IRIS AI
seeing beyond

Security of Iris AI Services

Preface

This document covers the security of Iris AI services. Currently Iris AI offers several distinct software products. While all Iris AI products share the same architecture and security measures, this document focuses on the services relating to the product Medusa Open-source Intelligence Tool (“Medusa”).

General

Subject

This security policy involves the security of Iris AI services. It consists of security objectives, guidelines for how to achieve them, overall security management strategy and implementation of policies on key security mechanisms in both every-day work and server security.

Scope

The security policy regulates interactions and relationships with the following subjects:

- partners
- customers
- subcontractors
- state agencies

It also covers the areas of server security, service development and practical security matters at Iris AI.

Goal of security policy

This security policy is the base for planning, design, execution and management of security at Iris AI. Its purpose and goal are to provide the means by which Iris AI can comply in security, actively develop its security practices and keep its customers’ data safe.

Security objectives

- Securing the confidentiality and integrity of the Customer’s data is the most important goal.
- Have the Customer data available at all times.
- Compliance to the security legislation (including copyright, personal information, state laws and regulations and workers’ health and safety requirements and fire safety requirements) must be ensured at all times.

All security measures must be economically justified and their disruptive effect to Iris AI operations and staff should be kept to a minimum.

Architecture and information structure

Customers' data

- Managing Customers' data is compliant with the General Data Protection Regulation (EU) 2016/679 ("the GDPR") and South Africa's data privacy legislation, notably the Protection of Personal Information Act, 3 of 2013 ("the POPI Act").
- The information is stored in database servers residing in South Africa. The data is used in South Africa.

Medusa

- Open source (internet resident) information relevant to Customer-specified concepts:
- Raw text of the associated data packet (such as Tweet, Facebook wall entry, RSS article etc)
- The place where the entry was made
- Time and date
- Additional metadata such as
- Co-ordinates where entry was made
- Author
- Classification fields relevant to the specific source
- User information (name, email, password)
- Basic information about customer usage profiles
- Any additional fields that the Customer may request via bespoke business cases.

The information is inputted to the system by the Customer, or by Iris AI on the behalf of the Customer, using for example .csv or .xls files. OSINT data is assimilated through a series of software services, called engines.

Master data resides at Iris AI SQL Server database servers in South Africa.

Access control principles

For each service Iris AI uses, only the absolute minimum number of users within the Iris AI organisation are granted admin level access rights (at the very minimum two people have access to given service).

- Only a few people in Iris AI organisation have access to the whole database. Most people have only partial access via user-profile controlled schema restrictions.
- Their computers' hard drives are encrypted, and password protected in order to avoid intrusion

if their laptops get stolen, using Steganos SAFE.

- By searching the logs, it is possible to track who has accessed a Customer's data and when this occurred.

Data isolation

Data is segregated using logical schemas, to avoid the cross-bleed of confidential data.

The database is accessed via a web service and does not allow direct links and cursors. Access to the web service is controlled by a security certificate. The web service resides on a separate physical machine.

Customers access the data through a browser client, which also consumed the web service. Access to the browser client is password protected.

