

# Security Overview

<b>Introduction</b>	<b>4</b>
<b>Information Security</b>	<b>4</b>
<b>Platform Overview</b>	<b>5</b>
<b>Certification and Accreditation</b>	<b>6</b>
<b>Physical Infrastructure</b>	<b>6</b>
Amazon Security Model	<b>6</b>
<b>Virtual Environment Security</b>	<b>7</b>
Virtualization	<b>7</b>
Network Security and Intrusion Detection	<b>7</b>
Data Storage	<b>8</b>
<b>Authentication, Authorization and Single Sign On</b>	<b>9</b>
<b>Application-level Security</b>	<b>9</b>
Logical Boundaries	<b>9</b>
Administrative Controls	<b>9</b>
Transport Layer	<b>10</b>
<b>Data Security</b>	<b>10</b>
<b>Organizational Security and Change Management Processes</b>	<b>11</b>
Access Policy	<b>11</b>
Code Development Review Process	<b>11</b>

<i>Design</i>	<b>11</b>
<i>Development and Test</i>	<b>11</b>
<i>Release</i>	<b>12</b>
<i>Incident Reporting and Response Process</i>	<b>12</b>
<i>Privacy Policy</i>	<b>12</b>
<i>Law Enforcement and Legal Access</i>	<b>12</b>
<i>Personnel Management</i>	<b>12</b>
<b>Data Backup, Disaster Recovery, and Data Disposal</b>	<b>13</b>
Data Replication Backup and Archiving	<b>13</b>
Data Deletion and Disk Destruction	<b>13</b>
<b>Regulatory Compliance</b>	<b>13</b>
Technical Compliance	<b>13</b>
<b>Appendix I: Control Objectives and Related Controls</b>	<b>14</b>
CONTROL OBJECTIVE 1: SYSTEM ACCESS	<b>14</b>
CONTROL OBJECTIVE 2: SYSTEM CHANGES	<b>14</b>
CONTROL OBJECTIVE 3: CLIENT SEGREGATION	<b>15</b>
CONTROL OBJECTIVE 4: INCIDENT RESOLUTION	<b>15</b>
CONTROL OBJECTIVE 5: BACKUP INTEGRITY	<b>15</b>
CONTROL OBJECTIVE 6: DATA INTEGRITY	<b>16</b>
CONTROL OBJECTIVE 7: OUTPUT INTEGRITY	<b>16</b>
<b>Appendix II: Customer Control Considerations</b>	<b>17</b>

## Introduction

Securing customer information is a critical component of GoodData's business value proposition. As an on demand service, security must be an integral part of both the design and operations of GoodData. In order to accomplish this, GoodData has included security in all levels of our technology and operations from the onset of the company.

GoodData's commitment is to invest in the technology, people, and process to ensure that data you have entrusted with us is safe, secure, and private.

This document describes how security has been designed into GoodData, as well as the operational controls we have put in place to ensure the security of both customer data and the GoodData infrastructure. Specifically this paper covers:

- **GoodData Platform:** an overview of platform architecture
- **GoodData Security:** an overview of security implementation across multiple layers - physical, virtual infrastructure, software infrastructure security as well as application and administrative security features
- **GoodData Operations:** a description of GoodData's operational security practices - organizational security and change management processes; data backup and disaster recovery; and compliance with industry regulations
- **Security Control Objectives:** GoodData's audited control objectives and description of related controls

