GoodData

GoodData Cloud

Security White Paper October 2023

Table of contents

This whitepaper is organized into four sections, starting with “Executive Overview.” After the overview, “GoodData Cloud Security” explains the security measures that are applied to the GoodData Cloud platform and provides guidance on the security concepts and techniques that GoodData customers should use to ensure security in the cloud.

The “GoodData.CN Security” section provides a quick security overview of the self-hosted version of GoodData Cloud. The fourth section, “GoodData Security Overview,” is intended for our customers’ security, compliance, and risk personnel. This section provides an overview of GoodData’s information security management system — built on the ISO 27001 standard — and explains our compliance against the 14 domains of the standard.

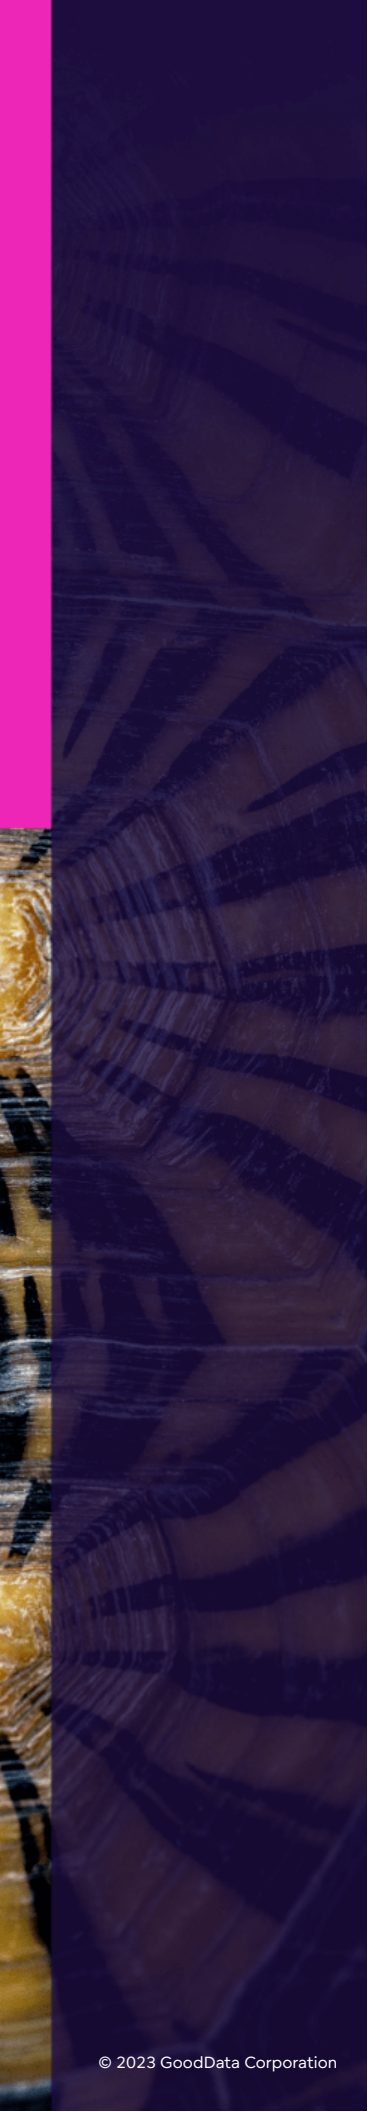

Executive Overview

The GoodData Cloud is a fully managed, API-first analytics platform. The platform helps enterprises as well as product and data teams securely build complex analytics solutions to deliver consistent insights to business users, customers, and partners at their point of work through a single source of metrics. GoodData Cloud offers developer-friendly features, such as declarative metadata and open API- and SDK-based integration, enabling the use of software development best practices in analytics development.

GoodData realizes that helping to protect our customers' data, mitigate any potential risks, and comply with relevant data protection laws, regulations, and standards is essential to building trust and delivering a high level of service. GoodData takes a risk-based approach to security, and this paper details the measures and technologies in place to protect our customers. It also outlines our internal security compliance standards to assure our customers about the diligence and robustness of our information security management system.

We adhere to the following certifications, frameworks, and best practices, demonstrating our commitment to data security and privacy:

- SOC 2® - SOC for service organizations: Trust Services Criteria for Security, Availability, and Confidentiality. Building on nine years of experience with the SOC report on our classic platform, we engaged with our long-term partner, EY, to expand the SOC 2 coverage to the GoodData Cloud offering as well.
- Compliance with the ISO 27001:2013 international standard for information security management systems and adherence to best practices documented in ISO 27002
- GDPR compliance
- California Consumer Privacy Act (CCPA) compliance
- Registered participant in the EU-U.S (incl. The UK extension) and Swiss-U.S. Data Privacy Frameworks
- HIPAA compliance

Defense in Depth

Like any other reputable SaaS platform provider, GoodData does not rely solely on the protection of its perimeter to safeguard customer data. Rather, it is a well-architected solution that considers every layer from the physical security measures at the data center to the access privileges that determine what data an individual user can access.

GoodData, as a best-in-class analytics provider, uses this approach to protect customer data. Leveraging our long-term experience in providing secure, reliable analytics platforms, GoodData Cloud has been built to align with up-to-date best practices for public cloud environments.

