PHILIPS

AI Manager

Datasheet

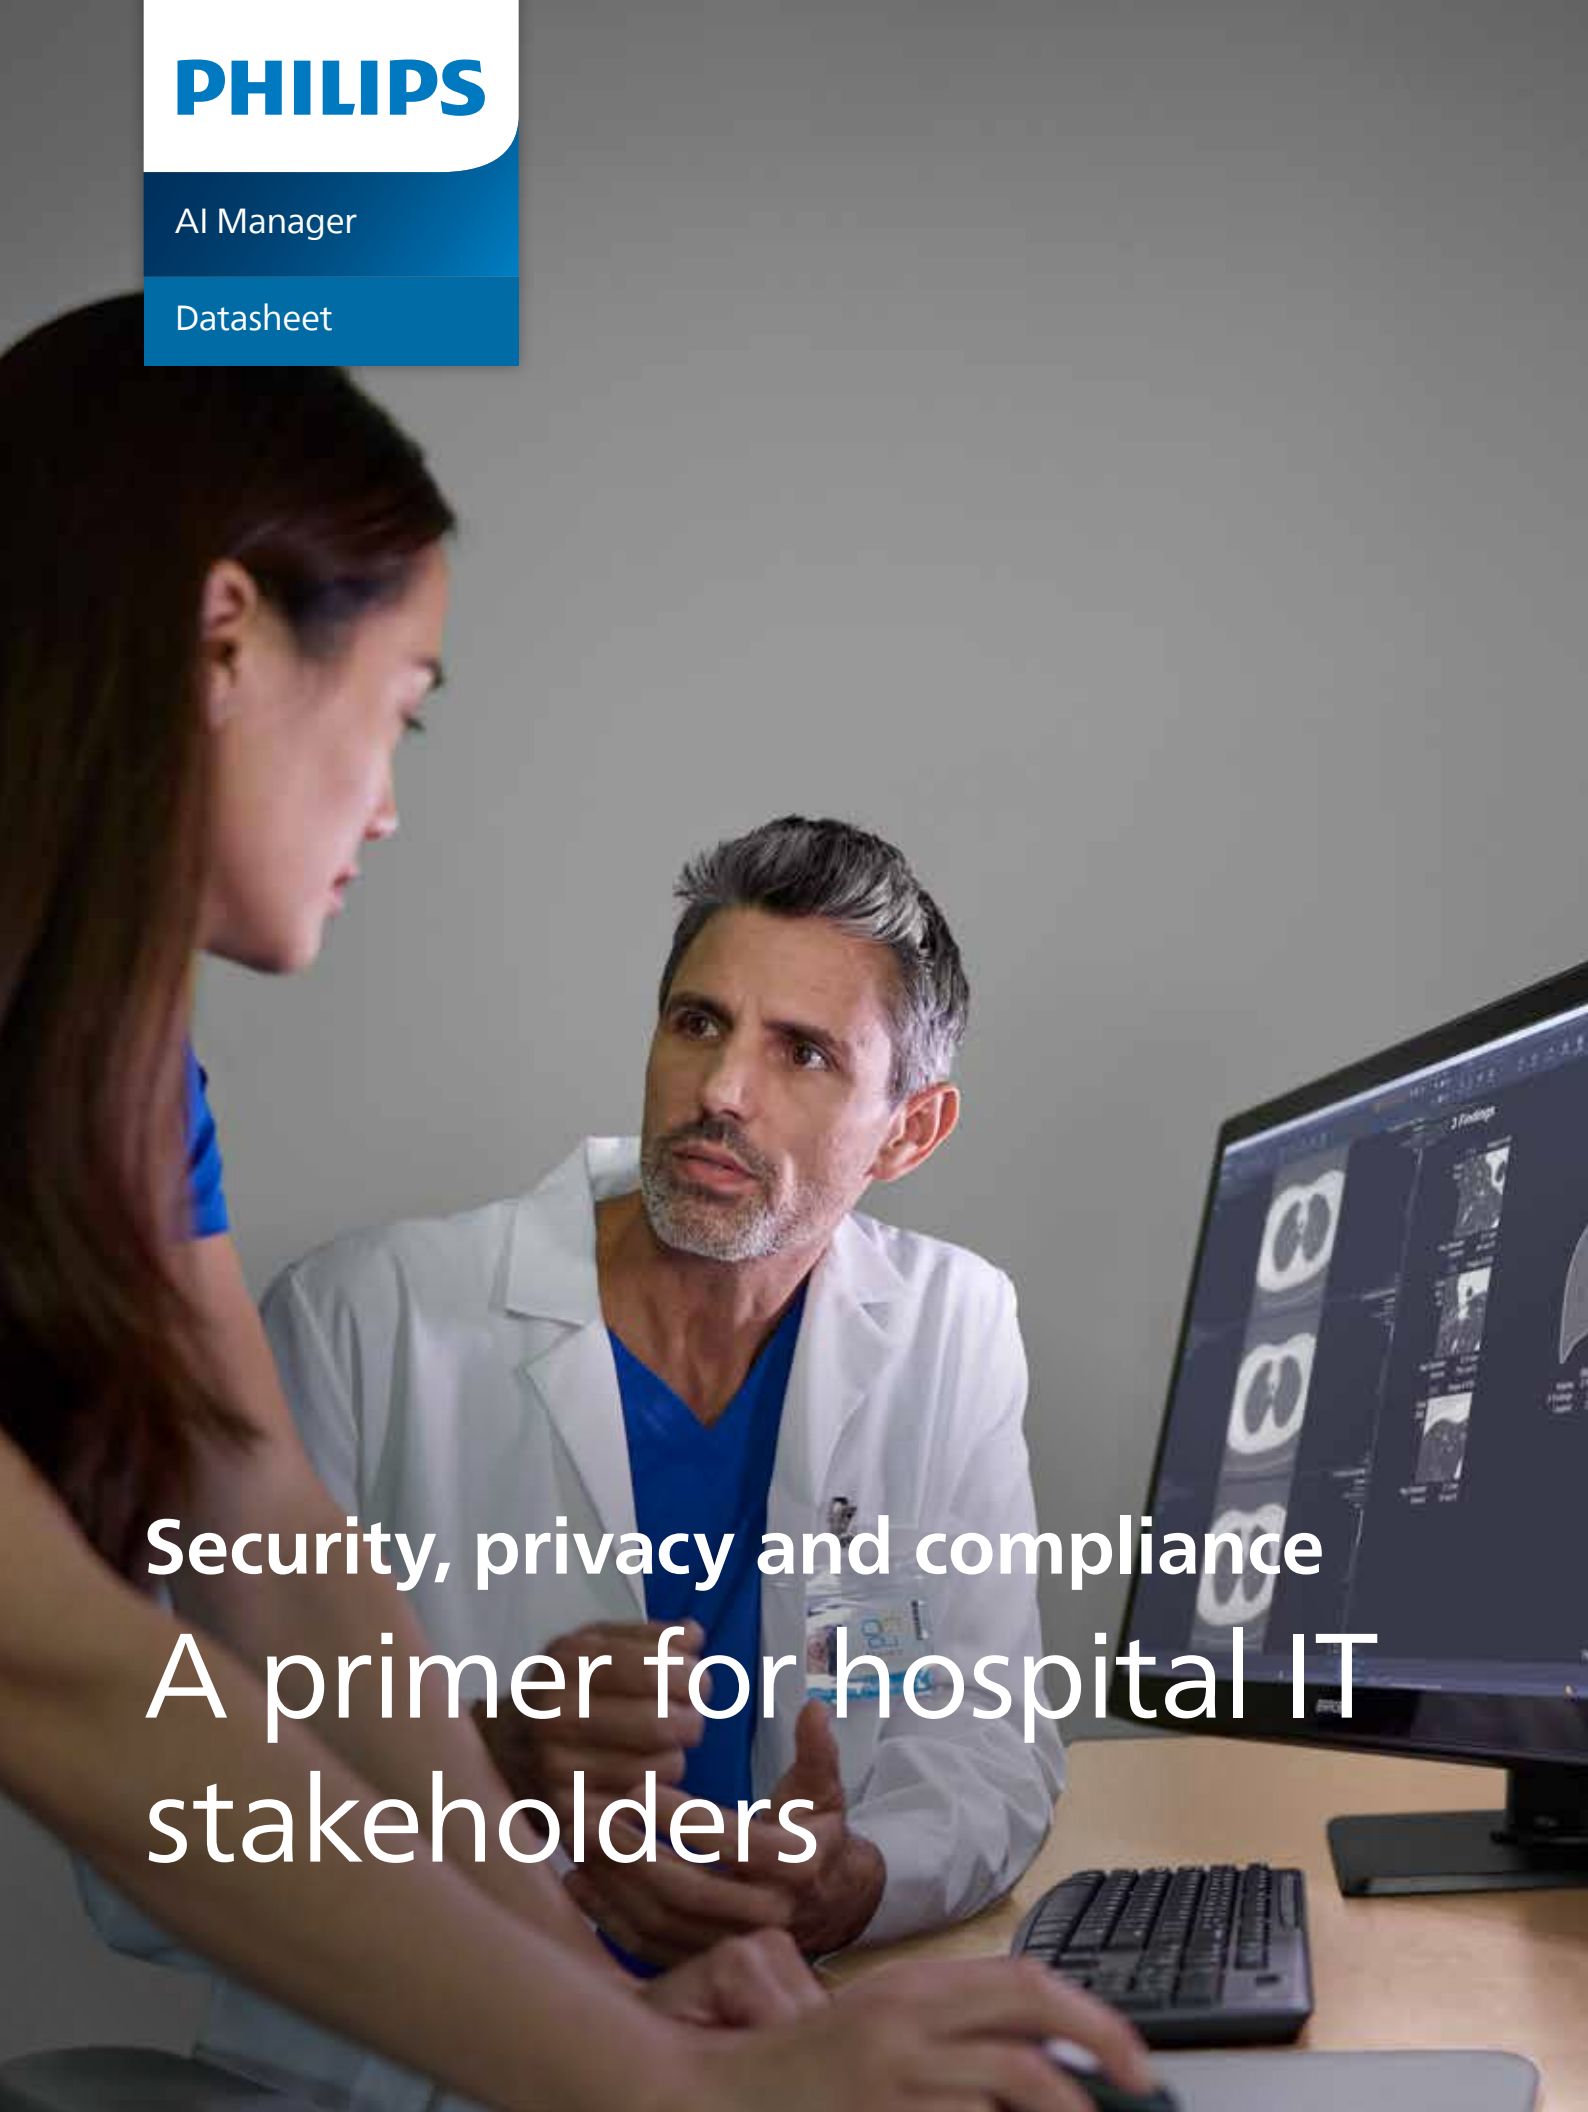

Security, privacy and compliance

A primer for hospital IT stakeholders

1. Introduction

Data governance must be a top priority for every healthcare organization when introducing a new application. Hybrid cloud solutions, such as the AI Manager, might add additional security and privacy considerations. This paper will address, high-level, how your sensitive data will be handled throughout the AI Manager.

Hybrid cloud architecture

Cloud computing can provide significant benefits for an AI solution, such as the flexibility to enable or disable AI applications without the need to adopt hardware requirements, or elasticity to handle unexpected peak-load at any time during operation. Alternatively, there might be also a desire to execute some of the AI applications on premises such as applications for urgency care or research driven in-house developed AI apps. Restrictive privacy rules might be another driver to prefer local executing of AI applications. AI Manager is a hybrid cloud solution supporting the execution of AI applications in the cloud or on premises, see flow below.

The Philips AI Manager is a platform installed and operated by Blackford[1] on behalf of Philips.
The platform itself runs on premises, mainly handling data routing, de-identification, and the AI result viewer. Only a dashboard for monitoring application execution is operated in the cloud.

The AI applications might run on premises or cloud-based, depending on the application and customer preference. Cloud-based applications often make use of large cloud providers like Amazon Web Services (AWS), Microsoft Azure or the Google Cloud Platform (GCP).