## Information Security

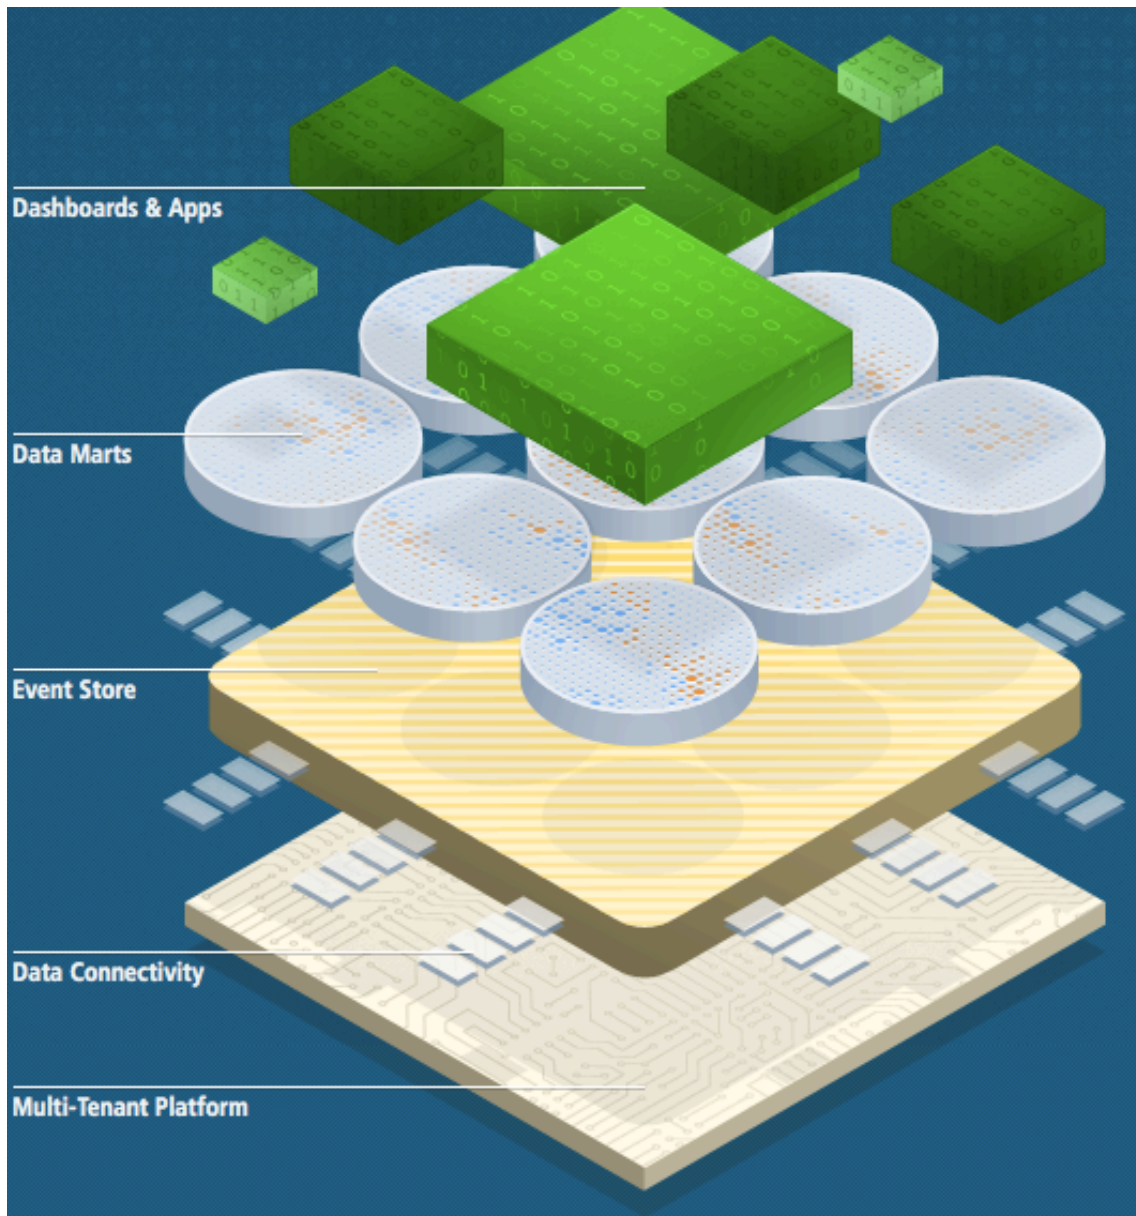
We believe that effective information security has much to do with people and processes as it does with technical protection measures. This document outlines the standards, processes and measures we apply to protect information entrusted to us.

Before digging deeper into GoodData's security, let's look at the overall points that guide our security strategy:

- Security is multi-layered - it must be addressed at all layers - physical, application, metadata, data, etc.
- GoodData is a service provider - the burden is on GoodData to provide a secure service for our customers
- We are built atop Amazon Web Services (AWS) - this means we inherit aspects of our security from Amazon
- We augment AWS security by applying selected technologies, such as key-based authentication, data encryption, platform monitoring and firewall configuration, as well as policies related to change and incident management
- Whenever possible, we make our security model open and pluggable to accommodate customer-specific requirements such as 3rd-party authentication, user account management or primary storage encryption

## Platform Overview

The architecture of the GoodData platform is designed to align the individual functional and security aspects in well-defined layers. This structure provides robust data processing and the privacy and security of the customers' data.