Regional Deployments

Data sovereignty is a complex issue that ranges from the technical arena to the regulatory arena and, at times, even the political arena. Understanding the complexities of this issue — and thanks to its standardized blueprint for a data center in a public cloud — GoodData is ready to offer additional data centers in AWS, Inc. worldwide according to the needs of its customers while maintaining our high standards for data security and privacy.

GoodData Cloud Security

Application Security

The GoodData Cloud not only provides customers with the ability to access their reports, dashboards, and data, but also enables direct integration with their other software so they can provide their clients with a seamless experience. The GoodData application employs many security measures to enable the secure flow of data from the customer's data sources to the end user's workspaces.

Integration and APIs

All integrations with the GoodData application programmatic interface (API) leverage HTTPS/TLS encryption. The user security model is enforced at the API level, ensuring that data retrieved through the API is still subject to user authentication and access privileges. For more information on this topic, please see the user security section below.

Connecting to Customer Data Sources

A data source is a logical object that represents the database where your source data is stored. To integrate your database into GoodData Cloud, you connect it to a workspace. The connection is protected via the same encryption as any other integration with the GoodData Cloud. To help securely connect your data source, you should create a separate set of credentials, whitelist them so that the database is accessible only from the secure environment of GoodData Cloud, and allow access to only the data that is required to be accessible from GoodData Cloud. Comprehensive guidance for setup is available in GoodData product documentation.

Physical and Logical Data Model

Logical Data Model (LDM) is an abstract view of your data in GoodData Cloud.

The LDM is connected to the underlying physical data model (which describes the tables of your database and represents how the actual data is organized and stored in the database), and serves as another layer of separation to mitigate potential security vulnerabilities such as SQL injections.

Multitenancy

GoodData Cloud allows GoodData customers to manage an environment with many tenants. In this multi-tenant environment, each particular tenant can access only the entities and data that they are enabled to access. Without appropriate permissions, a tenant cannot access those entities or data. Each customer has its own organization within a shared GoodData Cloud instance and is fully separated from the other customers on metadata level.

As an option, customers may choose a dedicated deployment hosting model, where a separate Kubernetes cluster is dedicated to each single customer. In this setup, additional technical security safeguards — such as establishing a private link between a dedicated GoodData Cloud cluster and the customer's VPC — are available. This option is suitable for customers with strict security requirements or those who expect that their solution will need to scale up/down dynamically.

Each object or data entity is tied to a specific organization, thus ensuring strong segregation between individual GoodData customers. Within the organization, the objects and entities can be further assigned to one or more tenants.

Workspaces

A workspace hierarchy in a multi-tenant environment defines how entities of a particular tenant (parent workspace) can be shared with other tenants (child workspaces) in read-only mode. The child workspaces use the parent workspace's LDM, analytical model, connected data sources, and so on. When the parent workspace receives a new entity, it becomes available to its child workspaces.

Child workspaces inherit entities from their parent workspace as well as that parent workspace's own parent workspaces — all the way up to the root workspace. The root workspace is the top-level workspace in the hierarchy which does not have a parent workspace. Customers may set up as many root workspaces as needed.

User Access

End users may access the data only through the application layer. Whether this access is through the user interfaces or through the publicly available API, it enforces user access controls to permit access to customer data only to authorized users and personnel.

For security purposes, GoodData does not provide end users with direct access to customer data sources or internal data caches. This approach prevents unauthorized services or systems from accidentally or maliciously retrieving or modifying customer data.

User Security

User security is enforced via a variety of security measures that allow authorized users to view only the strictly defined set of objects and data that are needed to perform their job.

Authentication

GoodData's architecture relies on a centralized authentication and authorization security framework to control access to services. We use OAuth standard so that our customers can use their own IAM to ensure seamless integration with their ecosystem.It also allows them to manage user access security, including the authentication mechanism, session expiration, etc., according to their company standards. OIDC tokens are stored in HTTP Cookies, and Authenticated Encryption with Associated Data (AEAD) with AES 256 encryption with Galois/Counter Mode is used. API tokens may be created for accessing GoodData Cloud from command line tools or for integration with other customer systems.

User Groups, Permissions, and Data Filters

By default, only the administrator starts off with the permissions necessary to view and modify the objects. To make your project accessible to other users, you must group users into appropriate user groups and assign permissions to these groups that are appropriate for their use case.

Permissions are organized into hierarchy and scopes. The scope is determined by an object type's unique features which can be further restricted or enabled via permissions.

In addition, data filters let you limit what data from a parent workspace is available to their child workspaces.

GoodData.CN Security

GoodData.CN is a self-hosted version of GoodData Cloud. It is purpose-built to scale with microservices and customers can deploy it in containers next to their data — whether it is in a public or a private cloud or on-premises. GoodData connects to customer data sources and can be integrated with the user authentication setup. Thanks to the declarative APIs, whatever customers do and build with GoodData can be easily stored in a version control system.

GoodData.CN follows the same development procedures as GoodData Cloud, including secure development policies and practices, requirements on access controls, segregation of duties, code review, static and dynamic code analysis, etc. The personnel involved in GoodData.CN development are subject to background checks and properly trained in GoodData policies, procedures, and information security best practices. They are also required to acknowledge the Employee Code of Conduct annually.

Infrastructure around the development, testing, and delivery of GoodData.CN is subject to the same controls and monitoring by GoodData personnel as the infrastructure for GoodData Cloud. The same requirements for any third-party tools or services also apply for the GoodData.CN as for the GoodData Cloud.