Architecture

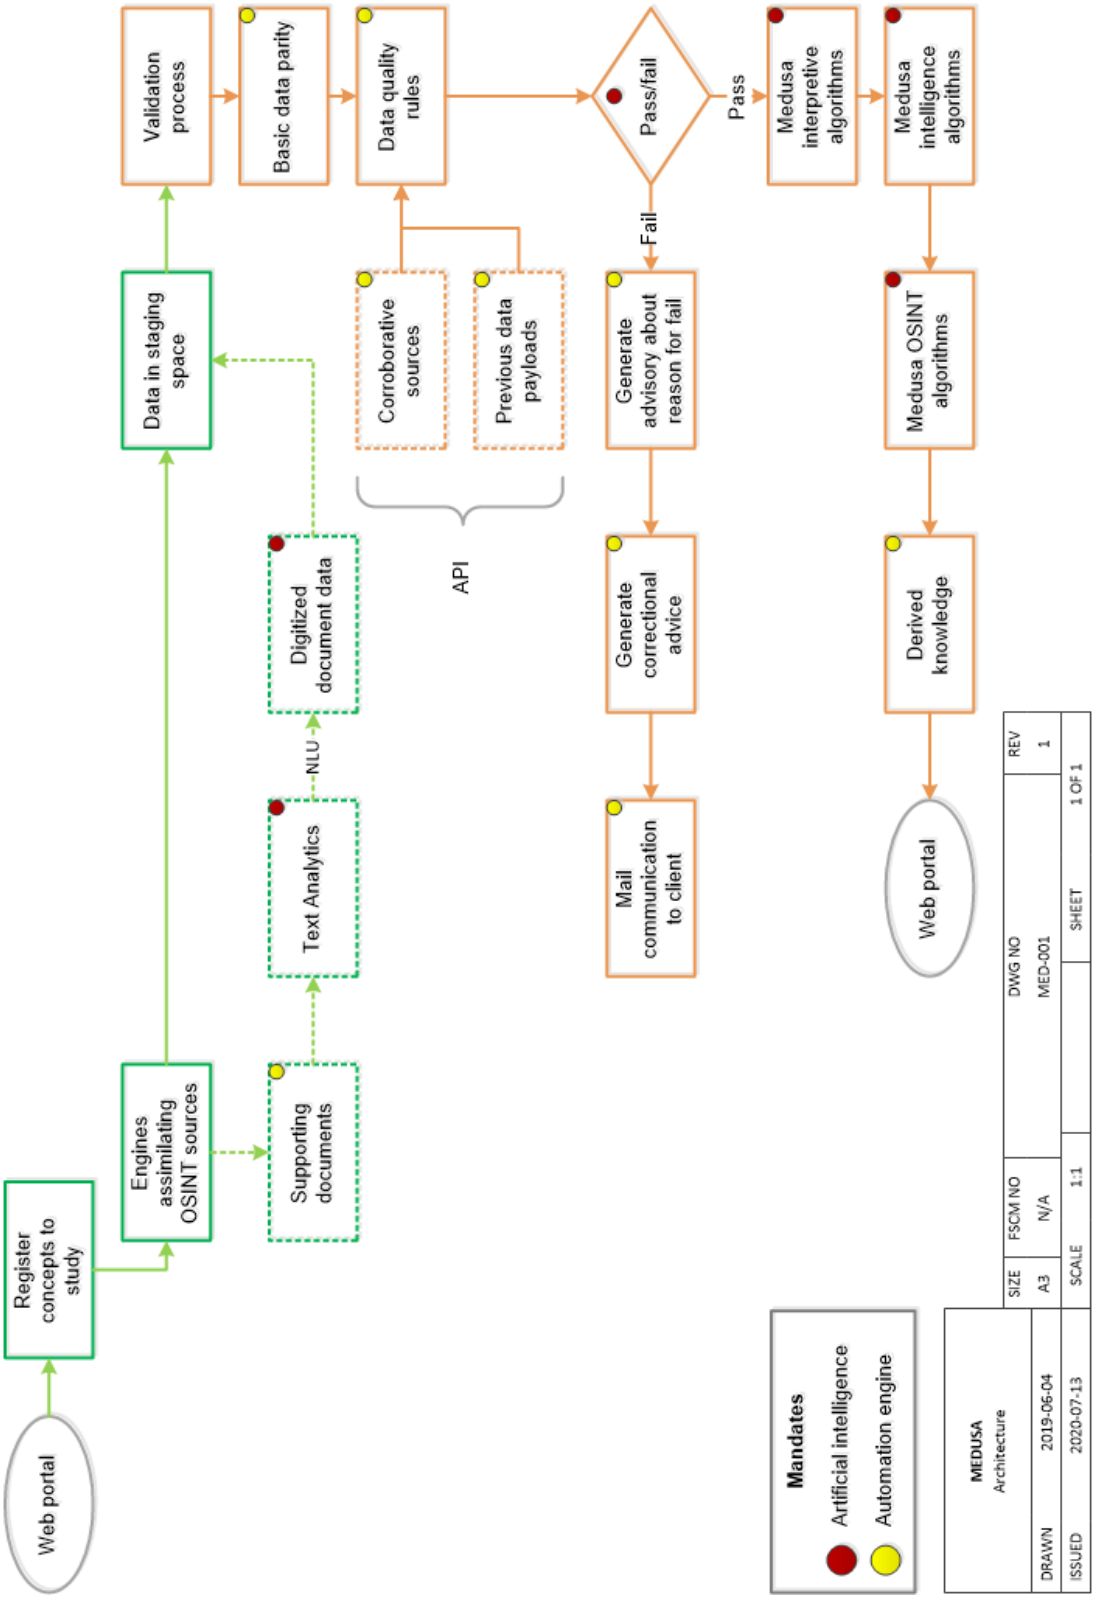

Application

The application is written using .NET Framework 4.6.1 programming languages and service frameworks.

The engines execute as instanced applications and consumes industry accepted data surfaces such as APIs or product protocols. Internet Information Services is used as the underlying HTTP server. The dominant sources are Twitter (api), Facebook (api), Instagram (api), Google (api), RSS feeds (RTSP), all rendering received data as JSON packets.

The engines are responsible for executing background tasks and use MSMSQ as a distributed task queue.

The application runs on Microsoft Windows Server 2019.

Test environments

All applications have separate testing environments. Testing environments do not contain production data. Test environments can be used to showcase upcoming features and to test new functionalities.

Resources

The application consumes the following resources:

- Internet Information Services1 is used to host the web services and web portal.
- SQL Server 20192 supports the Medusa databases.
- .NET Framework 4.6.13 is used to code all engines and structures.
- Python (VS 2019)4 is used to code some API interfaces.
- Power BI5 is used as a business intelligence tool.
- Crystal Reports6 is used to create any static reports required by customers.