The GoodData platform provides the data architecture to accommodate a large number of customers without requiring a separate instance for each individual customer. The foundation of the security constrains lies within the individual layers of the architecture.

User authentication and authorization on the web API layer ensures a valid identity is attached to each request and authorized for access to required resources.

Logical security measures and relationships between individual users, projects, meta-model and data stores are configured within the control layer. Multiple authentication providers are available to support various authentication methods enabling Single-Sign-On (SSO) and embedded applications.

Access to operational tools used to support customers' individual projects is restricted to allow processing of a single project at a time. Separation of client requests is maintained by the ROLAP engine such that each client request is separated into a set of tasks that are executed independently without sharing contextual or other information.

Logical separation of the meta-model and data is established on the storage cloud where each project is configured as a separate physical entity. Connection to the data store is restricted through access credentials configured and stored within the control layer.

Input and output is protected by SSL encryption technology.

## Certification and Accreditation

Participation in relevant industry certification and accreditation programs is intended to provide us, our investors, and our customers with the highest level of assurance regarding our operations, infrastructure and controls.

In January 2011 we attained a recurring Statement on Auditing Standards No. 70: Service Organizations, Type I ([SAS70](#) Type I) certification. We expect to obtain Type II certification by Q4 2011. Additionally, we've attained independent web application security certifications from [TRUSTe](#), [VeriSign](#), [GoDaddy](#), and [salesforce.com](#).

## Physical Infrastructure

GoodData's on-demand analytics platform is built and hosted on top of Amazon Web Services (AWS), leveraging its scalable distributed computing infrastructure. As a result of the relationship with AWS, GoodData can:

- Deploy the GoodData platform across multiple geographical regions and different availability zones (physical data centers) for redundancy and high-availability.
- Migrate infrastructure to any availability zone and geographical region.
- Instantly provision any type of server (node).

All GoodData platform servers are allocated to the respective security groups, characterized by specific security settings (TCP/IP level), supplemented by an individual instance level stateful firewalls. Linux operating system images are created using tools provided by AWS, with the regular patch management performed by and periodically reviewed for security vulnerability by a 3rd party as part of general security reviews.

These maintenance procedures over operating system images, patch management and security hot-fixes are subjected to GoodData's regular change management process.

## Amazon Security Model

Amazon maintains a number of data centers geographically distributed across the US, European Union, and Asia. The company applies a range of industry-leading measures to protect the physical infrastructure and the software systems that underpin the Amazon.com site and the EC2, EBS and S3 services.

Additional information regarding Amazon's security model is available from <http://aws.amazon.com/security/>, and in their AWS Security White Paper, which describes the provisions of AWS services from the following perspectives:

- Certifications and Accreditations
- Physical Security
- Secure Services
- Data Privacy

Service-specific technical details, applicable policies including data retention, the following services are mentioned in particular:

- Amazon Elastic Compute Cloud (EC2) Security
- Amazon Simple Storage Service (S3) Security
- Amazon SimpleDB Security

## Virtual Environment Security

### Virtualization

Amazon EC2 service is built on a large number of geographically distributed servers running a hardened version of the Xen hypervisor (virtual machine layer). Amazon administrators do not have access to our server's virtual images and cannot login to GoodData server instances.

Amazon recommends to implement either token or key-based authentication to access the EC2 virtual hosts. GoodData has chosen to implement the key-based authentication. In order to further protect the customer data we leverage the combination of pass-phrase protected RSA keys and EC2 security zones to implement security rules further described in the following sections.

### Network Security and Intrusion Detection

All network access to the Amazon EC2 servers is protected by a multi-layered firewall operating in a deny-all mode; Internet access is only permitted on ports we explicitly open. In order to reduce the network attack surface our virtual servers operate an enterprise version of the Linux with a minimum subset of services required.

We follow the established practice of implementing a 3-tier security model with the web servers at the front-line, application servers in the demilitarized zone and the database servers behind an additional firewall. The security tiers are implemented using AWS security zones. We audit the production settings on regular basis and archive the audit results in our internal document repository.

Permissions to carry out network configuration changes are limited to select GoodData administrative staff, require access to our protected certificate and the RSA private key, and are subject to ongoing review processes.

At the application level, all network data is examined for signs of exploitation of programming errors such as cross-site scripting. GoodData operations is notified immediately of any occurrence of such traffic.

