Stato e attività delle infrastrutture certificate

Barbara Martelli

EPIC
Enhanced PrIvacy and Compliance Cloud

Enhanced PrIvacy and Compliance Cloud is an ISO certified cloud platform	A region of INFN Cloud with a certified Information Security Management System
EPIC Cloud offers an IaaS Community Cloud for the communities of	Biomedical and genomic researchers Industrial researchers

Site locations: Bologna (active now), Bari and Catania sites will be added in June 2024 enabling for high availability and disaster recovery

Resource available today: about 700 TB of storage, 1440 cores, 10 TB RAM, 6 GPU A100

On going expansion with 3M euro of NRRP resources and 4M euro of funds from other projects

Progetti ospitati su EPIC

Progetto	HW (M€)	FTE
Harmony/Harmony+	Progetto in fase di conclusione, si discute della partecipazione alla fondazione	
DARE	3.2 (inseriti in gara TeRABIT)	4 TD CNAF, 1 TD Bari fino a dicembre ’26 (assunti)
ICSC Spoke8	0 (da richiedere a RAC)	1 TD CNAF, 2 TD Catania, 1 borsa Spoke0 dedicata ad attività di S8 (assunti)
HBD	6.5 (procedure d’acquisto da avviare, spalmati su 7 anni)	7/8 AR senior (1 assunto, 1 bandito e deserto, gli altri non ancora finanziati)
Sant’Orsola	0.34 (procedure d’acquisto da avviare, già disponibili)	2 AR senior (non assunti)
Sant’Orsola+	?	
IOR e AlmaHealthDB	?	3-4 ingegneri IOR-UniBO che lavorerebbero part-time al CNAF
Human Technopole	?	
Lafov PET	?	
UNCAN	?	
Numerosi IRCCS richiedono di utilizzare EPIC	?	
Intesa San Paolo (ICSC Spoke0) Smart City (ICSC Spoke9)	?	

Ulteriori dettagli sui progetti (aggiornato a Maggio ’23):
https://agenda.infn.it/event/34683/contributions/197354/attachments/105518/148360/20230523-datacloud-lifescience-v3.pdf

ISO/IEC 27001 27017 27018 certification

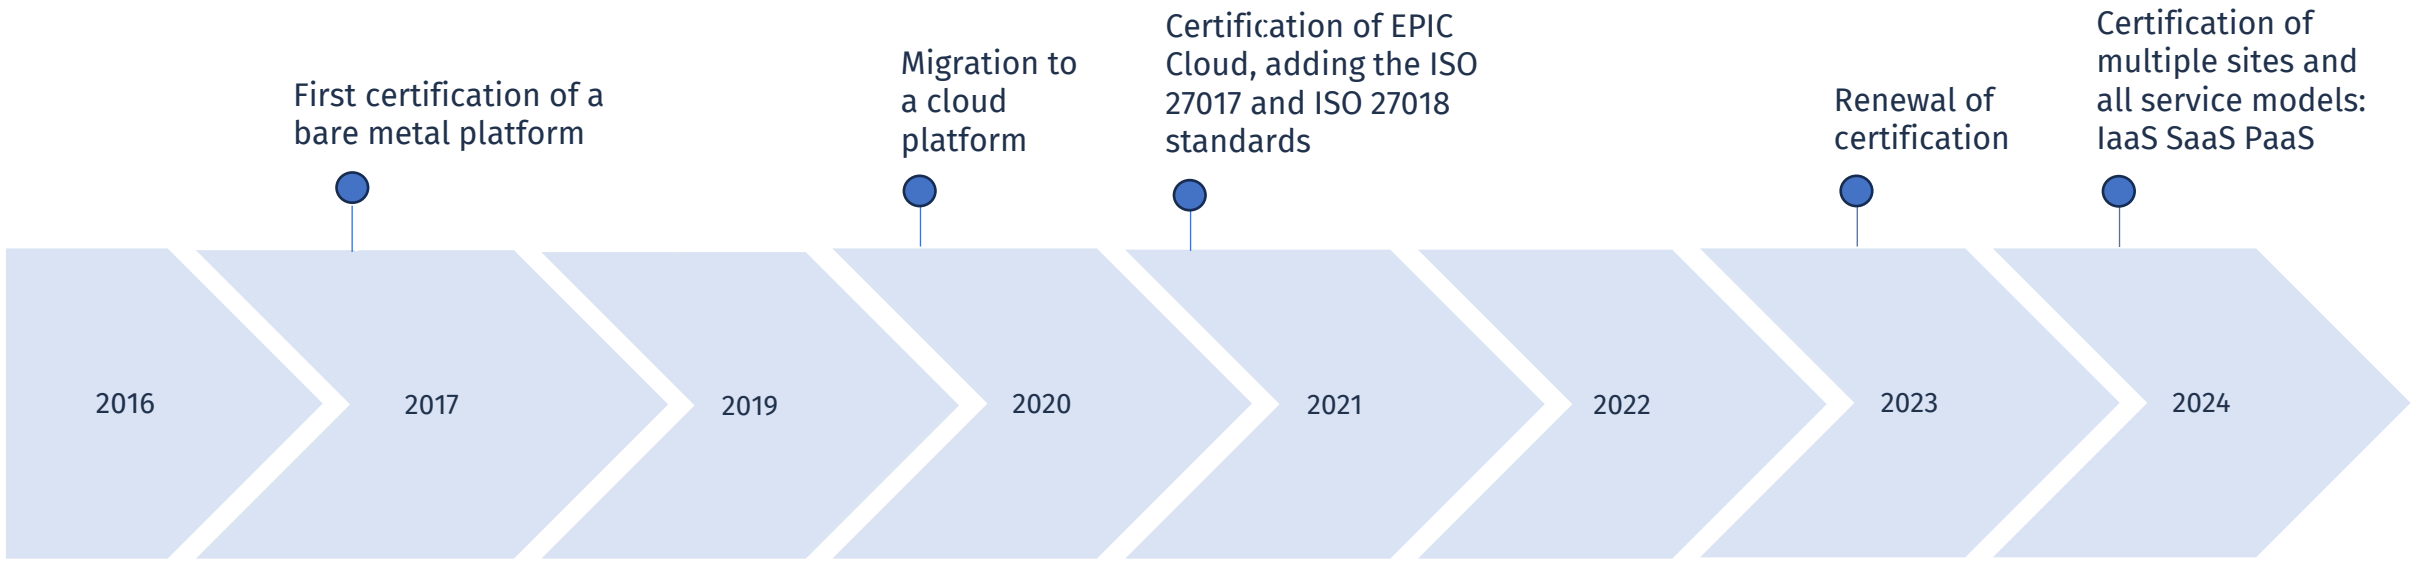

- In 2017 we started with a bare-metal that was ISO-27001 certified infrastructure
- In 2019 we moved to a Cloud-based infrastructure, and EPIC Cloud was born
- In 2021 we added the cloud certifications ISO-27017 and ISO-27018