AI Manager – Functional flow

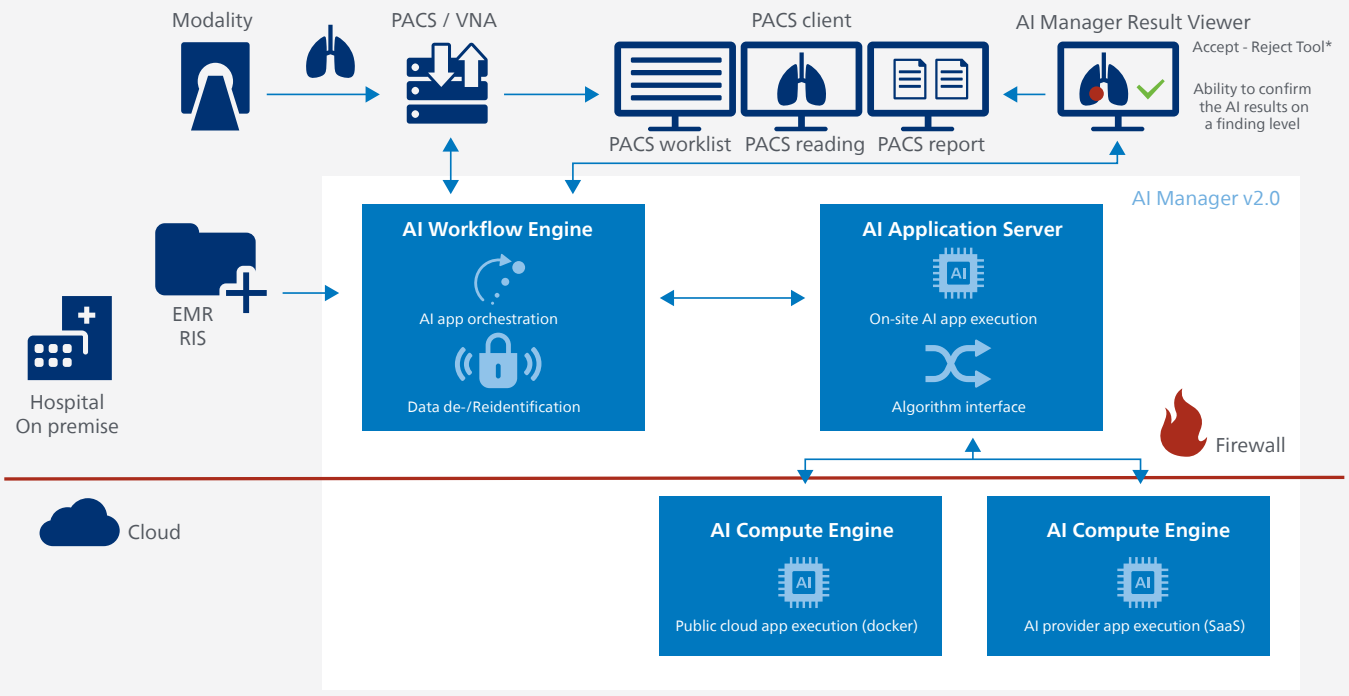

(*) Accept Reject Tool (Version 1.2) is a Class I Device in the EU (according to regulation (EU) 2017/745 Annex VIII classification rule 11).

2. Architecture, security, and privacy

The AI Manager solution includes the AI Manager platform as well as customer-selected applications from a large portfolio of third-party AI applications. The following section describes the architecture, with specific emphasis on security and privacy considerations.

On-prem platform

The AI Manager platform runs on two virtual machines (VMs) on premises. The Workflow Engine is a Windows server that handles data and application orchestration. DICOM data is de-identified here, before it is further processed. The Workflow Engine also hosts the optional AI result viewer (accept/reject tool). The Application Server is a Linux system running a Docker environment. This server can host applications on-site, and serves applications executed in the cloud via application-specific gateways. Communication between the platform servers (Workflow Engine and Application Server) is using DICOM C-Store through an SSH tunnel to ensure data integrity and confidentiality.

For AI Manager installation, maintenance, and service, remote access for Blackford service engineers is required.

AI Manager dashboard

The AI Manager dashboard is a cloud-based telemetry application, to monitor platform system resource usage and application execution. The dashboard is hosted in the Blackford cloud.

