INFRAVISION
SIMPLY EFFECTIVE

InfraVision Integration Service
Security & Compliance

Table of contents

1. Introduction

Service management systems are rarely used stand-alone. Often there is a desire (or need) to integrate with other systems and interact with various external data sources. Vendors often offer APIs to facilitate building these integrations. 4me is no exception, there are numerous APIs available that enable the customer to build complex integrations and automations.

The availability of APIs is only part of the overall solution however. Most integrations still need middleware to take care of workflow-specific triggers, data translation and transformation, logging.
Some organizations provide and support this middleware layer themselves, others rely on vendors to provide this functionality.

InfraVision created a purpose-built integration platform to quickly develop fully functional, reliable integrations. Most of these integrations process customer-specific data so security and compliance is extremely important.

This document describes the organizational and technical measures we have in place to maximize data security.

2. Definitions

The InfraVision Integration Service consists of a fully cloud-based infrastructure provided by Amazon Web Services (AWS) and application code, developed and maintained by InfraVision. Each integration runs in a separate containerized instance which guarantees that data is fully segregated between instances.

2.1. Infrastructure

All infrastructure is provided by Amazon Web Services (AWS). The configuration of the infrastructure is fully managed by InfraVision and the provisioning and maintenance is highly automated.

2.2. Instance

A containerized environment that runs integration code for a specific customer. Containers are volatile and are deleted and recreated every time the code is updated or when the integration job is completed.

The infrastructure for running the containers is automatically provisioned and decommissioned by AWS based on configuration settings managed by InfraVision. Production and QA instances run on different clusters in AWS.

2.3. Data controller

The owner of the data which is processed by an integration. The data controller is responsible for the quality and integrity of the data and meeting the requirements of privacy legislation in the jurisdictions in which they operate. In the scope of this document, the customer is considered to be the data controller.

2.4. Data processor

The data processor is processing the data provided by the data controller according to the functional specifications of the integration. In the scope of this document, InfraVision is considered to be the data processor.

3. Information Security Risk Management & Compliance

In this chapter we will describe the organizational measures which are in place to maximize the security of your data.

3.1. Frameworks

Security frameworks help to establish trust and demonstrate focus on information security. At the time of writing this document, InfraVision is working on getting an ISO 27001 certification. Our goal is to be certified before the end of July 2021.
InfraVision holds the Data Pro certificate which confirms that we fully comply with the GDPR regulations.

AWS has numerous security related certifications, including:

- SOC 1 (Audit Controls)
- SOC 2 (Security, Availability & Confidentiality)
- SOC 3 (General Controls)
- ISO 9001 (Global Quality Standards)
- ISO 27018 (Personal Data Protection)
- HIPPA (Protected Health Information)
- FIPS (Government Security Standards)
- FedRAMP (Government Data Standards)
- FISMA (Federal Information Security Management)

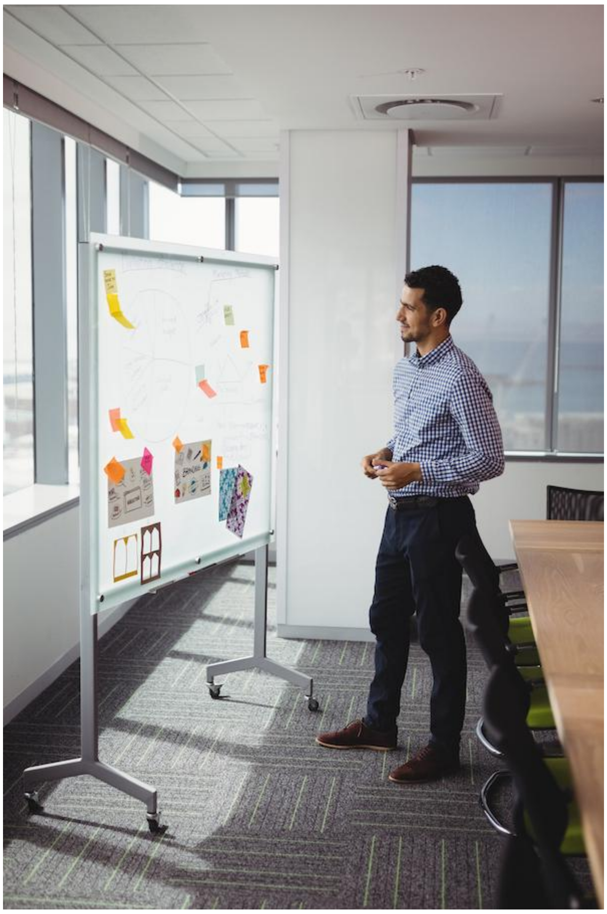

3.2. Risk Management

The process of identifying, assessing, and managing risks is a critical component of InfraVision's internal control system. The purpose of Infravision's risk assessment process is to identify, assess, and manage risks that affect the organization’s ability to achieve its objectives. The management of InfraVision also monitors controls to consider whether they are operating as intended, and whether they are modified as appropriate for changes in conditions or risks facing the organization.

