dynatrace

Data security and privacy in the cloud: A proactive framework for observability

What's inside

INTRODUCTION

Cloud done right

Successful digital transformation requires every application and digital service, and the dynamic multi-cloud platforms they run on, to work perfectly. We call this Cloud Done Right.

But these dynamic, highly distributed cloud-native technologies are fundamentally different than their predecessors. The resulting complexity brought on by microservices, containers, and software-defined cloud infrastructure is overwhelming at web scale. It's all beyond the limits of human teams to manage and scale on their own.

To understand everything going on in these ever-changing environments, all of the time, observability needs to scale.

CHAPTER 1

Data security and privacy at the forefront

Observability platforms collect and analyze significant amounts of data to give their customers the precise answers they need to simplify cloud operations and deliver flawless and secure digital experiences. Because such data may include personal data and confidential information, organizations are becoming more vigilant about the measures these platforms are taking to control access and the use of the data the platforms process. As a result, data protection has become a key requirement from customers for any SaaS observability and security platform. Every organization needs to carefully consider how to minimize the risk of unauthorized access and use of not just their own data, but data of their customers and users as well.

Both data security and data privacy need to be implemented as a primary mindset, and integrated into every step of developing, using, and supporting a SaaS platform. Protecting data as a priority:

- Provides customers with control and transparency over the management and use of their data.
- Enables security in depth with layered security controls.
- Enables customers to comply with their internal policies and legal requirements.

CHAPTER 2

Be in control of your data

When you want to maximize the value from observability and security platforms, it's important to enable them to process data at scale. However, you'll need to carefully control the processing of personal and other confidential data. To achieve that, you need to be able to:

- Easily configure data that need to be excluded from analysis
- Use strong default settings out of the box

Being in control of your data happens on three different layers.

You control the data sent to the secure SaaS cluster

The first step is defining what data can't leave your environment to ensure that you comply with your internal policies and external requirements. Masking at capture provides you with ease of configuration to exclude confidential data before it leaves your environment, and includes strong privacy by default to automatically exclude such data, including for example:

- Credit card numbers
- IBANs (bank account numbers)
- End user IP addresses

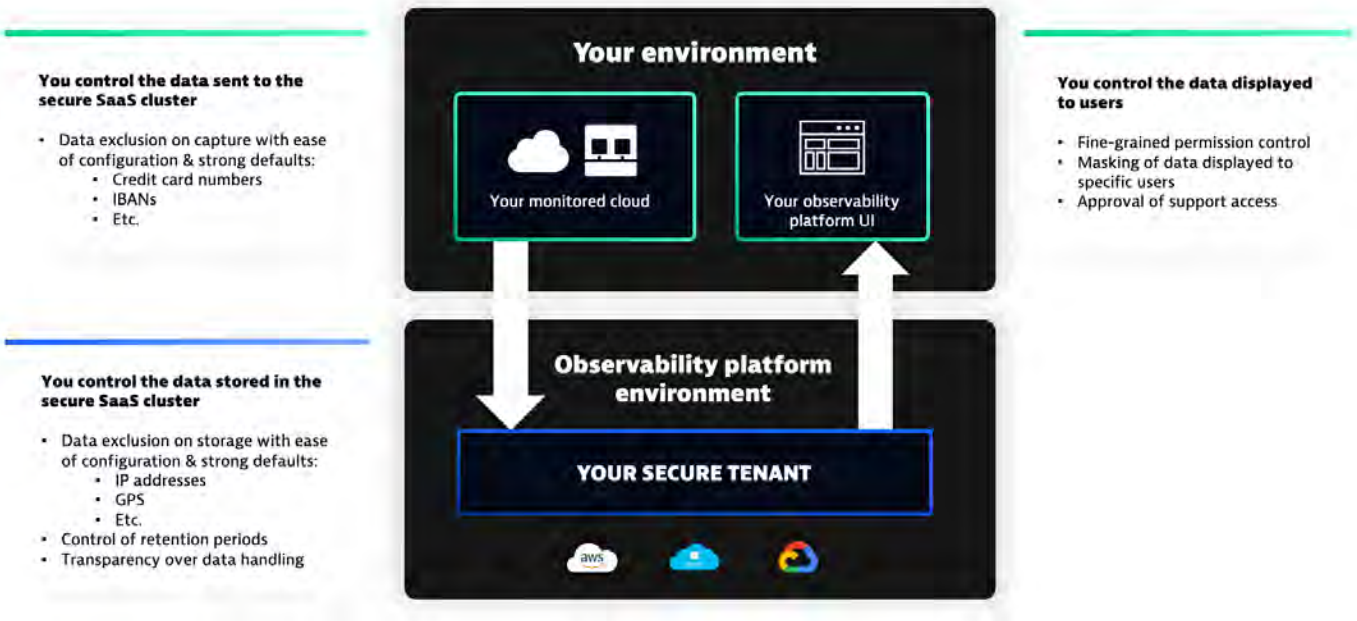

You control the data stored in the secure SaaS cluster

In order to get the full value from the collected data, it needs to be sent to a secure SaaS cluster. That way, the data can be analyzed, and value is generated from it for you. However, that doesn't mean that all original data has to be stored as well. With data masking at storage, the data is analyzed when it reaches your secure SaaS cluster, and masked before it's been stored. This way, such confidential data is effectively deleted and can't be reverse engineered.

Such data, that provide you with a lot of value but shouldn't be stored, might include:

- End user IP addresses
- GPS
- URLs.

In addition to that, it's important that you control the retention periods of your data, and decide how long they'll be stored in the secure SaaS cluster. This way you can fine-tune retention periods for your data to fulfil your security, privacy and legal requirements.

You control who can see the data

The last area where you need control of your data is the user interface. It's important that data is displayed only to the users who are authorized to see it, and no one else. You can easily achieve this through being able to configure roles and permissions to your user interface and data displayed.

In addition to that, you should always be able to mask certain data displayed to specific users who have access to various dashboards, to be able to provide them with more value while excluding the data that they aren't authorized to see.

Finally, you should be able to fully control external access, such as the access by Dynatrace support engineers, to your user interface and data displayed, as well as revoking such permissions at any time.

CHAPTER 3

End-to-end security

Being in control of your data is vitally important, but that alone isn't enough. You need to make sure that all your data, as well as the environments, are secure, and that it's safe to send and analyze your data. End-to-end security means that all components of the infrastructure need to be secured from unauthorized access from:

- the agents monitoring your environment
- clusters analyzing and storing the data
- the user interface providing you with the dashboard

Don't forget about the strong encryption of the data in transit between those environments too.

Secure by default

Many of the agents and technologies monitoring your environment need root mode privileges to be able to do their job properly. It might be easier, but root mode increases the risk of unauthorized access to your environment, in cases where the component has been compromised. Because of that, it's critical that such components and agents monitoring your environments adhere to the principle of least privilege (PoLP), ideally operating fully in non-root mode.

Even when operating in non-root mode, it's still important to ensure that the components aren't manipulated and potentially compromised in transit before they reach your environment. This can be easily achieved by automatically verifying their signatures, and preventing any actions in case the signature can't be verified.

Secure connection between the components monitoring your environment and clusters analyzing the data needs to be established. To minimize the risk of unauthorized access to the authentication tokens, it's good practice to include them directly in those components. Since they've been verified against any external manipulation, not having to copy, paste and store such authentication tokens externally reduces such risks significantly.

In addition to having been able to control which data is displayed to your users, access to this data needs to be secured.

Secure software development lifecycle (SDLC)

Secure software that has access to your data is the foundational layer of end-to-end security. Independent penetration testing validates that the software doesn't have any critical vulnerabilities. It also gives the software developer the opportunity to address vulnerabilities before they could compromise your production environment.

Independent audits, such as ISO 27001, validate that software developers have implemented the highest standard of security controls and processes to prevent introduction of potential vulnerabilities into their software.