AI Manager v1.1 – Network architecture

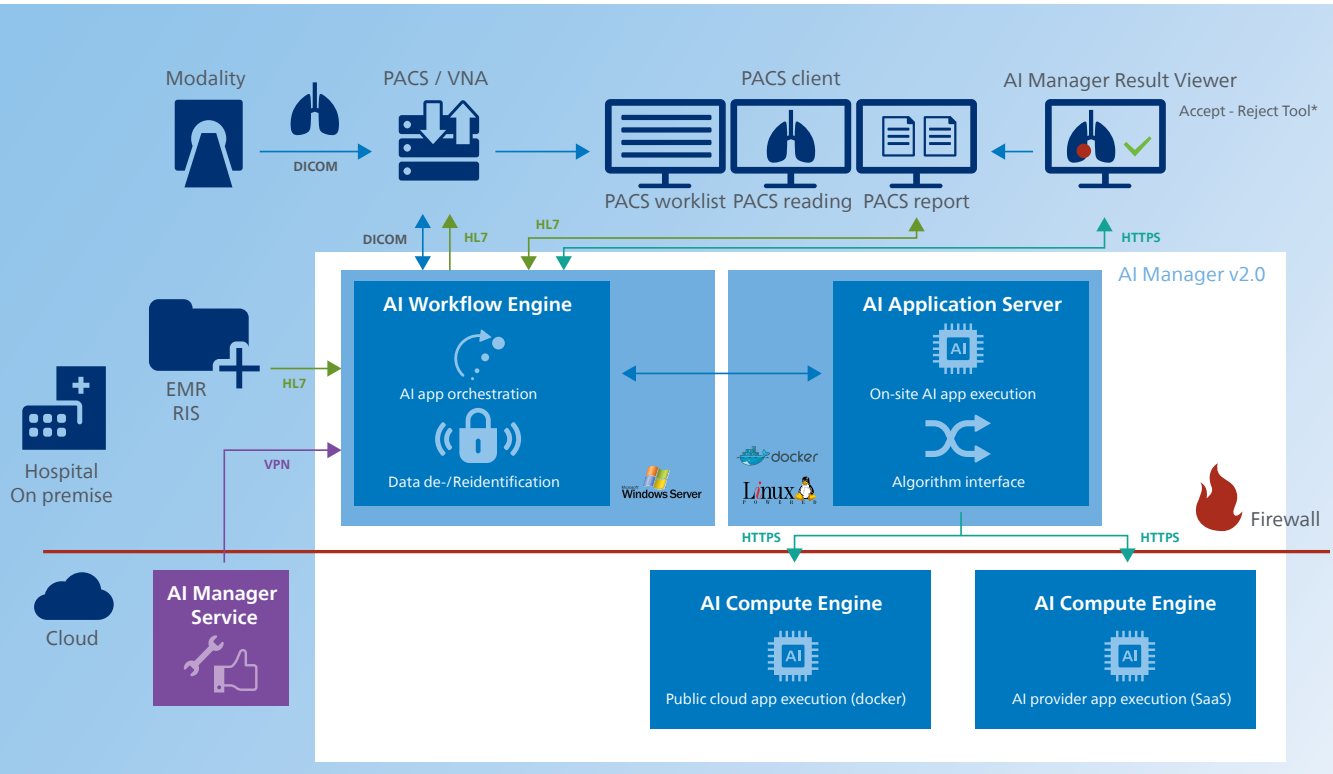

(*) Accept Reject Tool (Version 1.2) is a Class I Device in the EU (according to regulation (EU) 2017/745 Annex VIII classification rule 11).

Application execution

AI applications are executed from the Application Server, either as local Docker containers, or on the cloud using application-specific gateways. The gateways are outbound https connections and typically use REST APIs to exchange input and result data with the application. Cloud-based applications usually make use of large public cloud providers like Amazon Web Services (AWS), Microsoft Azure or the Google Cloud Platform (GCP). These also facilitate hosting of applications in a region suitable for the hospital. Alternatively, applications could be hosted in an existing hospital tenant of a public cloud.

Data flow

The AI Manager solution runs fully automatic in the background. The optional accept/reject tool is the only platform component visible to the end user.

The basic data flow is as follows (see diagrams above):

1. Relevant DICOM data is transferred to Workflow Engine
   a. Either using rule based DICOM forwarding from PACS
   b. Or DICOM query/retrieval from PACS by Workflow Engine, either based on HL7 message trigger or polling.
2. Workflow engine performs a more fine-grained match of data vs. requirements of given AI applications.
3. If data matches an application, the DICOM data is de-identified and sent to the application on the Application Server via C-STORE.
   a. For a cloud-based application, the gateway will transfer image data to the cloud.
4. DICOM result data is sent from application to Workflow Engine via C-STORE.
5. HL7 messages can be sent by Workflow Engine to PACS/RIS for workflow prioritization.
6. DICOM result data is sent from Workflow Engine to PACS (C-STORE).
   a. If accept/reject tool is configured for given application, DICOM result data will be held on Workflow Engine until accepted/rejected or until configured timeout.