EPIC Cloud Multisito

- il 26 ottobre scorso è stata firmata la disposizione del Presidente 26047, che istituisce lo steering committe di EPIC Cloud multisito.
- Lo steering è composto dai direttori delle quattro sedi coinvolte (Amministrazione Centrale, Bari, Catania e CNAF) ed ha il compito di allocare le risorse e definire gli obiettivi del Sistema Integrato di Gestione nazionale – quello che comunemente chiamiamo “certificazione ISO”.

Lo scope del nuovo certificato

Coprogettazione, sviluppo e manutenzione di soluzioni software di DataCloud per il settore della ricerca.

Erogazione di servizi di DataCloud IaaS, SaaS e PaaS in community deployment model.

Obiettivo da conseguire entro giugno ‘24

Stato delle attività

- Contratto di consulenza in fase di acquisto (RDA firmata prima di Natale su fondi Spoke0 assegnati a Bari)
- Necessario avviare la richiesta di offerte per l'audit di terza parte
  - Attualmente per 37 persone (~6 FTE), 1 sito e media complessità dell'SGSI spendiamo 6.5 kEuro per il primo anno e 2.6 kEuro per i due anni successivi
- Kick-off meeting in presenza il 19 gennaio in Presidenza
  - Impostazione delle attività dei process owner
  - Definizione delle policy da scrivere e approvare (link preliminare qui)
  - Discussione delle questioni tecniche ancora aperte (vedi prossime slide)

Processi e owner di processo (link)

Tipologia	Codice	Processo	Owner	WP
Management processes	MP 01	Information security governance/management interface process	Martelli	7
Core Processes	CP01	Security Policy management process	Bovo	
	CP02	Requirements management process	Donvito	3
	CP03	Information security risk assessment process	Carbone	4
	CP04	Information security risk treatment process	Ciaschini	4
	CP05	Security Implementation management process	V. Spinoso	4
	CP06	Process to control outsourced services	Monforte	
	CP07	Process to assure necessary awareness and competence	Costantini	2
	CP08	Information security incident management process	Peco	4
	CP09	Information security change management process	Stalio	1
	CP10	Internal audit	Belluomo	4
	CP11	Performance evaluation	Vistoli	7
	CP12	Information security improvement process	Lanzi	4
	CP13	Erogazione servizi DataCloud	Diego Michelotto	1
	CP14	Coprogettazione e sviluppo soluzioni software nell'ambito dei servizi DataCloud	Antonacci	5
	CP15	Manutenzione del deployment delle soluzioni sviluppate	Sergi	1
Support Processes	SP01	Records control process	Magenta	7
	SP02	Resource management process	Cesini	3
	SP03	Communication process	Rotondo	2
	SP04	Information-security customer relationship management process	Pellegrino	2

Project plan

Phase	Scope	Owner	Jul	Aug	Sep	23 Oct	Nov	Dec
			27 | 28 | 29 | 30	31 | 32 | 33 | 34| 35	36 | 37 | 38 | 39	40 | 41 | 42 | 42 | 43	44 | 45 | 46 | 47 | 48	49 | 50 | 51 | 52
Compliance	Assicurare il controllo e il soddisfacimento dei requisiti dello standard							
ISMS Organization	Definire l'assetto organizzativo del progetto: teams, ruoli, mansioni e deleghe. Per ciascun ruolo determinare le competenze minime richieste (conoscenze, abilità ed esperienza).							
Project tools	Identificare i tools che saranno utilizzati per pianificare, programmare e gestire tutte le attività di progetto comprese le informazioni documentate.							
Doc. Inf. Mngt	Tenere sotto controllo le informazioni documentate in termini di corretta identificazione, approvazione, distribuzione e conservazione.							
Process Approach	Definire l'approccio per processi utilizzato nello sviluppo del ISMS compresa la loro mappatura, descrizione e monitoraggio.							
Risk Mngt	Definire il nuovo Framework di gestione dei rischi.							
SOA Migration	Ri-mappare gli attuali controlli ISO27002:13 su ISO27002:22 aggiungendo e/o eliminandone altri secondo necessità e attribuendo a ciascuno di essi il livello di maturità conosciuto.							
Risk Assessment	Individuare, quantificare e valutare i rischi in linea con il Risk Framework							
Improvement	Pianificare e attuare le 3 linee di difesa del ISMS con l'emissione di almeno 1 piano di trattamento							

Da spostare in avanti di 6 mesi
(il consulente ci sta lavorando)

Gap analysis su sicurezza del SW DataCloud (1/2)

Prima analisi svolta da Marica Antonacci su orchestratore (link ai documenti qui)
In corso analisi sicurezza IAM (Ciaschini)

Work Package	Task	Estimated Effort	Resources
WP1 Governance & Compliance	T1.1 Define Security Standards & Policies	4 weeks	2 Security experts + 1 security champion for project team
	T1.2 Define Roles, Responsibilities, and Access Control Rules	2 weeks	1 Security expert + security champions
	T1.3 Security Training and Awareness	Ongoing through the project	Security champions + 2 security experts or training specialists (for outsourced courses)
WP2 Security Self-Assessment	T2.1 Create the projects inventory	3 weeks	Development leads + 1-2 developers for team
	T2.2 Initial evaluation and plan for improvement	3 weeks	Team developers (see the table below) + 1 Security expert (on-call)
	T2.3 Code review and initial plan implementation	16 weeks	Team developers (see the table below) + 1 Security expert (on-call)
WP3 Continuous monitoring	T3.1 Establish and maintain frameworks and processes for continuous monitoring	20 weeks	More work at startup with security champions and Datacloud operations support (if infrastructure setup is needed). Later: Automate most activities.
	T3.2 Metrics collection and stakeholder feedback	20 weeks	