[1] https://www.iis.net/
[2] https://www.microsoft.com/en-us/sql-server/sql-server-2019
[3] https://visualstudio.microsoft.com/
[4] http://www.python.org/
[5] https://powerbi.microsoft.com/en-us/
[6] https://www.crystalreports.com/

Information flow

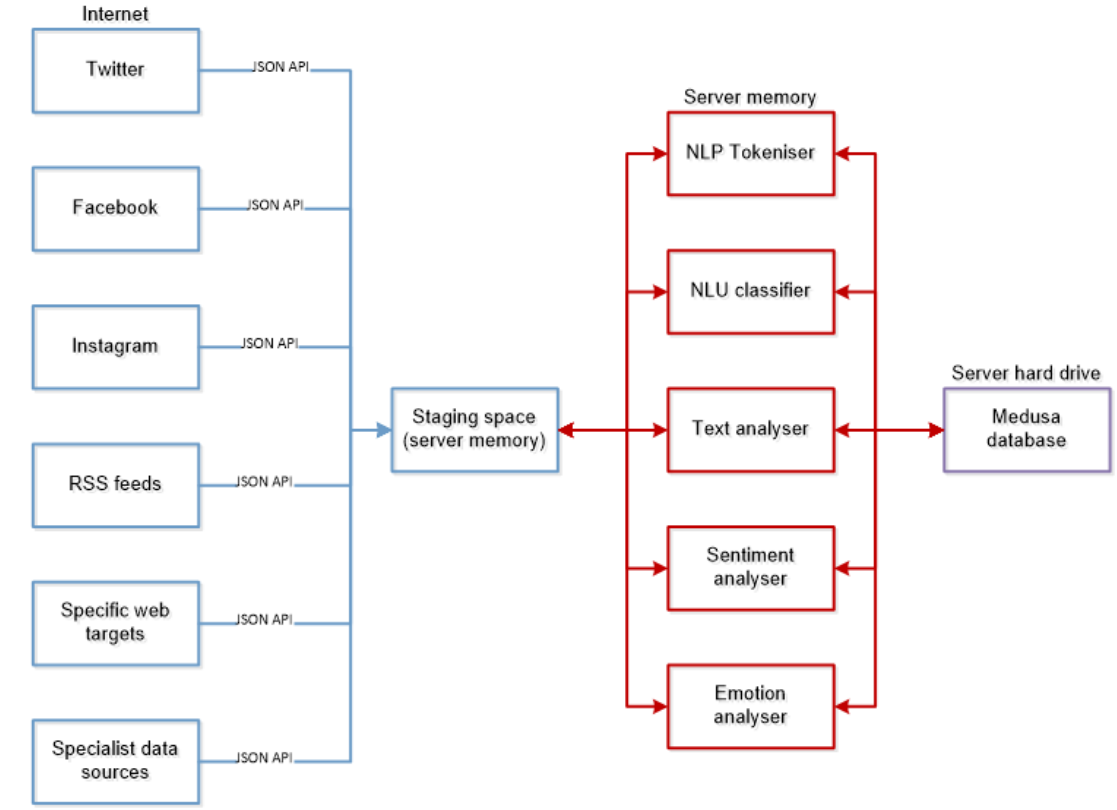

MEDUSA Information flow		SIZE	FSCM NO	DWG NO	REV
DRAWN	2019-06-04	A3	N/A	MED-002	1
ISSUED	2020-07-13	SCALE 1:1		SHEET 2 OF 2	

High level security considerations

- High network security with firewalls, DDoS mitigation techniques, spoofing and sniffing protections as well as port scanning
- The Service is fully isolated from other applications in the Iris AI data center.
- TLS 1.3 Encryption
- Access-controlled database with schema-based segregation
- Physical security ensured by CETA-certified data centres built at the Iris AI offices. A full security complement of digital and physical security measures protects the campus itself.
- Iris AI keeps the Customer up to date with any significant changes in the information security.

Service level agreement

Iris AI Service Level Agreement (SLA) characteristics are described below. Separate agreements may apply.

Server uptime

Servers should be up 99,7% of the time (~ 2h 9min per month). The calculation is done on a monthly basis.

Maintenance breaks

All maintenance breaks which take longer than 10 minutes are communicated to the customer admins at least 1 week in advance.

Support ticket response time

Iris AI will respond to service-related incidents and/or requests submitted by the Customer within the following time frames:

0 -8 hours (during business hours) for issues classified as High priority.

Support ticket resolution time

With any server issues we are constrained by the resolution time by Hetzner. Any critical bugs in the service are immediately assigned to a developer and if fixing them takes more than 24 hours the Customer is notified of this progress daily.

Support service availability

Coverage parameters specific to the service(s) are as follows:

- Telephone support: Monday to Sunday, 24 hours a day, 7 days a week
- Second tier support: On-site technical response team
- Email support: Monday to Friday, 08:00 to 17:00. SAST
- Emails received outside of office hours will be collected, however no action can be guaranteed until the next working day.

Data and privacy organisation of Iris AI

The employees of Iris AI have specific roles that limit their access to business-critical information and personal data. Access management controls have been implemented according to the documented roles and responsibilities. The roles are listed below.

Sales & Support

Responsibilities:

- Customer acquisition
- Customer onboarding and setup
- Customer support (can also make changes according to customers' requests)
- Transfer bugs and technical support requests to developers

Access rights in Iris AI's products:

- Can see all organisations and edit their settings
- Can see and edit all the users and their access rights – all actions are logged
- Can't see customer's observations or audits
- Customer's observations and audits might be visible to the employee during onboarding, when the employee is part of the organisation in Iris AI's product.