Ongoing monitoring procedures are built into the normal recurring activities of InfraVision and include regular management and supervisory activities.

Responsibility for identifying risks to the entity and monitoring the operation of internal controls is shared among InfraVision's CEO, CISO, Compliance Manager and the DPO.

3.2.1. Risk Management Plan

The purpose of this plan is to identify the processes, controls and assessments that are in place to identify and mitigate risks and vulnerabilities.
This plan applies to all staff that collect, process or otherwise handle personal information of InfraVision staff or customers and/or have access to the integration infrastructure.

3.2.2. Control Self-Assessment (CSA) meetings

These meetings are organized to identify and manage risks. A CSA meeting is held every 6 months but can also be scheduled ad hoc if requested by an employee or if specific changes occur such as a change in leadership for example.

3.3. Privacy and security by design

Security starts during the design phase of an integration. We always include security considerations into our designs. Our template for functional designs contains a separate section dedicated to security aspects of the design. Our templates for projects and changes contain dedicated security requirements specification and review tasks to ensure that to minimize the security risks of our integrations. All these steps are fully traceable and auditable for review and improvement purposes.3.4. Awareness and training
The InfraVision Staff Training Procedure dictates the training that each new staff member is required to complete, as well as the required periodic training of existing staff members. The progress that a staff member is making is tracked by the staff member's manager during the yearly performance review.

When developing integrations, test data is either randomly generated by InfraVision or provided by the customer. We never use production data for testing purposes.

Our labour contracts contain a dedicated section related to our information security policies and the importance of adhering to these policies. Bij signing the contract our employees commit to comply with these policies.

3.4. Incident handling and communication

We have a predefined, templated workflow in our Incident Management system that includes tasks for emergency responses to security incidents. This ensures that all steps related to minimizing the impact and resolving the issue are recorded and can be monitored and managed. Customers are always informed in case of security incidents when their data was in danger or was compromised. We also report (potential) data leaks to the appropriate authorities.

3.5. Physical security

All our integration infrastructure is located in the Frankfurt based data centers of AWS. We do not have physical access to those locations, nor can we provide access to third parties.

Our office is physically secured with an alarm and entrance is only possible by using a personal tag. Every entrance to the office is logged in our security system. We do not allow visitors into our office space. Visitors only have access to our meeting room, located on a different floor on our office building.

We do not store any customer information digitally in our office. All our back office systems are cloud based. Integration-specific data is only stored and processed in our AWS infrastructure.

4. Technical security measures

We have several technical measures in place to maximize information security. In this chapter we will describe them in an appropriate level of detail for the scope and purpose of this document.

4.1. Infrastructure as code

The infrastructure that we use for running integrations is provided and managed by Amazon Web Services, based on our specifications. All integration instances are based on software images which are used to launch a containerized environment for each instance. This means that for each instance, integration code, operating system, software and data are completely segregated.

When an integration processes data, this data is retrieved from an external system using SFTP, processed and stored in the target system or data source. When the processing is completed, the container, including all data that remained, is deleted.

Containers are volatile and can be deleted or renewed any time a new task starts or when the integration code is updated by InfraVision. This ensures that (traces of) customer data remains on any of our systems.

The 'Infrastructure as Code' principle is not only a very efficient method for deploying infrastructure, it also minimizes the need for manual configuration and possible security risks caused by configuration errors or inconsistencies.

4.2. Continuous integration

All integration code for a customer is stored in a separate Git repository. Every time new code is stored (committed), an automated process performs unit tests of the updated code and generates a new software image that is automatically deployed to AWS. As soon as the upload is completed, all containers using this image are automatically decommissioned and relaunched using the new image.
This means that there is no need to manually deploy new code. There is no way (or need) to log on to one of our production servers. The software images do not contain any customer data, only software and integration code.

The relaunch of running containers does not cause any downtime. The new container is launched in parallel with the existing one. The load balancer stops directing new traffic to the 'old' container and waits for all current connections to expire before the container is stopped and deleted.

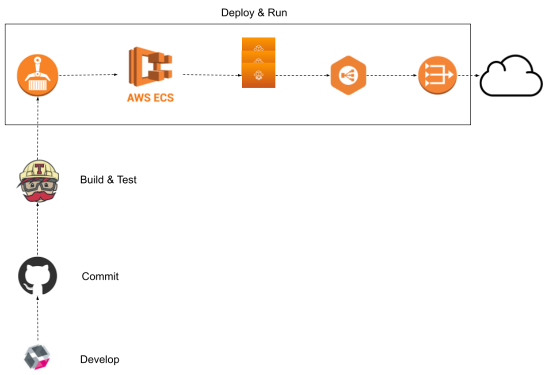

4.3. Monitoring and event correlation

Due to the volatile nature of containerized software we do not send log information to log files. Instead we send logs via an encrypted connection to a central monitoring system. This system generates alerts when errors appear in log files. A centralized log system also enables us to quickly check if exceptions also appear in other systems and makes it easy to spot trends across integrations.