Marica Antonacci

Gap analysis su sicurezza del SW DataCloud (2/2)

Project	LOC	Needed expertise (language, technologies, frameworks)	Estimated Resources	Current resources
Indigo-iam	To be completed	Java Spring Boot……. To be completed	To be completed	To be completed
orchestrator	27k	Java Spring Boot, Spring Security OpenID-Connect AuthN/Z SQL Database	4-5 skilled Java programmers	1 beginner Java developer
Cloud Provider Ranker	To be completed	Currently in Java – plan to re-write in Python	2 skilled python programmers	0
orchestrator-dashboard	4.7k	Language: Python, Javascript Flask, Flask-dance OpenID-Connect AuthN/Z SQL Database	2-3 skilled python programmers	1 python developer
Federation-registry & federation-feeder	7k	Language: Python FastAPI OpenID-Connect AuthN/Z Graph Database	2 python programmers	1 python developer
Slat	To be completed	Language: Python Flask, flaat Oauth2	2 python programmers	0
….				

Gruppo di sviluppo dello IAM

Marica Antonacci

30/11/23

Gap analysis sicurezza orchestratore e IAM

- In corso analisi sicurezza IAM (Ciaschini)
- Criticità:
  - le modifiche al disegno dell'orchestratore che stiamo proponendo in Spoke0 potrebbero non essere in linea con gli impegni presi da EPIC
  - In vista delle nuove normative sulla sicurezza (NIS2) gli adeguamenti a orchestratore e IAM non sono necessari solo a EPIC, ma a tutto DataCloud
  - A giugno certificheremo un sistema di autenticazione diverso da IAM (FreeIPA + Keycloak), ma questo non consente la federazione di risorse ed **è necessario sostituirlo con IAM il prima possibile**
  - Necessario scrivere/aggiornare/integrare circa 50 policy. Uno dei co-leader di WP7 è attualmente dipendente CNR, necessario negoziare il suo impegno su EPIC

Principali punti aperti (1/2)

- In vari contesti (ICSC-Spoke8, DARE, collaborazione INFN Sant'orsola, HBD) abbiamo promesso:
  - un sistema di autenticazione/autorizzazione
    - token based (integrato con le applicazioni life science)
    - in grado di federarsi con Cineca e in generale con EDUGAIN e IDEM. Su diversi tavoli viene richiesta l'integrazione con SPID
    - con multi-factor authentication
    - che rispetti gli standard di sicurezza richiesti dalla legge
    - Standard OIDC, ma anche in grado di integrare l'accesso via shell
  - Un sistema di orchestrazione che integri risorse cloud, HTC e HPC bubbles certificate -> *necessario unico orchestratore*
  - Un datalake sul modello WLCG (RUCIO + FTS) -> *necessario adattare questi software agli use case life science*

Principali punti aperti (2/2)

- In vari contesti (ICSC-Spoke8, DARE, collaborazione INFN Sant'orsola, HBD) abbiamo promesso:
  - La possibilità di trasferire grandi moli di dati in sicurezza -> *necessaria interazione con GARR*
  - La possibilità di estendere la cloud certificata per gestire decine di PB -> *necessario poter distribuire le risorse EPIC nel datacenter, superando il modello a silos attuale.*

- Effort rilevante richiesto a tutti i WP datacloud
- Sfatare il mito che vede la gestione di EPIC come un forte overhead, in realtà sarebbe sufficiente mappare le attività che già svolgiamo sull'SGSI. Commento dell'auditor a ottobre '23: *"ma se rispettate le leggi in tutto il datacenter, perchè certificate solo 3 rack?"*

Integrazioni e sinergie con altre comunità/progetti

- Alla plenaria di Spoke8 c'è stata apertura verso l'avvio di una PoC che testi l'infrastruttura ICSC utilizzando i dati Humanitas/Aosta/Cavalli
  - Occasione per lavorare su integrazione di policy e federazione con cloud certificata CINECA
  - Occasione per lavorare sulla rete (timore che il sustained throughput richiesto non sia raggiungibile verso INFN+CINECA)
- Elixir: volontà espressa da Pesole di costruire il nodo elixir-it su infrastruttura Bologna-Bari
  - Collaborazione con Casadio a Bologna da attivare (ferma solo per questioni di mancanza di tempo)
  - Attivazione di un nodo FEGA Italiano
- UNCAN: federazione europea di datalake per la ricerca oncologica. HBD DataCloud (-> cioè EPIC -> cioè ICSC) sarà il nodo Italiano. Le applicazioni ospitate sul nodo saranno quelle dell'ecosistema Elixir
- Utilizzare il blueprint AARC per IAM (sfruttare la collaborazione con GARR in TeRABIT, ICSC, AARC TREE, ecc)

Backup Slides

Technology

It is **based on the same technologies of INFN Cloud** (OpenStack, CEPH, IAM), with various enhancements introduced to meet higher security and privacy standards.

For example:

- The AuthN/Z system provides **2FA**, integration with web services, SSH and VPN (OpenVPN)
- **Network segregation** between OpenStack tenants is guarantee by **ACLs**
- At-rest and in-transit encryption
- Standard shared responsibility model:
  - User manages data, applications, runtime, middleware and OS
  - EPIC manages networking, storage, servers, virtualization
- Advanced logging and auditing services
  - **centralized syslog** server managed applying the ***segregation of duties*** principle

The EPIC Cloud infrastructure

- All services in HA (Availability is one aspect of security)
- It currently hosts 5 Projects
- ~ 1440 ~10TB RAM, ~700TB Disk
- Tenant/domain segregation
- Physical security (defined perimeters, controlled access to racks, TVCC)
- Network isolation from Tier-1 resources, Next-Generation Firewall (NGFW) in place