Access to customer information outside Iris AI's products:

- *Can have access to CRM and other marketing systems*
    - These systems can contain personal data of prospects and other customer contact personnel.

Developers

Responsibilities:

- Product development (planning, coding, bug fixing & code reviews)
- Help support & sales with technical support requests

Access rights in Iris AI's products:

- Developers do not have access to production's servers and services.

Access outside Iris AI's products:

- Access to full application source code
- Access to error tracking system
    - The system can contain personal data that is used to identify who encountered the error
- Access to the staging environments of Iris AI products. The staging environments are completely isolated from production and do not contain any customer data.
- No access to production database, or backups
- No access to make code deployments to production environment
- No access to production application log systems

Core developers

Responsibilities:

In addition to the responsibilities listed under Developers:

- Manage code deployments
- Manage and support custom integrations
- Manage third party services used by the products which contain sensitive information (email sending, etc.)

Access rights in Iris AI's products:

- Same as developers

Access outside Iris AI's products:

In addition to the access listed under Developers:

- Full access to production databases
- Full access to database backups
- Full access to production application log systems
- Full access to third party services used by the service

Members:

Due to the access to business-critical and personal data, the Core Developer-role is granted to a limited group of people.

Marketing

Access rights in Iris AI's products:

- None

Access to customer information outside Iris AI's products:

- Can have access to CRM and other marketing systems
    - These systems can contain personal data of prospects and other customer contact personnel

Data and privacy matrix

	Sales & Support	Developers	Core Developers
View & edit organisation settings	✓	✓	✓
View & edit user access rights	✓	✓	✓
View observations or audits	✓*		
Source code		✓	✓
Error tracking		✓	✓
Application databases			✓
Application deployments			✓
Application log system			✓

** Can view during the onboarding phase.*

Iris AI securityorganisation

Name	Role	Responsibilities
Gerhard Furter	Senior Engineer	Makes sure the development processes and software architecture adhere to the principles described within this document.
Quintin Smith	Core developer	Makes sure the development processes and software architecture adhere to the principles described within this document.

Physical security of Iris AI premises

Access control

- The Iris AI premises is located within a secure office park. The server room is situated adjacent to the Iris AI premises. Access is restricted to the Iris AI premises, office park and server room.
- Employees gain access to the Iris AI premises via FPX10, a biometric fingerprint and RFID reader.[7] Access is restricted to business hours. Any work to be conducted outside business hours will be treated as an exception, and the restriction will be lifted for the particular period of work. Otherwise, the biometric reader denies access to the Iris AI employees after-hours.
- Access to the server room is similarly controlled via an FPX10 biometric fingerprint reader. Access is restricted to 4 employees. One can only gain access to the server room through the Iris AI premises. This adds an additional layer of security.
- All visitors to Iris AI premises and office park must be accompanied by Iris AI employees.
- Visitors are welcomed in a reception area and hosted in one of the company's boardrooms or auditoriums. These facilities are all located in the office park, in a building separate from the Iris AI premises.

Premises Protection & Physical Security

- The Iris AI premises is located within a secure office park. The office park is well-lit, has good visibility and is surrounded by a 2m security fence. Electric fencing is installed along the fence line and linked to an alarm. The alarm, in turn, is linked to a third party armed response

[7] https://activetrack.co.za/products/

company, National Security & Fire, formerly known as CHUBB Security. [8] CCTV is installed to monitor the fence line as an additional measure.

- Access is restricted to the Iris AI premises and office park itself.
- Access to and egress from the office park is controlled and recorded using a real-time access management system.
- Access to the Iris AI premises is denied outside of working hours.
- Security officers of Maxi Security[9] are posted at the office park entrance. A security officer also carries out patrols along the perimeter of the office park.
- The security officer carries out his/her patrol with an Active Track device:[10] a hand-held guard monitoring device which tracks the movements of the officer in real-time. It also has panic buttons and distress signals that can be triggered if necessary. The Active Track devices are manufactured in Poland by EBS SP. z o.o.[11] and owned by entity Active Track[12] which is fully compliant with local privacy laws.
- CCTV is installed throughout the office park, as well as in the Iris AI premises and server room.
- A machine vision solution is integrated into the camera systems. The machine vision solution is provided by Activeye.[13] Activeye is fully compliant with local privacy legislation.
- The CCTV footage is managed and responded to by Maxi Security. Maxi Security's 24-hour control room is fully compliant with the standards prescribed by the South African Intruder Detection Services Association (SAIDSA).[14] Maxi Security complies with South Africa's privacy and security legislation, and by-laws.
- Doors to the Iris AI premises are self-closing and self-locking.
- The office park and Iris AI have appointed National Security & Fire as their armed response provider.

Security alarm system

- The physical alarm system is activated and deactivated manually by employees.
- The physical alarm system is tested regularly.
- Alarms generated pursuant to the Activeye-integrated CCTV are monitored in real-time by the

[8] National Security & Fire https://national.co.za
[9] Maxi Phumelela Security (Pty) Ltd, a security services provider registered in accordance with the laws of South Africa with registration number 2001/001526/07. www.maxisecurity.co.za
[10] https://activetrack.co.za/products/
[11] EBS SP. Z O.O. https://www.ebssmart.com
[12] Active Track (Pty) Ltd, with registration number 2013/229226/07. www.activetrack.co.za
[13] Activeye Wireless (Pty) Ltd, a software development company incorporated in South Africa with registration number 2017/331604/07. www.activeye.co.za
[14] http://www.saidsa.co.za/Bylaw6.pdf

Maxi Security 24-hour control room.