To reduce the risk of buffer overflow attacks, all software that processes Internet-based data undergoes a formal security review and audit process.

Our production architecture also incorporates a Host-Based Intrusion Detection System (HIDS) to provide visibility to any suspicious activities and provide the information necessary to respond to incidents.

GoodData production cluster EC2 hosts are running in three separate security zones:

- Tier 1: web servers (static content and balancers) publicly accessible via HTTPS
- Tier 2: application servers (dynamic content) (demilitarized zone)
- Tier 3; metadata and data warehouse servers

Only Tier 1 above is publicly accessible. No connections are possible between Tier 1 and Tier 3. There is no SSH/Telnet access to any zone (Tier 1 – 3). A dedicated management instance resides in a separate security zone accessible exclusively via SSH authentication using RSA key-pair (select GoodData super-admins). When accessing production nodes in Tiers 1 – 3 from the management instance, an additional pass-phrase protected RSA key is required.

Journalling is enabled on all production nodes and all activity is captured in log files.

## Data Storage

All access to the Amazon S3 storage service is virtualized. The AWS proprietary disk virtualization layer automatically wipes every block of storage used by the customer, and guarantees that one customer's data is never exposed to another. We are using Amazon Elastic Block Storage (EBS) as a primary data store. While we do not offer encrypted primary storage as a standard service due to performance considerations, it is available for customers on an as-needed basis. Backups are stored on Amazon S3 systems and encrypted using GnuPGP using 128 bit AES encryption and 64 bytes-long randomly generated pass-phrases (Gnome Password Generator).

## Authentication, Authorization and Single Sign On

GoodData architecture relies on a centralized authentication and authorization security framework to control access to services.

Our security framework enables us to enforce security policy with regard to:

- Password strength: algorithms to set minimum password length and complexity
- CAPTCHA filters that use human-readable images to reduce the risk of automated attacks against customer data

In communication between our virtual servers we rely on an additional set of authentication mechanisms and protocols to control access to customer data. As an example, access to any customer database is only permitted by a specified set of front-end servers. This is intended to prevent unauthorized services or systems from accidentally (or maliciously) retrieving or modifying customer data.

Unsurprisingly, we use our service to host our own business operations data and analytics. We face the same security risks our clients do, needing to protect our own customer, engineering and operational data.

## Application-level Security

The GoodData service provides a range of customer-level security mechanisms to allow customers to fine-tune their GoodData service to their specific requirements. Every granular action in our platform can be controlled by a permission. Permissions are grouped to roles and are always global.

### Logical Boundaries

Security and privacy is enforced at the project level. A project contains a data warehouse and its users. Users in a project can never see into other projects and each project has database instance affinity. User roles inside projects are either Admin or Editor. Admins have full control over projects, data and invitees. Editors cannot delete reports or dashboards, and cannot invite or remove users from a project.

### Administrative Controls

The GoodData Analytics Platform is built as a self-service Web 2.0 application enabling users to administer their own accounts and easily collaborate with the other users of the platform.

The following activities are completely self-service in the GoodData Analytics Platform:

- Account registration and activation
- Password reset
- Project (data mart) creation and administration
- Project invitations and sharing (project owner and certain roles only)
- Suspending user access to projects (project owner and certain roles only)

## Transport Layer

All user Web access relies on the HTTP/HTTPS protocols, and appropriate authentication. All traffic between the GoodData servers and the user's Web browser uses requires HTTPS and is encrypted using 128-bit SSL standards.

Programmatic access to the GoodData services requires the use of GoodData API Protocol, a REST/JSON interface for accessing and modifying data. The GoodData API interface and protocol are based on the Google GData API model.

All GoodData API access requires authentication, which can be performed by passing in user credentials (over HTTPS) Authentication tokens expire after a pre-defined period and the user needs to re-authenticate passing in user credentials using the internal GoodData authentication authority (login server). 3rd party authentication (Single Sign On) is also supported.

## Data Security

Platform architectural patterns are strategically selected around data confidentiality, integrity and availability. These include data segregation, consistency checks (MD5), log management and active monitoring using situational awareness algorithms. Data transport and long term storage is protected using industry standard methods of encryption (SSL/TLS, symmetric-key cryptography).