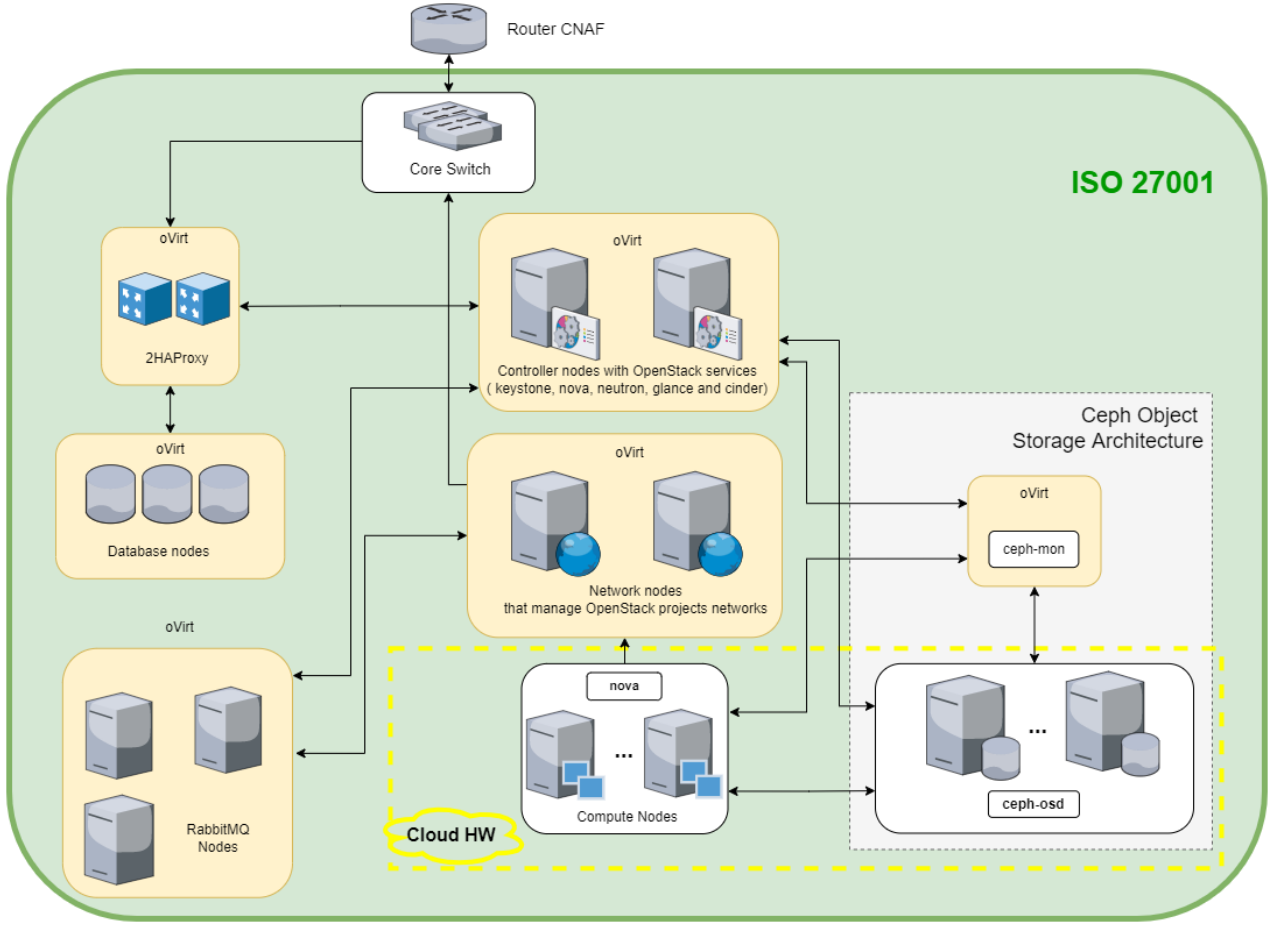

OpenStack Infrastructure

EPIC Cloud enables the federated data lake for health-related projects

- In general, we receive heterogeneous requirements form life-science communities: some need central repository, others need to store data locally
- Possible scenarios:
  1. Central harvesting of data collected remotely
  2. Edge-level anonymization and central ingestion and analysis of data
  3. Edge-level Feature Extraction and central ingestion and analysis of features
  4. Federated learning (training at local sites and algorithm publishing)