- A security officer on site is instructed by the control room operators to inspect any area on site where an alarm may be triggered.
- The officer will also be instructed to attend a particular area on site in real-time if Activeye determines that an unusual event may be taking place.
- The office park and Iris AI have appointed National Security & Fire as their armed response provider. National Security & Fire responds *inter alia* to fence alarms.

Fire Protection & Suppression

- The Iris AI premises is fitted with smoke alarms. Fire extinguishers are fitted in the Iris AI premises, and are serviced regularly by SANAS Accredited East Rand Fire.[15]
- When it comes to the server room, it was designed in line with a professional fire risk assessment. Consequently, the following measures are in place:
    - A CO_2 Fire Detection & Suppression system has been installed, fitted with sirens and strobe lights which are activated in the event of an alarm condition.
    - Server room is kept clean and well-maintained, free of combustibles, as well as of dust and debris.
    - Two air conditioner units are installed which maintains a constant temperature of 18 degrees Celsius, keeping the equipment cool.
    - 24-hour monitoring of the server room by CCTV integrated with Activeye machine learning technology. An alarm will be triggered in the security control room if an abnormal behaviour is detected.
    - Internal and external inspections of fire protection measures. East Rand Fire routinely inspects our CO_2 Fire Detection & Suppression System.
    - Hands-on training of personnel, and regular fire drills.

Power

The Iris AI premises and server room electrical power systems are designed to be fully redundant and maintainable without impact to operations, 24 hours per day, seven days per week. Uninterruptible Power Supply (UPS) units provide back-up power in the event of an electrical failure for critical and essential loads in the facility. The server room uses generators to provide backup power for the entire facility.

[15] East Rand Fire https://www.eastrandfire.co.za

Practical security matters at Iris AI

Password policy

- All Iris AI employees are required to use a password manager. Iris AI provides access to Steganos to all employees.
- Secure master password (at least 14 characters) should be used for accessing the password manager.
- The password manager should be used for generating a secure password for all applications and services.
- If some app has a default password it should be changed as soon as possible.

Connection security

- All Iris AI employees should always Iris AI main network inside Iris AI premises.
- Outside Iris AI premises employees should always use private networks or roaming instead of using public networks.
- Within Iris AI premises and the office park, visitors must always use the visitor network. In other words, the password of Iris AI network must never be exposed to anyone outside of the company.

Device security

Iris AI uses asset management to enforce device security. The asset management also includes an up-to-date inventory of device serial numbers. All devices used for work matters must be enrolled in the appropriate asset management system.

Laptop asset management

The laptop device security and asset management is managed through Steganos. The security policies in Steganos ensure that:

- Full disk encryption is enabled
- Password is required, and is at least 12 characters
- The laptop is locked after 3 minutes of inactivity
- Iris AI can remotely wipe the laptop

In addition, it is required that:

- All computers should be protected with appropriate Firewall software. Firewall
  should always be updated to the newest version.
- Automatic software updates should always be installed as soon as possible.

Portable device asset management

Portable device security and asset management is managed through Silo. The asset management ensures that:

- Iris AI Silo Suite products cannot be accessed without accepting the security policies
- Device encryption is enabled
- Iris AI can remotely wipe the mobile device
- Device uses password - PIN code, or a fingerprint unlock

In addition, it is required that:

- All mobile devices must enter sleep mode after 30 seconds of leaving the device unattended.
- Automatic software updates should always be installed as soon as possible.

Securely transferring information

Each employee is expected to handle sensitive, and business critical data with utmost care. This includes any information that a Customer sends in either physical or electronic format, or information that a Customer has provided to given contact person verbally. The following principles should apply:

- Employees do not store Customer information in their device hard-drives no longer than necessary.
- Any physical information provided by the Customer should be disposed of when it is no longer needed.
- Customer information within the applications should not be accessed unless absolutely necessary. Specific rules for accessing can be negotiated with the Customer.
- All Customer information should be stored within a single identifiable folder and the contents of this folder should be regularly checked for information that can be disposed of.

- Information and files containing sensitive, or business critical data should be shared between employees only when absolutely necessary. The information should be rather imported to the appropriate system right away, so that the safeguards of the target system will be enforced. If the information or file needs to be shared to an employee, it should be done through a safe channel, such as:
- Google Drive
- Gmail confidential email
- WhatsApp
- Unnecessary printouts of sensitive documents should be avoided. If a paper printout is necessary, it must be archived in a safe, or shredded as soon as possible.

Copies of production data

Sometimes it is necessary to test a certain feature or a fix for a bug with production data. The following principles should apply for testing:

- Features / bug fixes should only be tested with production data if absolutely necessary.
- If the feature / bug fix must be tested with production data, it should preferably be tested in a temporary testing server environment. Testing servers have the same security controls as the production server.
- Only if absolutely necessary should the developers take local copies of production data to their own working stations. In this case the production data should be immediately removed from the working station once the testing has been completed.