GoodData Security Overview

Information Security Policies

GoodData has established a comprehensive set of information security policies, processes, and standards. Our information security management system is based on the international standard ISO 27001:2013. We are building our security procedures and standards upon the National Institute of Standards and Technology's (NIST's) Special Publication (SP) 800 series, and our security controls are mapped against a wide range of standards, such as SOC 2, PCI, NIST SP 800-53 etc.

Our policies are owned and approved by appropriate senior management owners and communicated to affected internal and external personnel. They are reviewed on an annual or ad hoc basis in case of a significant business change to ensure ongoing suitability, adequacy, and effectiveness.

Organization of Information Security

GoodData has appointed a dedicated information security organization. The Head of Security & Compliance has the executive responsibility for information security across the corporation and leads the security and compliance department.

The Head of Security & Compliance also chairs the GoodData Security Council, a cross-functional group of senior stakeholders established for ongoing oversight of the GoodData information security program, both from a design and an effectiveness point of view.

The council's senior roles bring together a wide range of perspectives, ensure efficiency of the security program, and, last but not least, reinforce that information security is a business issue with involvement across the corporation.

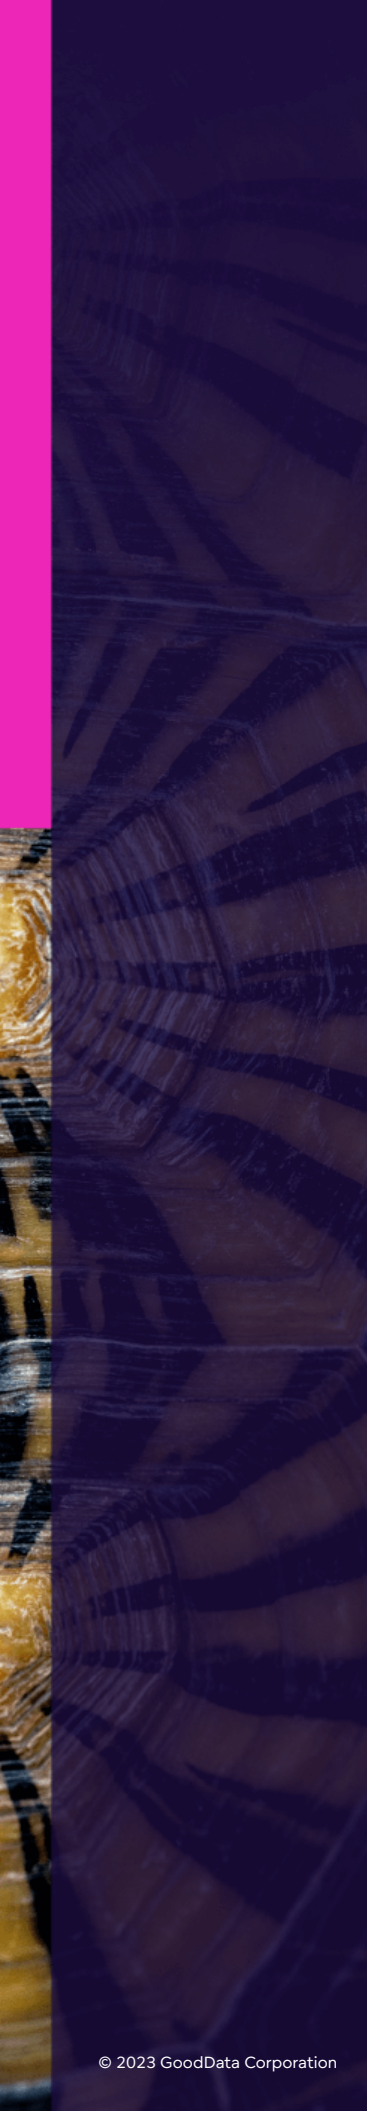

The council meets on a monthly basis to review security events and issues, discuss open and emerging security risks, and to otherwise ensure ongoing alignment between security and business objectives.

All new features and capabilities — from the development of a simple pluggable visualization to the building of a new data center — are managed as projects with the input of a security architect to maintain the integrity of security measures across all components. To ensure that security is built into all aspects of the GoodData Cloud, the GoodData engineering team follows DevSecOps methodology. Our software engineers and operations staff are trained on secure development practices and utilize a wide range of technical processes. These processes are built directly into the continuous integration infrastructure to address risks related to code flaws and vulnerabilities as well as to prevent promotion of changes without proper review and approval.

Events related to security are evaluated, investigated, and tracked to resolution by the Security team. The Security team is also responsible for developing and maintaining a comprehensive security monitoring and security response program both on the technical and organizational levels. They are also responsible for developing and maintaining the corporate patch and vulnerability management program.

GoodData's security and compliance department, together with the internal legal team, monitor the global regulatory landscape to identify emerging data security and privacy-related laws, standards, and regulations and ensure that customer data is protected accordingly.

Human Resources Security

All new employees around the world are subject to an industry standard background check. GoodData has established three levels of security clearance. The highest level, which has the most demanding background check requirements and which has to be regularly renewed, is mandatory for all key security-related roles as well as for personnel with the highest level of administrative access to the GoodData Cloud and critical internal systems.