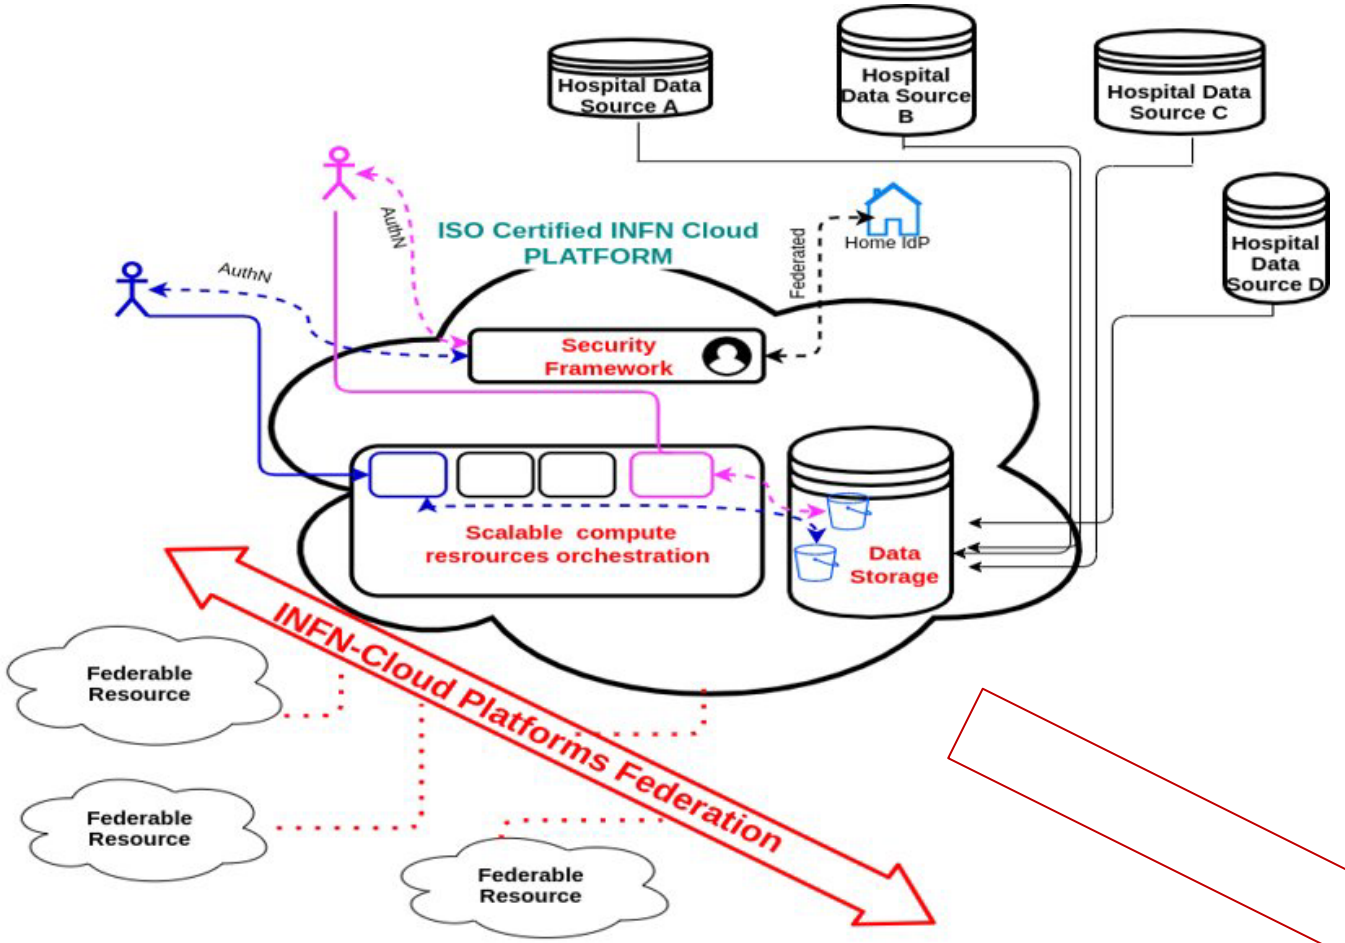

https://www.physicamedica.com/article/S1120-1797(21)00320-3/fulltext

Operational and support requirement to offer sensitive data management services within INFN EPIC Cloud

ISO/IEC 27001 27017 27018 certification

- This is a certification of the correct implementation of an **Information Security Management System** (**ISMS**)
- Information Security Management is about preserving the **Confidentiality, Integrity and Availability** (**CIA**) of information and associated facilities (systems, services, infrastructure or physical locations)
- It ensures **business continuity** by preventing and reducing the impact of security incidents
- Other properties can also be involved, such as **authenticity, accountability, non-repudiation, reliability** and **FAIRness**
- The **objectives** of the ISMS **are NOT fixed**, they depend on the context and are defined by the organization

ISMS: what's all about

Information Security Management System

- It is an **organizational framework** linking all the elements relevant to the information security, to assure that **policies**, **processes** and **security objectives** are implemented, communicated and assessed.
- It needs to **continually improve** -> **Deming Cycle**
- It is centered to the **risk assessment process** -> all decisions are based on the output of this process
- Goal: achieving the optimal **CIA balance**, i.e., ensuring Confidentiality of information, while still ensuring the information remains accessible to authorized persons and is not altered

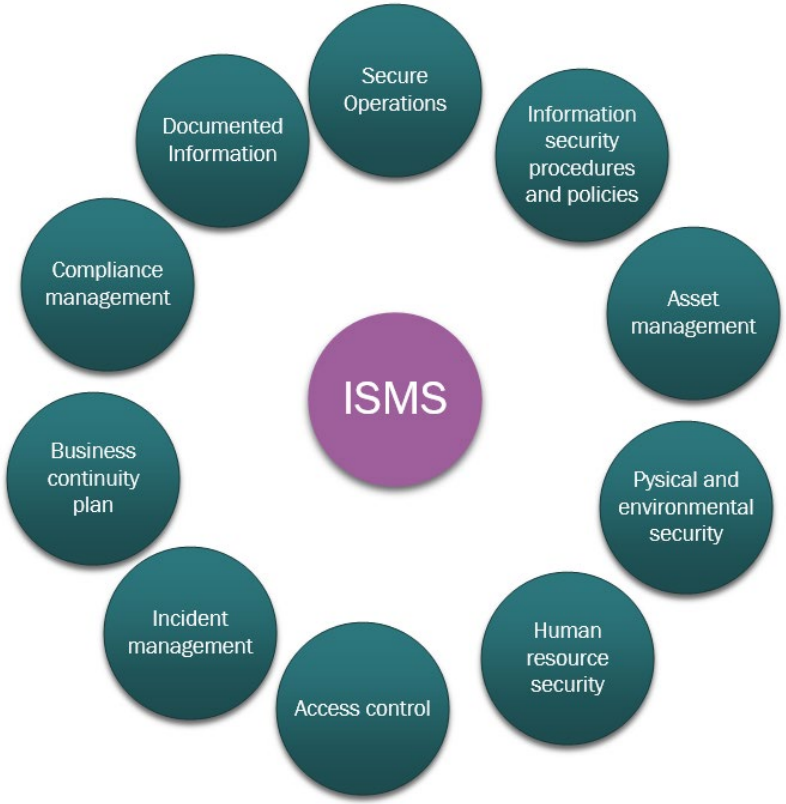

Deming Cycle and Risk Assessment

Plan
(context, leadership, planning & support)

Do
(risk assessment & operations)

Check
(performance evaluation)

Establish the organization

- Define the scope
- Analyze the context and define the security requirements and objectives
- Define the Security Policy
- Determine roles and responsabilities
- Define the risk assessment approach

Risk Assessment

Control Selection

Control Implementation (Operations)

Monitoring KPI

Internal Audit
External Audit

Effectiveness

Act
(continual improvement)

Risk Management Process in EPIC

- Iterative process aimed at supporting the decision-making process
- In EPIC is performed **once a month** and whenever a **relevant change** in the system **occurs**
  - ISO 27001 clause 8.2 requires to perform it *"at planned intervals or when significant changes are proposed or occur, taking account of the criteria established in 6.1.2 a). "*
- Established criteria to define Risk Owners