Patient data

Patient data is considered as high-business-impact data as it contains fields such as the name of the patient, the pathologies, the referring physician, etc. The AI Manager is using multiple techniques to protect data from the hospital:

- Data minimization: Only the subset of data that is required for processing will be transmitted to the cloud. Other parts of the data are stripped or replaced by a unique code such that it cannot be read once it leaves the hospital datacenter. The AI Manager platform operates as a transient data transfer device and does not store data long term.
- Data de-identification: Directly identifiable data (e.g., name, birth date, …) is replaced by a unique code such that it cannot be read once it leaves the hospital datacenter.
  - For cloud-based applications, before sending data to the cloud, DICOM data will be de-identified on premises according to DICOM Standard Part 15 confidentiality profiles[2]. More details can be found below.
  - For the dashboard, DICOM StudyInstanceUID and accession number might be processed in the cloud by the AI Manager platform. These can be de-identified before transfer to the cloud.
- Encryption: data in transit between the hospital datacenter and the cloud will be transmitted through encrypted channels to avoid a man-in-the-middle attack. Stored data like DICOM UIDs for the dashboard can on request be stored encrypted to be fully unidentifiable.
- Authorization: Only authorized users can access data as required to deliver the AI Manager service. For the accept/reject tool, radiology users can authenticate via single-sign-on. The dashboard can be accessed by authorized AI Manager personnel. For remote service and maintenance, only authorized AI Manager personnel has access. All user accounts have 2-factor authentication with time-based one-time password.

3. Shared responsibility

The AI Manager solution must be considered as SaaS solution (Software as a Service). The operation of the VM's on premises is managed by the hospital. Underlying technologies and tasks for cloud operation, and the protection of the data in the cloud will be executed by Blackford and the individual application providers. This is also reflected in the contractual setup, see below

All involved in the operation and development of the AI Manager and the AI applications as well as the public cloud providers invest significantly in security technology and highly qualified IT-staff to deliver a service to the highest security standards. However, using a cloud solution is always a shared responsibility between the providers and the user. Some security, data governance and identity management tasks are the responsibility of the hospital to ensure correct and secure handling of the data throughout the AI Manager solution.