Contractual agreements include confidentiality clauses as well as the responsibilities for information security and ensure that the relevant employee responsibilities (including the non-disclosure clauses) remain valid after job termination. During onboarding and on an annual basis, employees familiarize themselves with company Code of Conduct which includes ethical behavior standards, employee compliance requirements, guidance for safeguarding intellectual property and maintaining confidentiality, as well as reporting of violations and whistleblowing.

Management is responsible for security compliance in their areas of business. Documented job descriptions further outline specific security-related rights and responsibilities for all roles.

All internal and external employees must complete security awareness training as part of their onboarding and then on an annual basis. Additional role-specific trainings are required for privileged access and for work with highly regulated or sensitive data.

GoodData has an established disciplinary process. Management reviews all compliance violations, and sanctions are taken in the event of high-risk violations. All employees have to sign an acknowledgment of the possible consequences of policy violations, which include loss of access, employment termination, and/or criminal prosecution.

Asset Management

GoodData maintains inventories of relevant IT assets and has established responsibilities and assigned ownership. Our internal data classification policy defines five levels of data classification as well as mandatory data protection requirements on systems that process particularly classified data.

Customer data has the two highest levels of protection and is classified as either “restricted” or, in the case of data subject to strict corporate and/or regulatory requirements such as ePHI under HIPAA, as “highly restricted.”

All internal systems, as well as GoodData Cloud components, are labeled in line with data classification rules to ensure enforcement of adequate data protection.

Procedures for handling of restricted and highly restricted customer data are documented, communicated to all personnel with access to such data, and strictly enforced. GoodData personnel never access customer data without proper business justification and procedures, and technical safeguards ensure that customer data is never stored outside of the GoodData Cloud. GoodData does not use customer data for development purposes, and as a policy rule, customer data is never loaded to employee devices or removable media.

Following the end of a customer contract, GoodData follows a documented procedure to ensure that all customer data is properly removed and, if applicable, the media sanitized and/or securely disposed of. Upon written request, the GoodData security team will provide written attestation of data deletion.

GoodData employee laptops and BYODs follow strict security rules, and centralized monitoring is established to ensure ongoing compliance.

All MacOS- and Windows-based laptops are equipped with centrally managed antivirus protection. All laptops are protected by a firewall, hard drives are fully encrypted, user accounts are protected by strong passwords, and session timeout with screen lock is activated. Additionally, endpoints of non-technical personnel have an MDM solution installed and are fully managed by the internal IT department. Acceptable use rules are documented and communicated to all employees.

Upon termination of employment, laptops are collected and wiped before reuse, and BYODs are deauthorized from company systems. Whenever possible, remote wipe is enabled and triggered upon the report of a lost or stolen device.

Access Control

GoodData has implemented access control policy and enforcing mechanisms which comply with industry best practices. These rules are applied across all internal systems and the GoodData Cloud to ensure that only authorized users with proper business justification have access to internal and customer data. We enforce the principle of least privilege for all systems with data classified as confidential, restricted, or highly restricted.

We apply industry standard password policies and, whenever possible, we use single sign-on for access to internal company systems. Two-factor authentication is enforced for administrative access to the GoodData Cloud and key internal systems. We review access entitlements across all company systems on an annual basis at minimum.

Before being granted privileged access, employees must complete the security training as well as role-specific training related to their access. Based on the sensitivity of access, security clearance of level two or three is mandated (security clearance of level one is mandatory for all employees). The Head of Security & Compliance reviews and approves requests for highly privileged access (including “super admin” and other level three access rights), and the Security and Compliance team monitors ongoing business justification and reviews all privileged access entitlement and usage on a quarterly basis.

For details around end user access control, please refer to GoodData Cloud security section.

Encryption and Cryptography

GoodData uses state-of-the-art cryptography technology to achieve protection of data in transit and at rest and has documented cryptographic policy and standards.