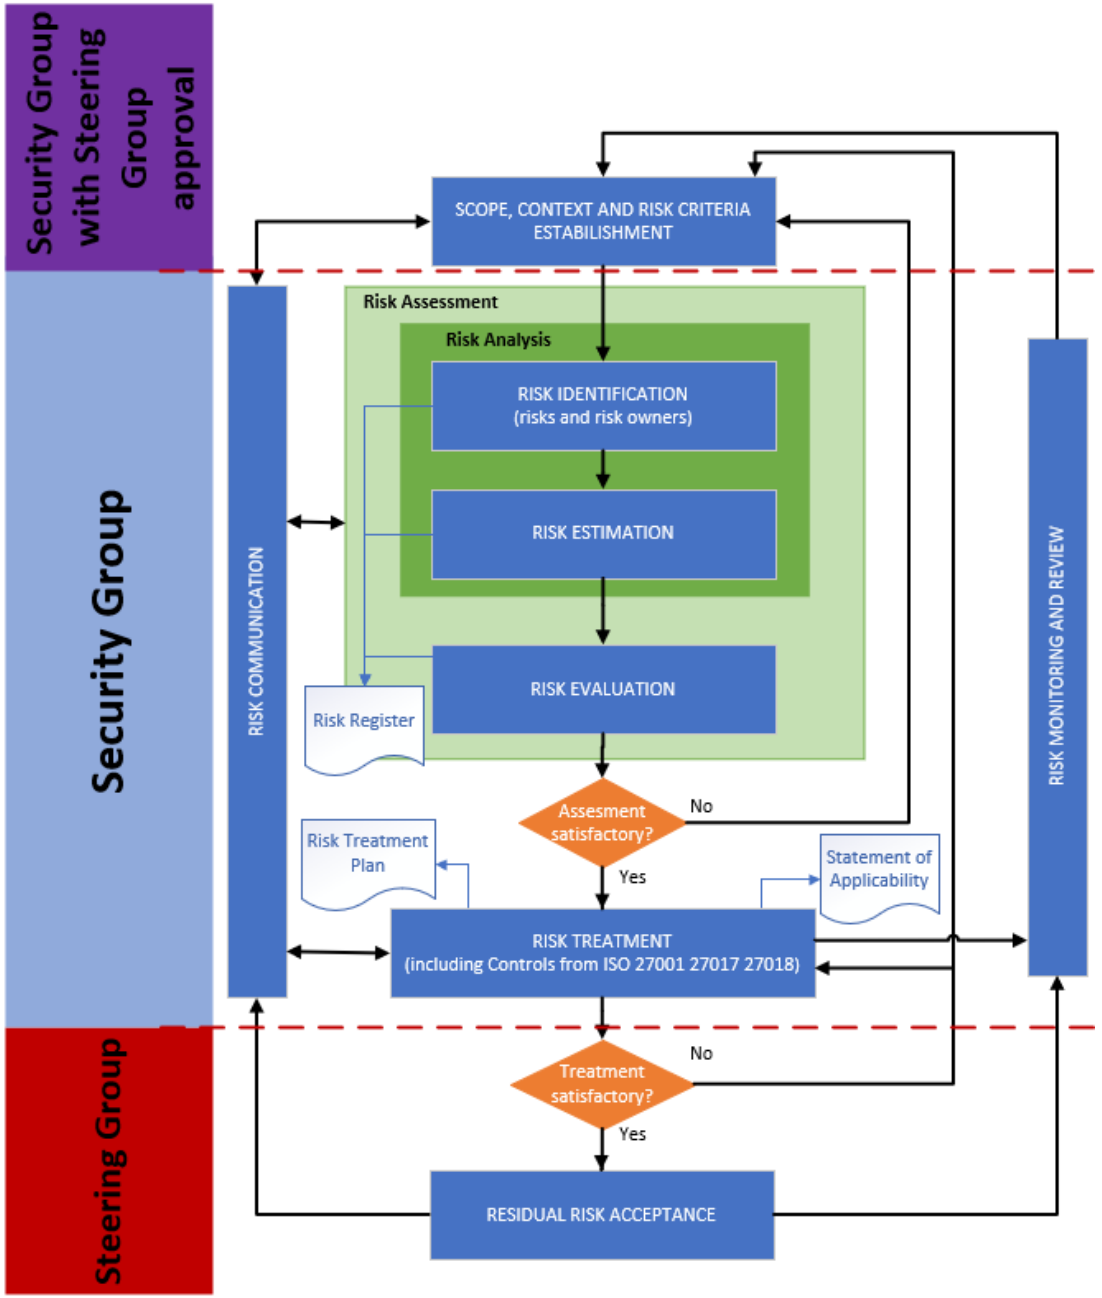

Security Incident Management

- Security incidents impacting CIA are managed following a documented security management process
- In case of high risk, the incident is escalated to INFN **CSIRT** (**Computer Security Incident Response Team**) https://csirt.infn.it
- The CSIRT team performs:
  - Incident identification
  - Incident categorization
  - Incident prioritization
  - Incident response
  - Incident closure

- A root cause analysis is also performed to prevent that the same incident occurs again

Required Document Information

8.1 Evidence Operational Planning and Control

A.5.1.1 Information Security Policy

9.2 Evidence of Audit Programme and Results

4.3 Scope

A.13.2.1 Information Transfer Policy and Procedures

A.10.1.2 Key Management Policy

9.3 Evidence of Results of Management Reviews

5.3 Organization Chart

6.1.3 Risk Treatment Process

A.11 Physical and Environmental Security Policy

A.12.2.1 Controls Against Malware Policy

A.6.2.1 Mobile Device Policy

6.2 Security Objectives

A.8.1.3 Acceptable Use of Asset Policy

8.3 Information Security Risk Treatment

A.12.3.1 Backup Policy and Procedures

ISO/IEC 27001

A.10.1.1 Cryptographic Control Policy

A.14.2.1 Secure Development Policy

A.11.2.9 Clean Desk Clean Screen Policy

A.15.1.1 Information Security Policy for Supplier Relationships

A.6.2.2 Teleworking Policy

A.12.6.2 Restrictions on Software Installation

A.14.1 Information Security Requirements Analysis and Specification

A.8.2.1 Information Classification Policy

7.2 Evidence of competence

A.16.1.4 Assessment of & Decision on Information Security Events

A.9.1.1 Access Control Policy

A.8 Asset Management Policy

A.17.2 Redundancies Policy

10.1 Actions Resulting from Nonconformities and Corrective Actions

6.1.2 Risk Assessment Process

8.2 Risk Assessment Results

A.9.3 Password Policy

9.1 Evidence of Monitoring and Measurement Results

Use case examples

HARMONY / HARMONY PLUS

Healthcare Alliance for Resourceful Medicines Offensive against Neoplasms in Hematology

Two IMI-2 European Projects. HARMONY PLUS, build up on the success of HARMONY, involves 39 partners and 8 Associated Partners from 10 countries.