Strict process separation (sealed) is a built-in design feature of all GoodData software development and operational life-cycles. Multi-tenant security patterns employed provide effective isolation and sealing of data and metadata even while sharing the same physical storage grids. Continuous monitoring and situational awareness lets us analyze and log known data movements, but more importantly, lets us quickly identify anomalies and outliers for immediate reaction.

GoodData is database technology-independent since users are interacting with a logical data model (LDM) that defines attributes, facts and their correlations rather than the physical data layer (PDM). All metrics and reports are defined at the LDM layer and correlate to the underlying physical data model. Data visibility can be restricted using mandatory filters and via metadata security. For instance queries for a user or group can be restricted to a specific region, or access to sensitive datasets may be restricted.

## Organizational Security and Change Management Processes

To be effective, security must be embedded into the organizational culture and everyday business processes. This is even more important for software-as-a-service providers such as GoodData.

### Access Policy

GoodData policy is to provide system access only to appropriately trained staff who require a specific level of access to perform authorized tasks. Our internal systems enforce unique user ID's and strong passwords, and limits password reuse. We rely on industry-standard security systems and standards including LDAP, Kerberos and RSA to manage access. Using a combination of these we ensure that only authorized users can gain access to servers, logs, customer information and system configuration information.

Logical access to the production environment by our employees is limited to the core operational personnel only, using encrypted session (SSH) and public/private key cryptography. All keys are stored within a credentials vault. Access requests, grants and revocations are periodically reviewed. All changes to access rights are logged and are based on roles and job responsibilities; the approval process maintains audit records of all changes.

Access to the production infrastructure servers for our platform is restricted on the network level. Each server is accessible only from one access node which can be accessed only by authorized GoodData operations personnel. A specific set of credentials is required for authentication; therefore access to the access node server does not automatically enable access to production servers. The Director of Operations monitors the revoking of access to employees who become inactive or change job roles.

## Code Development Review Process

### Design

Software design in GoodData is a two phase process. The first phase is the requirements analysis phase which includes both functional and non-functional requirements document. The second phase is the technical analysis phase which results in a detailed technical specification document. Both documents require a three-way sign-off between the product management, engineering and operations.

Once a conclusion is reached between all three stakeholders, engineering proceeds with the technical analysis phase which results in a detailed technical specification. The technical specification is again reviewed and eventually approved by the product management from a budgeting perspective and by operations from a security and operating cost perspective.

During both phases, the engineering and operations teams carefully consider the impact of the newly introduced features or changes on GoodData platform security.

### Development and Test

All source code and other artifacts which are a part of the product (a release bundle) are subject to version control and are stored in a centralized version repository.

Upon completion, the new code artifacts need to pass the following quality controls before they are allowed to the main product code base: (1) code review by a peer and the team lead, (2) sufficient test coverage as specified by the applicable QA policy 3) automated test-suite pass covering both the new developments and the existing code.

The main product branch is a subject to continuous integration (automated testing) so that any regressions not captured by the checklist above are discovered and corrected as soon as possible. The continuous integration process includes the full cycle product build, packaging and deployment in order to simulate the actual production deployment as closely as possible.

The development cycle reaches the QA phase when all the features approved for the upcoming release have reached the main product code base. One or more release candidates are subsequently built from the main product code base and are subject to extensive manual testing. Each release candidate test cycle has its own test plan and a written record of passed and failed test cases linked to the defect tracking system is kept.

### **Release**

The release candidate which reaches QA acceptance is subsequently scheduled for a production release. The entire process (including production data migration) is first tested in the pre-production environment which has the deployment topology identical to the production environment and is completely separated from the development and QA environments.

If the result of the test upgrade passes all the prescribed tests and validations routines, the release is subsequently applied to the production environment. A deployment plan together with the deployment log is kept for each production deployment. The GoodData operations staff is required to comment on and explain all the manual steps taken during the deployment, which are specific to that particular release.

### **Incident Reporting and Response Process**

We proactively monitor the platform for security incidents including alert notifications generated by our systems, alerts generated by Amazon, open source and industry alerts, and community alerts.

When an alert is raised we first assess the risk level. Based on this assessment we select and launch the prescribed response process. Documented escalation procedures and communication protocols clarify when and how an escalation takes place, and who is notified.

### **Privacy Policy**

GoodData maintains a strong privacy policy to protect customer data. Our privacy policy is accessible at <http://www.gooddata.com/privacy-policy>.

### **Law Enforcement and Legal Access**

We are obligated to protect access to customer information while also abiding by the law. We do not release information about our customers or customer data to 3rd parties.