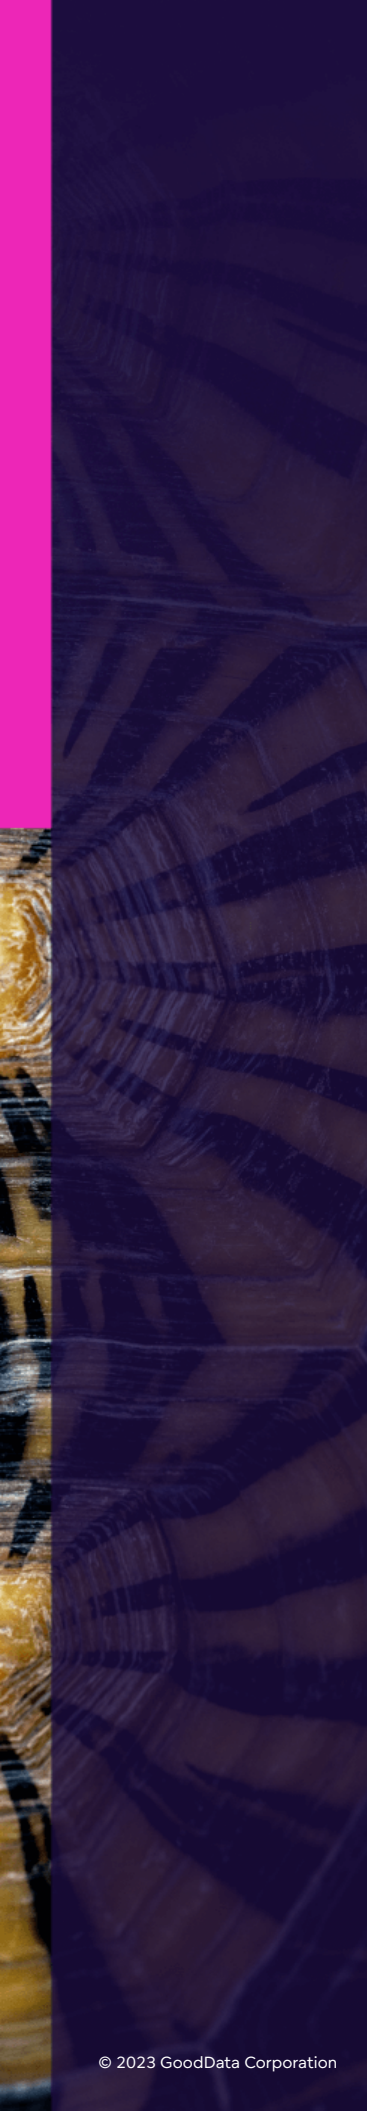

All traffic outside of our data center is encrypted in transit, and we use TLS 1.2 or higher and AES-256 by default. While we may support some of the legacy protocols and cipher suites for compatibility reasons, we systematically deprecate older versions and disable those that have known weaknesses. Our servers enforce HSTS and offer forward secrecy as well as a strong key exchange.

The entire platform infrastructure is encrypted at rest on the file system level and leverages the industry-standard AES-256.

Backups stored outside of our primary data center are encrypted using AES-256-based symmetric cryptography on the client side before being stored in the off-site, encrypted at-rest file system.

For passwords, we use glibc crypt(3) SHA-2-based scheme with an increased number of rounds to mitigate offline password cracking attacks. Administrative sessions are protected via SSH protocol.

Physical and Environmental Security

GoodData Cloud runs on top of AWS infrastructure.

AWS has obtained a wide range of security certifications and conforms to compliance standards, including ISO 27001:2013, SOC 2 Type II, PCI-DSS, HIPAA, and GDPR. GoodData personnel review their audit reports and certificates on an annual basis to ensure ongoing compliance with GoodData physical security requirements.

All data centers also feature at least N+1 redundant HVAC and UPS, diesel-powered generators, and multiple internet connections by independent Tier-1 providers. The physical security adheres to the best practices in the industry and include:

- Keycard protocols, biometric scanning protocols, man traps, review of door logs, and round-the-clock interior and exterior surveillance
- Access limited to authorized data center personnel — no one can enter the production area without prior clearance and appropriate escort
- Assurance that every data center employee undergoes thorough background security checks

All decommissioned hardware is securely disposed of and industry standard media wiping procedures are applied in line with NIST SP 800-88 requirements.

For additional detail on physical and environmental security of AWS, please refer to their data center security overview.

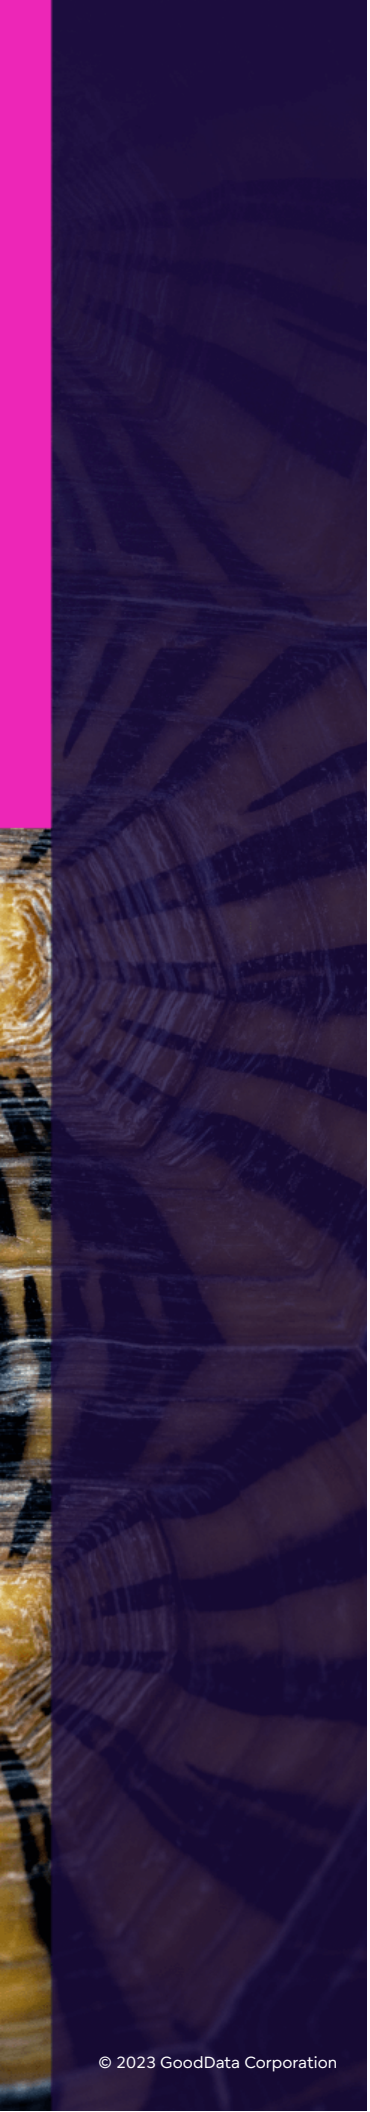

Even though we are a cloud company and do not host any data internally, GoodData protects its offices by industry standard means including key-cards and CCTVs. All visitors must sign an NDA and must be accompanied by GoodData personnel at all times. We implemented a clean desk and clear screen policy to address risks related to undesired exposure of sensitive information to external parties, and we train our employees on security while working remotely or during travel.