CHAPTER 4

Data security

Data security safeguards data to protect it from unauthorized access, maintain its accuracy as well as to ensure the correct use of information and its availability.

Secure SaaS cluster

SaaS cluster security builds upon the already high security that the hyperscalers (usually AWS, Azure and GCP) provide. However, such high security should be extended by additional controls in order to provide the best-in-class security, particularly in scenarios such as:

- Protecting from attacks and blocking them before they even reach the environment
- Automatically identifying and remediating common vulnerabilities and exposures (CVE) in production

Data separation
To prevent any unwanted mix-ups, your data needs to be safely separated from the data of your software provider's other customers in transit, during analysis and at rest.

Secret and token management
To avoid copying and pasting sensitive information, such as usernames and passwords, they'll need to be stored in (and ideally connected to) your secret vaults.

Encryption
Your data needs to be encrypted with a unique key to ensure that only you have access to your data.

Transparency
Audit logs and full transparency of what's happening with your data are necessary, so that you can see how it's been used and who has accessed it.

API access management
Effective and secure management of access tokens to your environment ensure that integrations and automation scripts follow the least privilege access principle.

In the worst-case scenarios, such tokens might be leaked by your engineers during development. Automatic scans of public code repositories, such as Github, for potentially leaked access tokens can notify you about such findings.

CHAPTER 5

Secure platform apps

Apps that extend the observability platform and bring customized answer-driven automation to your specific use cases have the potential to add great value to the tools you already use. In such cases, it's important to ensure that custom-made apps are built securely, and won't introduce additional security risks to the whole system. Addressing the following four areas should give you confidence that these apps are safe to use:

1. **Vulnerability scanning.**
   Because the apps often live on your SaaS cluster and will have access to your environment and data, they also need to be scanned for vulnerabilities in the same way as the core platform.
2. **Tight permission controls.**
   Similar to agent permissions, you should apply the principle of least privilege to the apps to ensure they have access only to the data they require.
3. **Secure connection to your environment.**
   When the app connects to the services in your environment for the custom use cases it's been built for, using a secure channel ensures access only to the intended service.
4. **Strict browser security policies.**
   Apps may also have access to the information from the other apps in the user interface. This is why the user interface security should be extended, and apps need to have tight browser restrictions to be able to only access the data they're supposed to.

CONCLUSION

Data security and data privacy by design

Addressing data security and data privacy well is a crucial requirement for all SaaS vendors. Our approach to incorporate data security and data privacy into everything we do gives us the opportunity to look far beyond the minimum legal requirements when it comes to protecting your data.

Data security
Dynatrace provides you with the security controls to prevent unauthorized access to your data, maintain data accuracy and ensure the correct use of information and its availability.

Data privacy
Strong privacy by design and privacy by default enable you to configure Dynatrace to maximize value while complying with your legal requirements.

Compliance and certifications
Independent verification ensures that the data security and privacy controls implemented by Dynatrace follow best practices to address your compliance requirements.

Automatic and intelligent observability for hybrid multiclouds

We hope this ebook has inspired you to take the next step in your digital journey. Dynatrace is committed to providing enterprises the data and intelligence they need to be successful with their enterprise cloud and digital transformation initiatives, no matter how complex.

If you are ready to learn more, please visit www.dynatrace.com/platform for assets, resources, and a **free 15-day trial**

Learn more

Dynatrace (NYSE: DT) exists to make the world's software work perfectly. Our unified software intelligence platform combines broad and deep observability and continuous runtime application security with the most advanced AIOps to provide answers and intelligent automation from data at enormous scale. This enables innovators to modernize and automate cloud operations, deliver software faster and more securely, and ensure flawless digital experiences. That's why the world's largest organizations trust Dynatrace® to accelerate digital transformation.

Curious to see how you can simplify your cloud and maximize the impact of your digital teams? Let us show you. Sign up for a free 15-day Dynatrace trial.

dynatrace

dynatrace.com blog @dynatrace

04.21.23 BAE4983_WP_USLet_mt