Memory sticks and portable storage devices

The use of memory sticks and other removable media is discouraged. The use of such devices for storing personally identifiable information is prohibited.

Policy violations

If any Iris AI employee violates the documented security policies, the following actions should take place:

- Director, and relevant manager, of Iris AI investigate what policy was violated and how serious the violation was.
- Director and relevant manager take the necessary actions with other executives

of Iris AI.

Employee hiring and contract termination

- Each new employee gets introduced to the security matters at Iris AI.
- Each new employee that handles customer data signs a personal confidentiality and non-disclosure agreement.
- Customer can order background checks for Iris AI employees.
- Once an employee's contract is terminated, his or her access to all Iris AI services and applications is revoked.

Employee competency management

- Annual security training is mandatory for all employees
- Best practices are regularly and non-formally shared among teams
- Findings in external security audits are shared and discussed among the whole company

Risk management plan

Risk	Probability	Impact	Risk factor	Mitigation
Bankruptcy or other failure of main server provider	Very unlikely	Severe	Medium	Documented steps on how to install Iris AI applications to other cloud platforms.
Employee business continuity. For example, a key employee leaving Iris AI that has lots of tacit knowledge.	Possible	Moderate	Medium	Culture that endorses openness and sharing information. Making sure that in all business critical processes the bus factor is > 1.
The bankruptcy or other failure of other business critical SaaS providers.	Unlikely	Moderate	Medium	Iris AI has listed alternatives for all SaaS providers. Iris AI has changed many providers already while tendering their prices. Iris AI reviews yearly all its SaaS providers.

Hacked employee devices	Unlikely	Significant	Medium	Following industry best practices for securing employee devices. Using G Suite for centralized control of employee mobile devices.
Hacked servers	Very unlikely	Severe	Medium	Following best practices for password management. Choosing only the most trusted and best possible server providers.

Security by design

Software development processes

- Test driven development
- Automatic code coverage checks of unit tests
- Continuous integration
- Automatic code style checks
- Every change goes through extensive code review

Using open source libraries and components

Iris AI uses lots of open source in its applications. Each open source library and component gets reviewed before it is added to the application. The license of given library / component is also reviewed, so that it is compatible for commercial purposes.

Before updating any component, the changelog of given component is reviewed and checked.

Patch management

Microsoft Windows handles operating system updates and database security patches automatically.

Choosing SaaS partners

Iris AI uses only industry leading SaaS partners for providing its service. Each partner is

carefully hand-picked, with special attention given on security. When selecting SaaS partner Iris AI checks the following things:

- What alternatives are there?
- How does the potential SaaS partner handle security?
- Size, revenue, ownership structure (if possible)

Annual security training

Iris AI offers annual security training to its employees. Attendance on these training sessions is mandatory and tracked.

Annual security audits

Iris AI conducts annual internal security audits that cover various aspects from physical security of our office premises, to security of our everyday processes.

Service security testing

Testing of security architecture is automated with comprehensive unit tests and continuous integration (CI). The CI server runs the tests after each commit in codebase. The unit tests are also run locally by the developers. On a typical day this means dozens of test runs.

Also, each line of code gets code reviewed by at least one developer before going to the production server.

SAML SSO integration

Medusa offer SAML SSO integration for additional authentication security. This ensures users cannot login when they are removed from Customer's internal system (for example Microsoft Active Directory). Also, the policies used there (password complexity and password change interval) are automatically reflected to Medusa.

Malware scanning

As an additional security measure PB services support malware scanning of attachments via Malware Bytes. Customers may choose whether or not they want to activate this option.

Compliance: GDPR & POPI Act

Iris AI ensures that all its applications are compliant with the provisions of both the GDPR and POPI Act. Iris AI ensures that all the server infrastructure and component providers are also GDPR and POPI Act compliant.

Iris AI maintains up-to-date documentation and practices regarding GDPR and the POPI Act in a workshop held every 6 months.

Data classification

While article 30 of the GDPR does not apply to Iris AI, the company maintains a data inventory for purposes of completeness. The Iris AI data inventory contains:

- The name of each controller on behalf of which Iris AI is acting
- The categories of processing carried out on behalf of each controller
- Confirmation that personal data is transferred to Iris AI in South Africa, and the Iris AI documentation specifying the suitable safeguards implemented
- A general description of the technical and organisational security measures referred to in Article 32(1) of the GDPR.

Data retention

General data deletion policy

Customer can issue deletion of specific database record or ask Iris AI to delete some range of records and their associated log entries. The records are permanently destroyed after one year, when the database backups containing those records are destroyed.

Customer specific data deletion policies

Iris AI allows customers to set their own data deletion policies. These policies can be set on user, observation category and observation question basis. The following use cases are possible:

- Setting the references for a given user role to be deleted from the system after a given time interval.
- Setting attachments for given concepts to be deleted from the system after a given time interval.

- Setting answers for given concepts to be deleted from the system after a given time interval.

Data retention after end of contract

Iris AI will retain the customer's data for 180 days after the end of contract. During this time period, the Customer can request a data handoff. After the 180 day time period Iris AI will remove the customer's data from all associated systems.