Operations Security

A formal change control process minimizes the risk associated with system changes. The process enables the tracking of changes made to the systems and verifies that risks have been assessed, interdependencies explored, and necessary policies and procedures considered and applied before any change is authorized.

The production environment may be accessed only by authorized personnel and when adequately justified by business needs. Operations personnel have administrative access only to the subclusters and services they are responsible for, and all access is fully logged. Access to the infrastructure is controlled via a separate network which is physically isolated from the GoodData corporate network. This ensures that only personnel authorized to access the data center may do so. In addition, the privileged access entitlements are not granted permanently but the authorized personnel may assume it for the period necessary for the completion of the assigned task.

A limited number of key personnel have “super admin” access to the entire platform, which may be used in emergencies. Such access triggers an alert for immediate independent review. All privileged session logs are subject to ongoing monitoring by a session audit tool with 200+ custom alerts for high-risk events and triggers for non-standard activity. The Security team reviews the logs on a daily basis.

Development, testing, and production environments are strictly separated both on the logical access level and on the network level to reduce risks related to unauthorized or unexpected changes to the production environment.

The GoodData Cloud is protected both internally and externally by firewalls and security groups. We use industry standard hardening procedures, including building minimized Docker base images and running the images with the least privilege possible, changing default system passwords or disabling implicitly created accounts, using AWS IAM roles and trust policies rather than credential-based access whenever possible, and making sure that firewalls let through only explicitly allowed traffic.

We apply infrastructure-as-a-code and configuration-as-a-code principles to ensure consistent application of our security standards as well as for in-time monitoring and alerts in case of unintended or unauthorized changes.

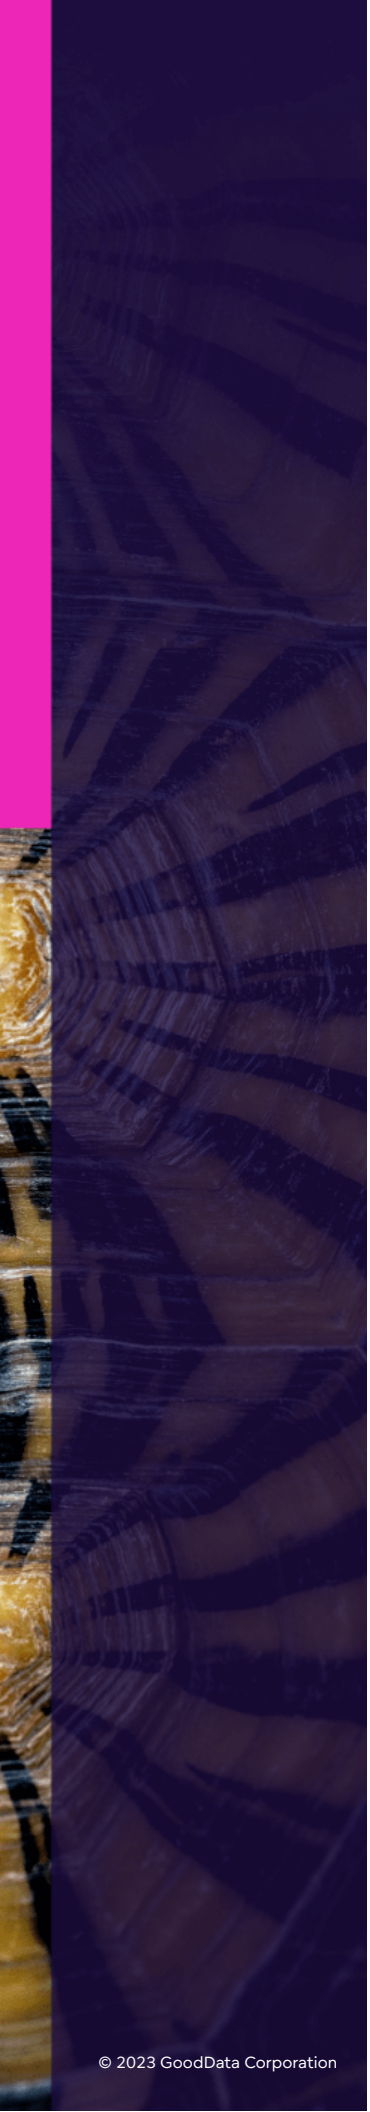

The entire production infrastructure as well as all platform components are monitored, and alerts are addressed by operations personnel 24x7x365. The platform team is responsible for capacity monitoring and planning to ensure the timely provision of new hardware as our customers' usage of the platform grows.

The log management system is set up in line with NIST SP 800-92 recommendations. Logs are securely transferred to the centralized log management system and protected from unauthorized access. All systems have synchronized clocks via NTP. Logs are available for 90 days in a SIEM and then for a year in secure offline storage.