Sensitive data is never included in logs.

Availability and response times are continuously monitored. In case of an outage our support center is notified immediately.

Exceptions in integration code are caught by exception handling routines in the code itself. When exceptions occur, an incident is automatically created in our support system to ensure traceability and resolution within the agreed timeframe.

4.4. Security patching and updates

The infrastructure that is running our integrations is kept up to date by AWS. Software patches that are relevant for our container images are automatically included every time a new image is generated. In case of emergency patches we can easily trigger a new build for all images that are in use.

4.5. Minimal Data principle and SFTP

Source data is preferably retrieved from a customer system at runtime using SFTP. When this is not possible we provide a SFTP server that can be used to store data. This server is only accessible from whitelisted IP addresses and only via the SFTP protocol. We use key based authentication, using 2048 bits keys.
For each integration we determine the minimal required data to run the integration and ensure that only this data is included in the source files. Together with a customer we define a data retention policy. Data older than the maximum age defined in the policy is automatically deleted from the SFTP server.

Customer data is never stored on any other system that our SFTP server.

4.6. Data encryption in transit

Data in transit is encrypted using TLS V2.0 when communicating with external systems. SFTP transactions are encrypted using RSA with SHA-256 hash.

4.7. Availability and Regions

AWS is providing data centers in several regions all over the world. Every region has several availability zones: separate data centers in the same region. Our integrations are running in the EU-Central region of AWS which is located in Frankfurt, Germany. Critical infrastructure is replicated across all 3 availability zones in the Frankfurt region. In case of an outage, containers are automatically redeployed to a different availability zone.

Your data will never leave our infrastructure and will therefore never leave the European Union.

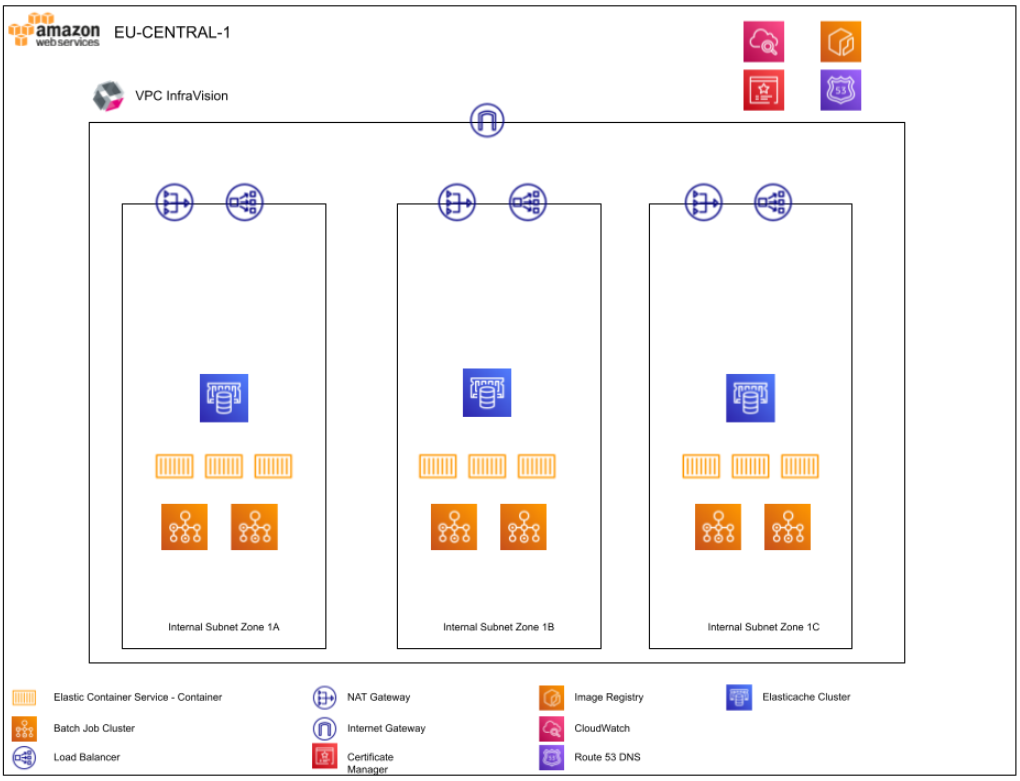

4.8. Notifications and email

Integrations can send notifications by email. We use the Amazon Simple Email Service for outbound notifications. Amazon uses a strict policy for acceptable bounce rates and spam reports. Our DNS records include SPF, DKIM and DMARC records to ensure email is delivered correctly.
All SMTP transactions are encrypted using TLS 1.2.

5. Conclusion

At InfraVision, we take the security of your data very seriously. We put a lot of effort in keeping our integrations secure and making sure that people in our company are constantly aware of the security aspect of all our activities.

If you have any questions about our security policies and efforts, do not hesitate to contact us.