Data Protection Impact Assessments (DPIAs)

Iris AI assesses the need for a Data Protection Impact Assessment for a processing activity by referring to the GDPR Article 35:

> *Where a type of processing in particular using new technologies, and taking into account the nature, scope, context and purposes of the processing, is likely to result in a high risk to the rights and freedoms of natural persons, the controller shall, prior to the processing, carry out an assessment of the impact of the envisaged processing operations on the protection of personal data. A single assessment may address a set of similar processing operations that present similar high risks.*

The need for a DPIA is assessed when the following activities occur:

- Implementation of integrations
- Customer onboarding projects that are different in terms of personal information handling from normal onboarding projects
- New features that involve processing activities of personal information
- When deciding to take a new data source

Existing DPIAs

Based on our assessment, Iris AI services do not have any processing activities that would fulfill the description of the GDPR article 35.

IT-security

Hardening

Medusa is operated on Iris AI cloud application platform.

System Configuration:

System configuration and consistency is maintained through standard, up-to-date images, configuration management software, and by replacing systems with updated deployments. Systems are deployed using up-to-date images that are updated with configuration changes and security updates before deployment. Once deployed, existing systems are decommissioned and replaced with up-to-date systems.

Vulnerability Management

Our vulnerability management process is designed to remediate risks without customer interaction or impact. Iris AI is notified of vulnerabilities through internal and external assessments, system patch monitoring, and third party mailing lists and services. Each vulnerability is reviewed to determine if it is applicable to Iris AI's environment, ranked based on risk, and assigned to the appropriate team for resolution.

New systems are deployed with the latest updates, security fixes, and Iris AI configurations and existing systems are decommissioned as customers are migrated to the new instances. This process allows Iris AI to keep the environment up-to-date. Since customer applications run in isolated environments, they are unaffected by these core system updates.

To further mitigate risk, each component type is assigned to a unique network security group. These security groups are designed to only allow access to the ports and protocols required for the specific component type. For example, user applications running within an isolated dyno are denied access to the Iris AI management infrastructure as each is within its own network security group and access is not allowed between the two.

Medusa Application Security

We undergo penetration tests, vulnerability assessments, and source code reviews to assess the security of our application, architecture, and implementation. Our third party security assessments cover all areas of our platform including testing for OWASP Top 10 web application vulnerabilities and customer application isolation. Iris AI works closely with external security assessors to review the security of the Iris AI platform and applications and apply best

practices.

Issues found in Medusa applications are risk ranked, prioritized, assigned to the responsible team for remediation, and Iris AI security team reviews each remediation plan to ensure proper resolution.

Network security

The application is operated on Iris AI's application platform. Iris AI ensures network security with firewalls, spoofing and sniffing protections, and by prohibiting port scanning.

The protection mechanisms are described below in more detail.

Firewalls

Firewalls are utilized to restrict access to systems from external networks and between systems internally. By default, all access is denied and only explicitly allowed ports and protocols are allowed based on business need. Each system is assigned to a firewall security group based on the system's function. Security groups restrict access to only the ports and protocols required for a system's specific function to mitigate risk.

Host-based firewalls restrict customer applications from establishing localhost connections over the loopback network interface to further isolate customer applications. Host-based firewalls also provide the ability to further limit inbound and outbound connections as needed.

Spoofing and Sniffing Protections

Managed firewalls prevent IP, MAC, and ARP spoofing on the network and between virtual hosts to ensure spoofing is not possible. Packet sniffing is prevented by infrastructure including the hypervisor which will not deliver traffic to an interface which it is not addressed to. Iris AI utilizes application isolation, operating system restrictions, and encrypted connections to further ensure risk is mitigated at all levels.

Port Scanning

Port scanning is prohibited and every reported instance is investigated by our infrastructure provider. When port scans are detected, they are stopped and access is blocked.

Secure connections

All traffic between Customer and the cloud service is always transmitted over HTTPS. All connections are secured by either HTTPS or SSH.

Cryptography

The application uses PBKDF2-HMAC-SHA512 hashing algorithm for storing the hashes of users' passwords and authentication tokens.

The cloud service does not encrypt any data with public-key or symmetric-key algorithms.

SSL versions lower than 3.0 are prevented from accessing the service. At the moment the service supports encryption protocols TLS 1.2, TLS 1.1 and TLS 1.0. SSLv3 and SSLv2 are all supported. TLS 1.0 support will be removed by the end of 2020.

Data loss prevention

From Iris AI's security documentation:

> Continuous Protection keeps data safe on SQL Server 2019. Every change to your data is written to write-ahead logs, which are shipped to multi-datacenter, high-durability storage. In the unlikely event of unrecoverable hardware failure, these logs can be automatically 'replayed' to recover the database to within seconds of its last known state. We also provide you with the ability to backup your database to meet your own backup and data retention requirements.]

Logging

Iris AI uses IIS Logs for searching and archiving logs. All HTTP requests, SQL queries taking longer than expected and executed worker tasks are logged.