We have established an industry standard patch and vulnerability management procedures. The docker containers are scanned both at build time and in runtime. Our operations personnel monitor relevant security groups, upstream software providers, as well as hardware vendors for patch and vulnerability notices, and we have defined SLAs for remediation. Critical patches are handled via incident management procedures. Compliance with SLAs is monitored by the service delivery function and is reviewed by management on a monthly basis.

Network Security

GoodData Cloud containers run in kubernetes cluster with hardened and secured network configuration, including strict separation of networks and firewalls that ensure that only designated externally facing microservices are reachable from the load balancers.

Consistent with our DevSecOps approach, we maintain a configuration-as-a-code approach for all configurations, including deployment schema, network security and firewall rules, and have alerts for any discrepancies between the approved configuration and production settings.

System Development, Maintenance, and Acquisition

We follow industry standard secure development life cycle practices. A formal change control process minimizes the risks associated with system changes. The process enables tracking of changes made to the systems and verifies that risks have been assessed, interdependencies explored, and necessary policies and procedures considered and applied before any change is authorized. We have integrated static and dynamic security testing in our CI/CD infrastructure, and peer code review includes secure development considerations.

Development of new features follows agile project management principles. Security architecture considerations are part of architecture design and reviews. Before roll-out to production, we review the implementation against design and for new or significantly modified components, and we also execute external penetration tests.

Vendor Management

All vendors with access to the GoodData Cloud are rigorously reviewed for security and compliance practices, and we have contractual arrangements in place to ensure their ongoing compliance with our security requirements. We review the contractual performance of these vendors and their adherence to security and compliance requirements annually.

Security Incident Management

GoodData has established an industry standard security incident response plan. We train our staff to ensure that all potential security incidents are identified and reported in a timely manner. Our security operations team is on call 24x7x365, and we have defined protocols and escalation trees for the handling of security incidents and, when required by the nature of the incident and applicable contractual commitments and regulatory requirements, for the notification of the affected parties as well as the authorities. Procedures for collection of evidence ensure chain of custody.

Following the resolution of a security incident, the GoodData security team conducts root cause analysis and, as applicable, implements changes to its technology and procedures to prevent regressions.

GoodData maintains industry standard commercial insurance covering cybersecurity incidents and has engaged external breach services to assist in case of a major security incident.

Business Continuity and Disaster Recovery

GoodData follows international standard ISO 22301 for its business continuity and disaster recovery management. We have completed business impact assessments for all key corporate processes and have documented business continuity and disaster recovery plans.

To achieve the committed platform availability SLA and reduce the impact of failures, GoodData applies high-availability architecture principles on the software and hardware level and ensures its data center providers have adequate redundancies in the infrastructure.

While GoodData Cloud does not store any customer data, we still perform daily backups of all customer configurations and move them regularly to a secure, highly available, and durable off-site storage. Our disaster recovery plan addresses major disruptions to GoodData facilities, key internal systems, and the GoodData Cloud, and we can restore production operations at another public cloud data center. We monitor the backup process, test our ability to restore the backups regularly and conduct a disaster recovery test on an annual basis.