- Use of big data analytics to accelerate blood cancer research
- Budget of 42.3M€ for HARMONY and 11.8M€ for HARMONY PLUS
- Over 100.000 patient datasets
- Harmonization of datasets
- Open and Standard interfaces

https://www.harmony-alliance.eu/

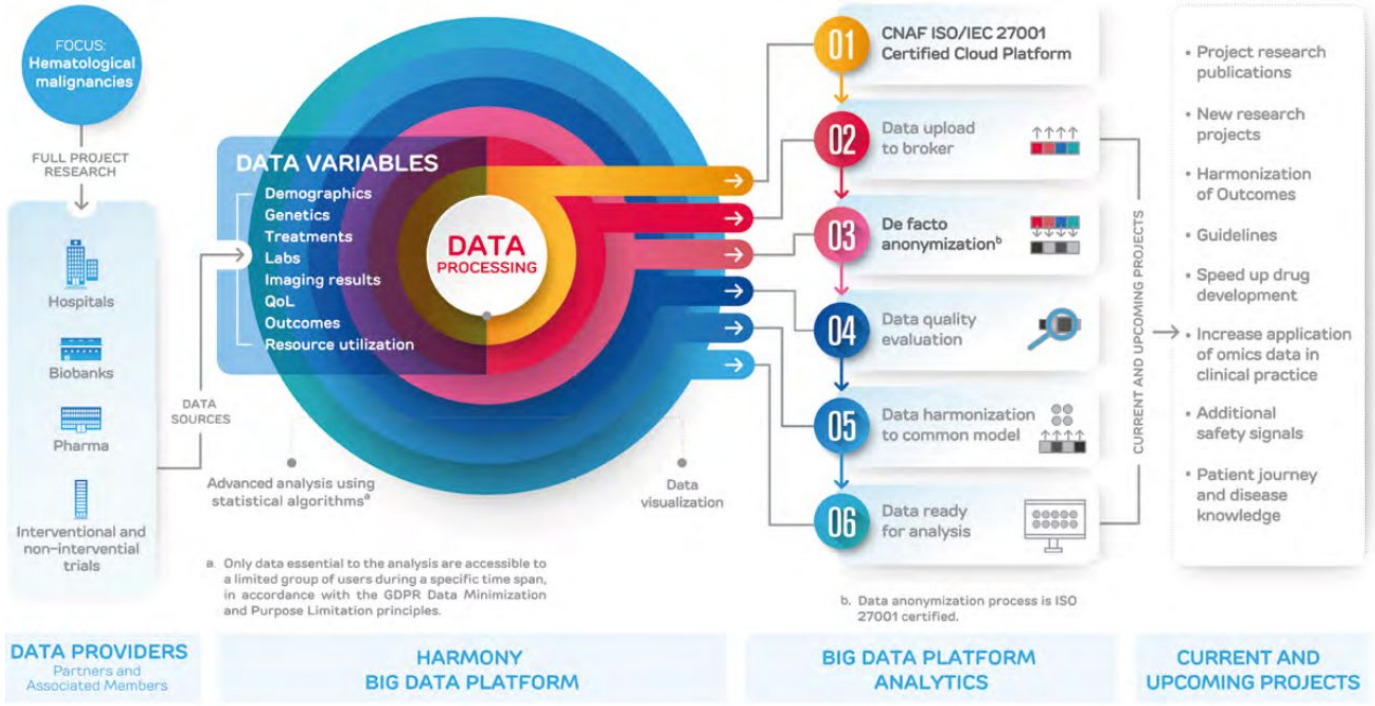

Figure 1 – Image based on an original idea by HARMONY (https://www.harmony-alliance.eu/).

HARMONY / HARMONY PLUS

Healthcare Alliance for Resourceful Medicines Offensive against Neoplasms in Hematology

- De facto anonymized data
  - GDPR does not apply, but the de-facto anonymization procedure assumes that the technology providers are ISO 27001 certified
- AgID measures (Standard and some of the Advanced)
- IaaS Service Model
  - HARMONY is the Data Controller
  - We are responsible for the infrastructure

https://www.harmony-alliance.eu/

Alleanza Contro il Cancro - ACC

The National Oncology Network founded in 2002 by the Ministry of Health, joined by 51 IRCCS, ISS, AIFA, INFN and Politecnico di Milano and several patients' associations to perform translational research in the field of cancer research.

- Genomic pseudonymized data
- GDPR applies
- AgID measures apply
- Italian Data Protection Authority rules apply

https://www.alleanzacontroilcancro.it/en/

ACC - Services

Alleanza Contro il Cancro

Two separate OpenStack projects:

- **ACC-Test**: services have been configured and tested and every change in configurations has been validated
- **ACC**: where services have been configured and tested and every change in configurations has been validated