Information can only be obtained from us through a valid legal process such as a search warrant, court order, or subpoena. If legally permitted, we will notify the organization whose information is being sought and allow them 21 days to respond.

### **Personnel Management**

Our hiring practices ensure that all staff are qualified for their functional responsibilities and hold appropriate certifications or accreditation, if required. At a minimum, these practices include verification of the individual's education and previous employment as well as a reference check. Based on statutory environment and position being applied for, additional background checks may be performed.

On acceptance of employment, all staff are required to sign a confidentiality agreement and acknowledge the receipt and conformance with the Employee Handbook. Security and privacy of customer information and data are highlighted during the employee orientation and in the Employee Handbook.

The employee on-boarding process includes a mandatory security orientation session during which they are instructed about our security policies and procedures. All employment contracts include a clause clarifying staff member's responsibility to communicate significant issues to GoodData's management team.

## Data Backup, Disaster Recovery, and Data Disposal

### Data Replication Backup and Archiving

GoodData Analytics Platform uses Amazon Elastic Block Storage (EBS) as its primary data store. If changes are detected in a particular data set, a backup is created on Amazon Simple Storage Service (S3). The backups are archived for minimum of 1 year.

A basic level of data redundancy is built into the Amazon EBS and S3 services. The virtualized storage provider ensures that there are no single points of failure, and that mechanics of replication and automated provisioning are managed behind-the-scenes. Since GoodData provides data analysis functionality, customers can and should maintain their own backup of their source data.

### Data Deletion and Disk Destruction

In case a customer closes an account with GoodData, we maintain the backups and archives for a period granted according to the service plan effective at the date of termination. A customer may request complete and permanent deletion of its data by contacting GoodData support. The unit on which a data destruction can be requested is an entire project (a data mart). GoodData does not support data deletion on individual report or data load level.

## Regulatory Compliance

Many of our customers operate under a complex statutory environment that governs retention and management of data. It is GoodData's intention to comply with our customers' specific regulatory requirements. We have attained a recurring Statement on Auditing Standards No. 70: Service Organizations, Type I ([SAS70](#) Type I) certification. Additionally, GoodData is a licensee of the [TRUSTe® Privacy Program](#) and abides by the EU Safe Harbor Framework as outlined by the U.S. Department of Commerce and the European Union.

### Technical Compliance

GoodData provides a range of technology tools and measures to assist our customers in meeting their obligations. These include data and transport encryption technologies, data access API's and administrative controls.

- **Data archiving:** information managed by the GoodData infrastructure can be retrieved by customers using GoodData API's; this API can be used to export data (including collaboration data) periodically
- **Encryption:** all data can be encrypted in transit. Certain regulations, for example HIPAA, require that encryption be used between network endpoints to prevent network sniffing

## Appendix I: Control Objectives and Related Controls

As part of the Statement on Auditing Standards No. 70: Service Organizations, Type I ([SAS70](#) Type I) certification attained by GoodData in January 2011, PricewaterhouseCoopers LLP audited the following GoodData control objectives and related controls.

### CONTROL OBJECTIVE 1: SYSTEM ACCESS

Controls provide reasonable assurance that logical access to system resources (software, system files, data, and operating systems) is restricted to properly authorized individuals.

1. Segregation of duties is achieved by organizing operations into functional groups. Responsibilities are segregated between Product Management, Development, Operations, and Release Management / QA.
2. Privileged access is restricted to authorized personnel in the GoodData platform.
3. Procedures exist to ensure adherence to policies for requesting, granting and terminating access to permissions on the GoodData platform to GoodData employees based on position and job function and manager approval.
4. Permissions and privileges assigned to GoodData platform are reviewed at least annually.
5. Only authorized personnel have access to grant and revoke user access right and permissions to the individual systems according to requirements specified by Information Security Policy.
6. An Information Security Policy is in place and the policy is reviewed and approved annually by an executive management team member.

### CONTROL OBJECTIVE 2: SYSTEM CHANGES

Controls provide reasonable assurance that changes to systems are authorized, tested, approved, documented and implemented appropriately to production.