The logged data for HTTP requests may include:

- user IP address
- time when the request was made
- execution duration
- size of the request content in bytes
- URL of the requested page

- HTTP status code of the response

For all HTTP POST requests an additional log entry is created that may contain:

- user ID (if user is logged in)
- all POST data excluding passwords and other sensitive information

Each database query taking longer than expected is logged with the following information:

- time when the query was executed
- query execution duration
- SQL statement

All errors are logged using Windows Event Logging. Each logged error contains the following information:

- time when the error occurred
- full error stack trace
- HTTP request details
- user who made the request (if any)
- IP address where the HTTP request originated from

the logs are archived for 1 year.

Windows Events error logging service does correlation analysis on logged error messages. It groups similar errors together and prioritizes the errors based on their occurrence.

All user activity which affects database is also stored in separate database version history. The version history can be used for checking which user was responsible for a certain change.

Log inspection and alerts

Iris AI has several alerts set up for our logs for certain unwanted scenarios. In some cases, these alerts also execute mitigating actions (for example, if a user has too many failed login attempts, the system will temporarily block succeeding attempts).

Availability

DDos attacks

The following quotation from Iris AI Security documentation describes how to mitigate DDoS

attacks:

> Our infrastructure provides DDoS mitigation techniques including TCP Syn cookies and connection rate limiting in addition to maintaining multiple backbone connections and internal bandwidth capacity that exceeds the Internet carrier supplied bandwidth. We work closely with our providers to quickly respond to events and enable advanced DDoS mitigation controls when needed.

In addition, Slowloris based DDoS attacks are mitigated with strict timeouts to keep connections open no longer than necessary and to reduce resource usage.

Disaster recovery plan (DRP)

> Our platform automatically restores customer applications and Medusa SQL databases in the case of an outage. The Iris AI platform is designed to dynamically deploy applications within the Iris AI cloud, monitor for failures, and recover failed platform components including customer applications and databases.

Backups

Automatic backups of the SQL databases used in Medusa are taken daily, weekly, and monthly. Daily backups are archived for one week, whereas weekly and monthly backups are archived for one month. Database backups will always be taken before deploying a new major version of the cloud service.

In addition to database backups, the cloud service keeps a version history for all successful database transactions.

Business continuity

Data centre failures

From Iris AI security documentation:

> The Iris AI platform is designed for stability, scaling, and inherently mitigates common issues that lead to outages while maintaining recovery capabilities. Our platform maintains redundancy to prevent single points of failure, is able to replace failed components, and utilizes

multiple data centers designed for resiliency. In the case of an outage, the platform is deployed across multiple data centers using current system images and data is restored from backups. Iris AI reviews platform issues to understand the root cause, impact to customers, and improve the platform and processes.

Bankruptcy

It is possible to copy Customer owned information from the cloud database to Customer managed storage. This can be done manually by downloading a database dump of Customer's own data.

Physical security

The Hetzner facilities the cloud service uses adhere to ISO27001 and FISMA certificates.

Customer Applications and Databases

Our platform automatically restores customer applications and SQL Server databases in the case of an outage. The Medusa platform is designed to dynamically deploy applications within the Iris AI data center, monitor for failures, and recover failed platform components including customer applications and databases.

Iris AI Platform

The Iris AI platform is designed for stability, scaling, and inherently mitigates common issues that lead to outages while maintaining recovery capabilities. Our platform maintains redundancy to prevent single points of failure, is able to replace failed components, and utilizes multiple data centers designed for resiliency. In the case of an outage, the platform is deployed across multiple data centers using current system images and data is restored from backups. Iris AI reviews platform issues to understand the root cause, impact to customers, and improve the platform and processes.

Integrations

Iris AI offers various ways for integrating 3rd party software to its products.

SAML / ADFS integration

SAML / ADFS SSO integration is available for all of our products.

Data warehouse for BI tool integration

Iris AI may provide customers with credentials to a separate data warehouse database. The data in the warehouse database is updated once per day, and it can be analysed in the customer's own BI tool.

REST JSON API

All products have REST APIs. All the APIs are for our internal use only unless otherwise explicitly said, and they may change without any notice.

PB Gateway

PB Gateway is a separate service for updating user access rights based on user roles within a customer's own HR system. For large organisations PB Gateway can be essential for keeping access rights up-to-date.

Custom integrations

Iris AI has done various custom integrations in the past to different systems. These include integrations to property management systems, as well as SSO solutions.

Incident management

Customer notices serious security breach:

- The Customer contacts cyber security contact point either by email or by telephone.
- Iris AI incident management process.

Iris AI notices serious security breach:

- Iris AI informs the Customer.
- Iris AI incident management process.

Iris AI incident management process:

- Iris AI detects, records and classifies the incident.
- Iris AI gives initial support.
- Iris AI analyses the situation and either fixes the situation quickly or makes a proposal of possible mitigation procedures.
- Iris AI further analyses the incident and tries to think of possible scenarios of similar

kind and how these could be mitigated as well.

Iris AI documents all security breaches and unplanned service disruptions to Medusa.

References

- Iris AI https://irisai.co.za/
- National Security & Fire https://national.co.za
- Maxi Phumelela Security https://www.maxisecurity.co.za/
- Active Track products https://activetrack.co.za/products/
- EBS Poland https://www.ebssmart.com
- Activeye software development company http://activeye.co.za/
- SAIDSA by-laws http://www.saidsa.co.za/Bylaw6.pdf
- East Rand Fire https://www.eastrandfire.co.za/