Contract Framework

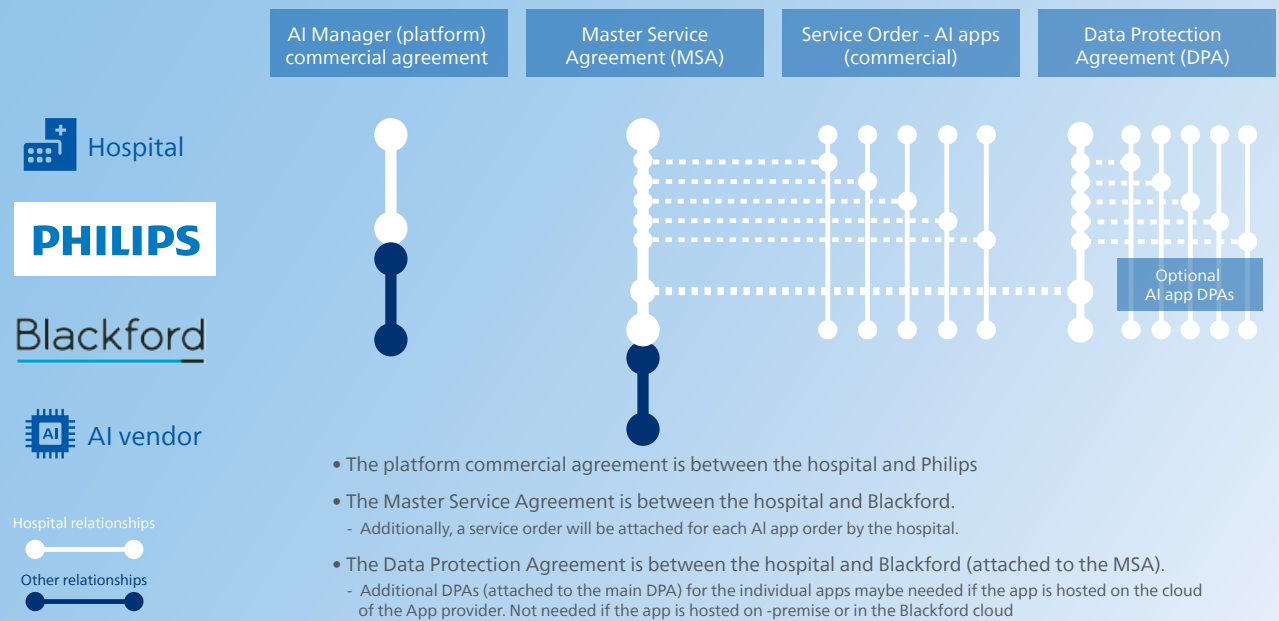

- The platform commercial agreement is between the hospital and Philips
- The Master Service Agreement is between the hospital and Blackford.
  - Additionally, a service order will be attached for each AI app order by the hospital.
- The Data Protection Agreement is between the hospital and Blackford (attached to the MSA).
  - Additional DPAs (attached to the main DPA) for the individual apps maybe needed if the app is hosted on the cloud of the App provider. Not needed if the app is hosted on -premise or in the Blackford cloud

3.1. Hospital roles

On premises security controls

The hospital is responsible for the hospital infrastructure, e.g. physical access of servers and network, the security of host machines running virtual servers and the operations executed on the hospital infrastructure (e.g., the virtual machine server patching, the firewall and network configuration, the security of the client machines accessing the services, anti-virus, etc.).

Identity management

Furthermore, the hospital is responsible for the management of the identities and linked authorizations for the staff who will have access to the AI Manager services. Roles and credentials must be defined for people with IT responsibility to manage the infrastructure of on-prem deployment, and remote access needs to be granted to Blackford service engineers for installation, configuration and service of the platform and applications on premises.

Privacy configuration

The hospital holds the responsibility for requesting customized configuration of de-identification which is a critical component to protect the patient data. For the dashboard, DICOM StudyInstanceUID and accession number might be processed in the cloud. These can be de-identified before transfer to the cloud by encryption (hashing).

For AI applications, DICOM attributes containing personal data are replaced by dummy values. By default, dates are shifted (DICOM Retain Longitudinal Temporal Information Modified Dates Option) and DICOM UID's are retained. Details can be found in the DICOM Conformance Statement of Blackford. Other de-identification profiles are available on request. Re-identification of AI application results is based on the executing workflow on premises, not on a mapping of DICOM attributes.

Depending on the AI application, attributes such as body part, age, gender, slice thickness, etc. must be enabled to allow correct operation of the AI algorithms. For images to be processed by cloud-based applications, it is the responsibility of the hospital not to include burned-in annotations (text elements written directly in the image) as this is also potentially sensitive patient data.

3.2. Philips roles

Commercial contact for the AI Manager platform

All commercial quoting and payment processing for the AI Manager platform itself will be handled by Philips. Note that Philips acts only as an agent for the AI applications. Quotes for the AI applications must come from Blackford who is the responsible party for the app sales and operations.

Project management

The local Philips organization can offer specific project management services in line with the complexity of the project to be deployed (e.g. in combination with a PACS upgrade).

Support

Philips serves as initial point of contact for customer service. Philips channels AI Manager service requests to Blackford, the delivery of service is performed by Blackford. Philips as a 'data processor' is limited to triaging customer issues.

3.3. Blackford roles

Commercial contact for applications
All commercial quoting and payment processing for the AI applications is handled by Blackford. Philips acts as an Agent for the Blackford AI ecosystem portfolio of AI applications.

Installation and support for platform and applications
Blackford is responsible for installation and technical maintenance of the AI Manager platform (development of patches or new releases potentially with improved security mechanisms). Further, Blackford is responsible for deploying, updating, and configuring on-prem components of AI applications. For these purposes, Blackford service engineers need remote access to the AI Manager platform.

Initial installation of the AI Manager platform by default is executed from the USA, with no access to patient data. Optionally, installation can be executed from other countries, if required.

While Philips serves as initial point of contact for customer service, the delivery of service is performed remotely by Blackford. Blackford also carries out customer training.

Secure configuration and cloud processing
Under GDPR[3], Blackford acts as a 'data processor' for platform and applications, acting on behalf of the hospital who remains the 'data controller'.

Blackford processes patient data in the cloud for telemetry purposes only. No image data is processed, only meta data like DICOM UID's or accession number. On request, this data can be stored encrypted to be fully unidentifiable. The telemetry dashboard is hosted on Google cloud and operated by Blackford.

Blackford undergoes penetration testing annually and is certified according to ISO/IEC 27001:2013, as well as ISO 13485:2016 and ISO 9001:2015.

