers succeed in a future where every business—from apparel to energy—is being rewritten by software. From planning to development to management to security, at CA we create software that fuels transformation for companies in the application economy. With CA software at the center of their IT strategy, organizations can leverage the technology that changes the way we live—from the data center to the mobile device. Our software and solutions help our customers thrive in the new application economy by delivering the means to deploy monitor and secure their applications and Infrastructure.



About Sweetcode.io

Sweetcode.io is a site owned and managed by Fixate IO. Its purpose is simple: To give techies a place to share what they know and impact the market with high-value, practitioner-generated technical content. Sweetcode is committed to publishing tactical content that supports the growth of seasoned developers, while also serving as a platform for those who are just starting their careers. All of the content on the site is practitioner-created, and Sweetcode invests in sourcing content from under-represented coders to give them a place to share what they know.





Sweetcode.io is a site owned and managed by Fixate IO. Its purpose is simple: To give techies a place to share what they know and impact the market with high-value, practitioner-generated technical content. Sweetcode is committed to publishing tactical content that supports the growth of seasoned developers, while also serving as a platform for those who are just starting their careers. All of the content on the site is practitioner-created, and Sweetcode invests in sourcing content from under-represented coders to give them a place to share what they know.

sweetcode.io

Email: hello@sweetcode.io