1. A formal technology roadmap is maintained and updated at least semi-annually, aligned with business objectives.
2. Each change to the GoodData platform is authorized, tested and approved by a person other than a developer.
3. Development and testing environments are segregated from the production environment on the GoodData platform.
4. Regression tests are performed prior to promotion of each change to production. Output of the new version is compared to output of the old version. Deviations from expected output are researched and resolved.
5. A configuration management system is in place that requires all changes to be authorized by Director of Operations.
6. Access to the configuration management system is restricted to authorized personnel.

### **CONTROL OBJECTIVE 3: CLIENT SEGREGATION**

Controls provide reasonable assurance that client access to other client environments is segregated both for processing and backup applications.

1. GoodData platform client user access permissions are assigned to the individual client environments within the GoodData platform using automated authentication and authorization mechanisms.
2. GoodData client user access to individual client projects within the GoodData platform is restricted from accessing other client projects.
3. GoodData platform system maintenance tools can access only a single project at a time.

### **CONTROL OBJECTIVE 4: INCIDENT RESOLUTION**

Controls provide reasonable assurance that operational incidents are identified, recorded, analyzed, and resolved in a timely manner.

1. Alerts are configured within the GoodData platform to ensure minimum availability and services levels.
2. GoodData platform service owners are responsible for event classification between problem and incident management and their severity level.
3. All incidents are logged, assigned to responsible person, tracked, and resolved within required time based on the severity level.
4. All incidents are logged with one of the three severity levels and assigned to a responsible person from the incident response team within GoodData platform business operational hours. Notification are sent to the operations personnel, customer support and company management. Incidents are tracked, reviewed and resolved by the respective service owners in accordance with GoodData incident response procedures.

### **CONTROL OBJECTIVE 5: BACKUP INTEGRITY**

Controls provide reasonable assurance that systems are configured to perform backups on a periodic basis and that procedures are employed to maintain integrity of stored information.

1. GoodData platform is configured to enable a periodic backup of client's data and environment.
2. Restoration of GoodData platform data is tested on a quarterly basis.
3. Data mart referential integrity is ensured by automated checks and checks running on-demand both GoodData and user initiated.
4. Scheduled operations are monitored, failures are automatically evaluated against set thresholds, and breaches result in alerts that are followed up by the respective service owners.

**CONTROL OBJECTIVE 6: DATA INTEGRITY**

Controls provide reasonable assurance that client-managed data load to the GoodData platform is complete and accurate.

1. Client upload utility provided by the GoodData platform generates an error message when data transmission fails to upload completely and accurately.
2. Data load structure and consistency is checked by GoodData platform upon the upload request against the data model. Inconsistencies are reported in an error message to the GoodData platform client users.
3. Client- managed data load is always transmitted using SSL encryption.

**CONTROL OBJECTIVE 7: OUTPUT INTEGRITY**

Controls provide reasonable assurance that output is complete and accurate and the information processed and produced by the system is complete and accurate.

1. Automated integrity checks run in the GoodData platform that monitor customer report requests and return an error message in case the intended report would be based on inappropriate data fields.
2. The GoodData platform application output displays error messages if the report cannot be computed completely or within a reasonable time.
3. GoodData platform output data are transmitted to the client using SSL encryption.
4. Client specific emails with reports are generated automatically by the GoodData platform. Outgoing email schedules are evaluated every 30 minutes.

## Appendix II: Customer Control Considerations

GoodData's platform is designed and operated with the assumption that certain policies and controls are implemented by its customers. The following controls should be in place at user organizations to complement the controls at GoodData. The list of customer control considerations does not represent a comprehensive set of all controls that should be employed by the customer. At a minimum, the following controls should be in place:

- Controls to provide reasonable assurance that information input into the GoodData platform including but not limited to sensitive, confidential, personally identifiable or other information which requires protection based on the user organization's requirements, is approved prior to upload.
- Controls to provide reasonable assurance that information uploaded onto the GoodData platform is complete, accurate and valid and that reports and other outputs from the GoodData platform are reconciled to inputs.
- Controls to provide reasonable assurance that access to the GoodData platform is restricted to authorized personnel of the user organization.
- Controls to provide reasonable assurance that administrative access to projects is delegated to authorized personnel of the user organization.
- Controls to provide reasonable assurance that user roles that enable changes to information uploaded into the GoodData platform are restricted to authorized personnel of the user organization.
- Controls to provide reasonable assurance that data transmission from user organization systems are properly configured to support SSL encryption.
- Controls to provide reasonable assurance that scheduled maintenance, updates to the GoodData service and unplanned service periods are monitored and evaluated for their impact on user organizations' controls.