3.4. AI application provider roles

AI applications that are available for the AI Manager platform are sold by Blackford (acting as principal), while Philips has the role of an agent.

Legal manufacturer and sub-processor
The third party AI application providers are responsible for the performance of the medical device AI applications itself. The provider of the application will identify itself as the legal manufacturer of the licensed AI application. For cloud-based AI applications, the application provider will act as data sub-processor to Blackford for the specific AI application. Necessary data privacy provisions for sub-processors are included in a data processing agreement between Blackford and the hospital.

Service and Support
In case certain abnormalities are found regarding algorithm performance, the hospital must resolve support or liability questions with Blackford as the primary contact. Blackford will support the hospital by triaging any issue to identify the root cause and hand further resolution to the hospital and the AI application provider if it is related to the third party medical device AI application. The legal manufacturer is responsible to act in accordance with all complaint handling regulatory requirements.

Cloud-based third party AI applications are typically hosted by large cloud providers like Google or Amazon Web Services. The public cloud provider is a sub-processor to the application provider and is responsible for the security of the cloud (compute, networking, availability zones, edge locations, …).

Third party AI application providers may support Blackford in providing customer training and support for the applications if needed.

4. Step-by-step guidance to get started with the AI Manager

4.1. Preparation

1. Select clinical use cases of interest based on the needs of the hospital
2. Identify AI applications from the portfolio to fulfill the needs
3. Check how these apps can be deployed: on premises, public cloud hosted in the hospital tenant, hosted as a service
4. Complete the Data Processing Agreement (DPA) for the selected applications in close cooperation with Blackford, who holds the responsibility of the AI applications.
5. Prepare remote access for Blackford deployment engineers.

4.2. Installation

1. Two servers (typically virtual machines) are needed in the local datacenter:
   - Workflow engine: A Windows-based machine to allow workflow and data orchestration, including data de-identification. A typical setup requires 10 cores, 24GB RAM and 500GB disk space. Final requirements will be determined based on the expected peak load.
   - Application server: A Linux (ubuntu preferred) based machine to execute AI applications. The sizing of this virtual machine depends on the selected AI apps, the related hosting option, the complexity of the AI app and the expected workload. The typical setup ranges from 4 cores, 4GB RAM, 32GB disk space (cloud-based algorithms only) to 16 cores, 64GB, 500GB disk space to also support local algorithm execution.
   - If a locally operated algorithm requires a GPU, the application server can be realized as a physical server. One workflow engine can operate multiple Application servers. The local algorithms may be distributed across these according to hardware and processing load needs.

2. Configure the firewall to allow:
   - Remote access to the VMs for the AI Manager setup.
   - Remote access to the VMs for maintenance, support and upgrade of the solution
   - Reporting to the AI Manager dashboard (from the Workflow engine)
   - Access to cloud-based AI application services depending on the chosen AI apps (from the Application server)
   - All required external/internet connections are outbound only.

3. Request Blackford to enable the AI applications of choice, typically first in a trial version limited in time and limited in studies (depending on the conditions of the AI app supplier). If the AI results satisfy the expected outcome, a purchase order must be submitted to Blackford to enable an annual subscription for specific applications and agreement for apps must be signed between Customer and Blackford.

4. Request PACS administration support to establish DICOM connectivity between the PACS and the Workflow engine server. Specifically
   - the PACS needs to send (automatically route) studies to the Workflow engine server
   - the Workflow engine server needs to send algorithm results to the PACS
   - if longitudinal analysis algorithms are to be operated the Workflow engine server needs to query the PACS for prior studies
   - the connection details (IP addresses/hostnames, ports, AE titles) are to be provided to Blackford for the configuration of the Workflow engine server

5. [Optional] Request IT administration support to integrate the result viewer with the local user management infrastructure. This involves certificate generation.

6. [Optional] Request RIS/Reporting system administration support to establish and configure HL7 connectivity between the corresponding systems and the Workflow engine server.

7. Identify test cases and data to perform a complete operational test once the platform is installed and fully configured. This includes a resource that can send data from the PACS and assess the reception of the results.

1 Blackford Analysis Ltd, Edinburgh, Scotland
2 https://dicom.nema.org/medical/dicom/current/output/html/part15.html
3 Regulations (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016. (General Data Protection Regulation)

00000434-00-00 * JUN 2024

How to reach us
Please visit www.philips.com/AI-manager