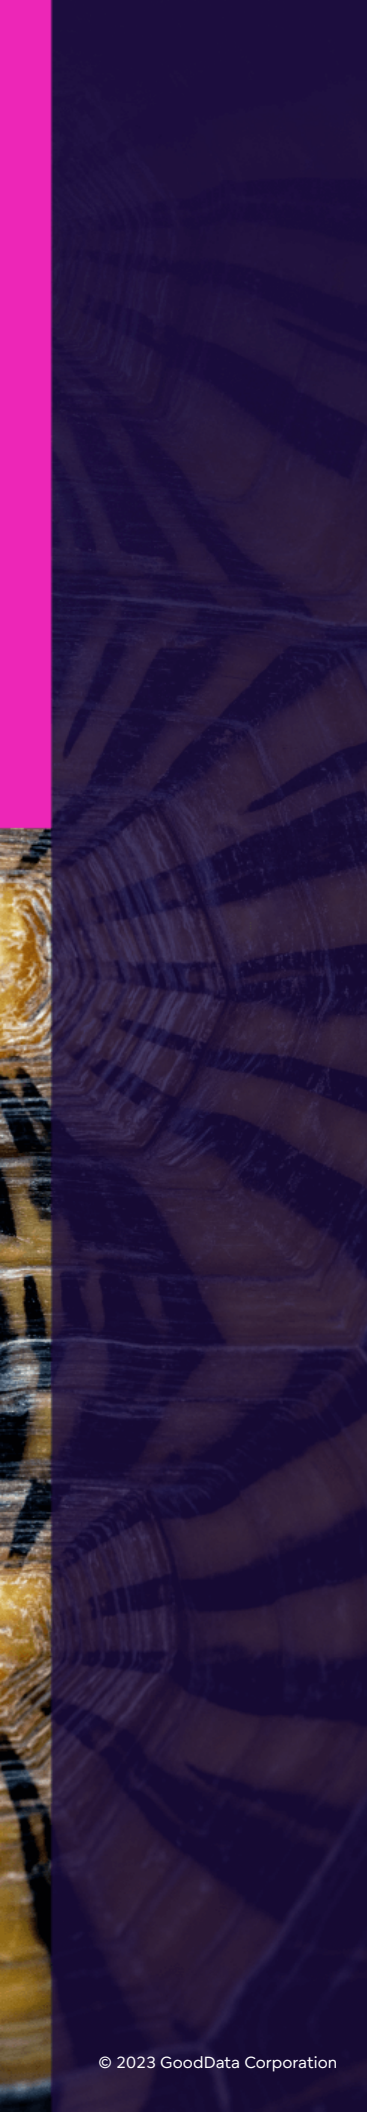

Compliance

GoodData complies with a variety of data protection standards. We undergo an annual SOC 2 Type II audit review by an independent third party, maintain compliance with the ISO 27000 standards family, and build our security practices upon industry standards, including applicable NIST and OWASP standards and recommendations.

We have policies and procedures to ensure appropriate protection of PII and personal data both in the platform and as part of our business operations. We comply with GDPR, CCPA, HIPAA and similar privacy regulations globally, and we offer signing data processing agreements with our customers.

We monitor the emerging legislation and standards to maintain compliance and achieve best-in-class security of our products.

We continuously monitor and regularly review and audit our security compliance. We do this on the policy and on the technical level both internally and by using external penetration and vulnerability testing providers and auditors.

On an annual basis, external reputable penetration testers conduct a comprehensive penetration test of the complete GoodData Cloud API set. The entire GoodData infrastructure (including GoodData's office network) is subject to quarterly "weakest link" penetration tests and vulnerability scans. We partner with two independent penetration and vulnerability test providers who alternate on the different test types to achieve above-standard coverage and depth of testing.

Conclusion

Here at GoodData, we pride ourselves on the vigilance we employ to protect our customers' data assets, and we continually stress that a mature security organization requires coordinated dedication across technology, policy, procedures, and people. This dedication is underscored by the risk-based approach laid out in this document to demonstrate strength at every layer of security, minimizing any potential vulnerability or weakness.

We want our customers to know that their data is adequately protected by this approach, and we welcome the opportunity to discuss these practices and approaches further. We also encourage our customers to consider the criticality and sensitivity of their usage of the GoodData Cloud and, in line with the recommendations provided in this whitepaper, to implement adequate technical and administrative safeguards to achieve the desired level of security. The GoodData security team is looking forward to assisting with the implementation.

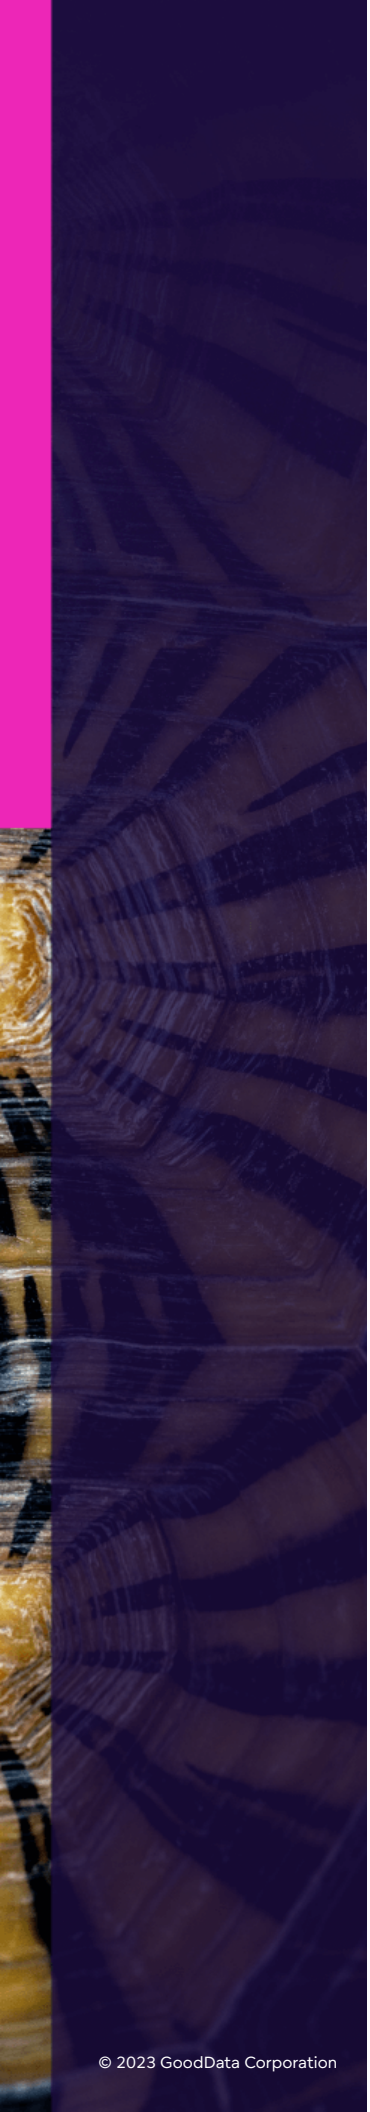

GoodData Corporation reserves the right to amend, modify, delete or remove this Security White Paper, at its sole and exclusive discretion, at any time. All information contained herein is provided “as-is”, and GoodData disclaims all liability for itself and its affiliates, licensors and suppliers, with respect to the descriptions, statements and contents of this Security White Paper.