Each VM is hardened according to ISO 27001 OpenSCAP profile

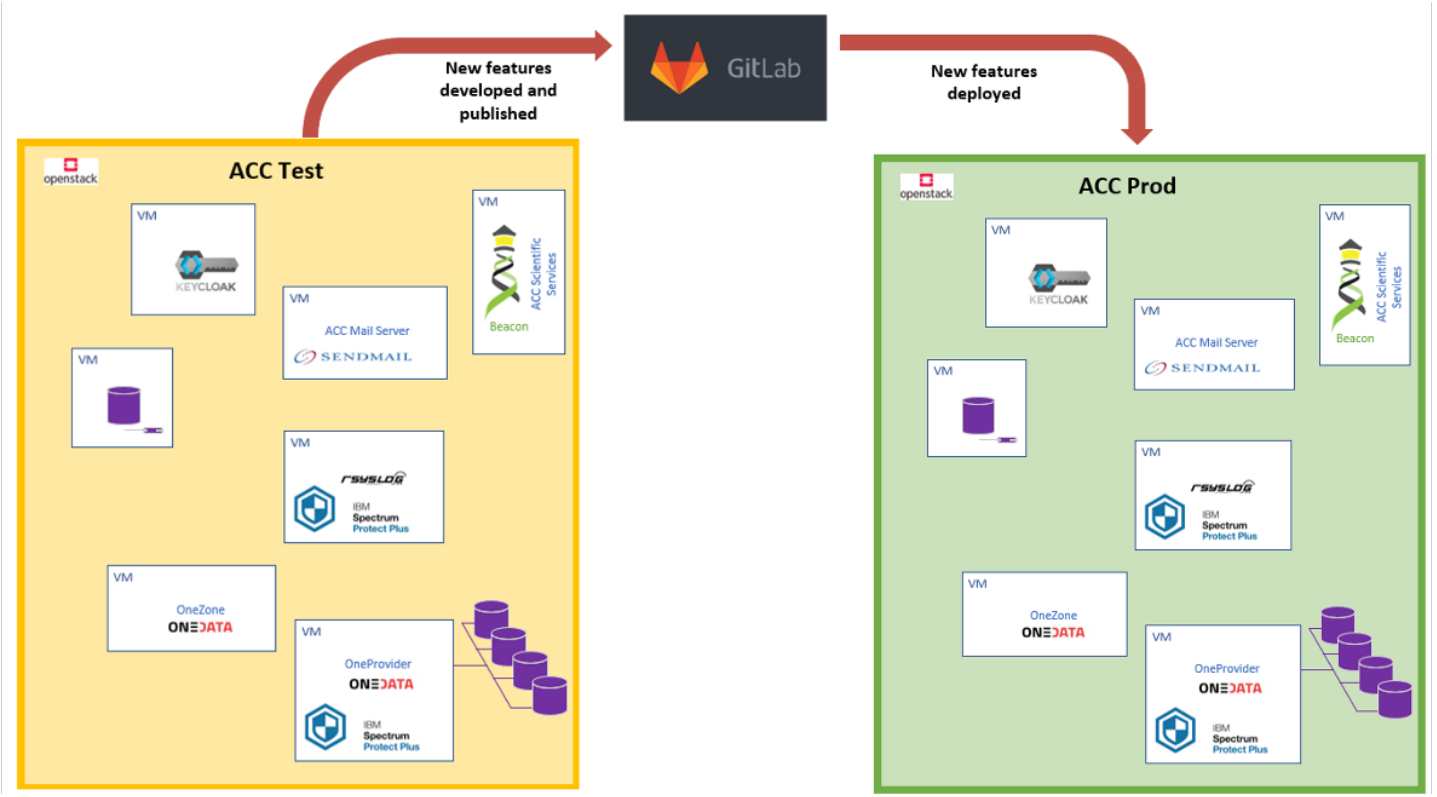

https://www.alleanzacontroilcancro.it/en/

PLANET

Pollution Lake ANalysis for Effective Therapy

An INFN funded research initiative aiming to implement an observational study to assess a possible statistical association between environmental pollution and Covid-19 infection, symptoms and course

- Data:
  - Pseudonymized clinical data (Covid-19 and Electronic Health Records from several hospitals)
  - Atmospheric data, population density, urban vs rural environment, mobility, socio-economic conditions
- Regulated by GDPR, Italian Data Protection Authority and ISS-INFN Convention
- Originally deployed on Cloud@CNAF

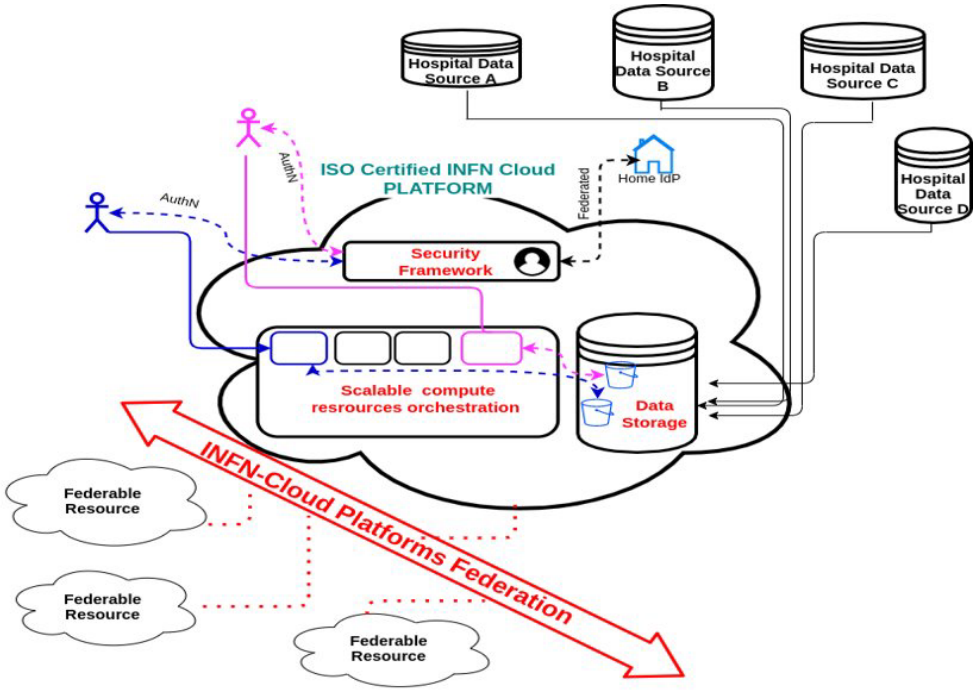

Our aim is to make this use case the first EPIC SaaS certified service (ambition/need to expand the scope of the ISO certificate to include SaaS services)

Health Big Data (HBD)

- Health Big Data is a 10-years project funded by the Italian Ministry of Health aiming at the creation of a federated and integrated big data platform for the health research at national level
  - 4 research networks: ACC, RIN, Cardio, IDEA
  - Research objectives: preventing diseases, personalizing treatments, improving the quality of life of patients
  - Budget: 55M€

Health Big Data (HBD)

- INFN is in the managing board of HBD. Its tasks include the definition of an integrated national platform and contributions to several Work Packages.
- The HBD architecture will provide solutions for several scenarios:
  1. Central harvesting of data collected remotely
  2. Edge anonymization, followed by central ingestion and analysis of data
  3. Edge feature extraction, followed by central ingestion and analysis of features
  4. Federated learning based on edge-based training, followed by publishing of the trained methods and by inference performed either centrally or at other edge locations

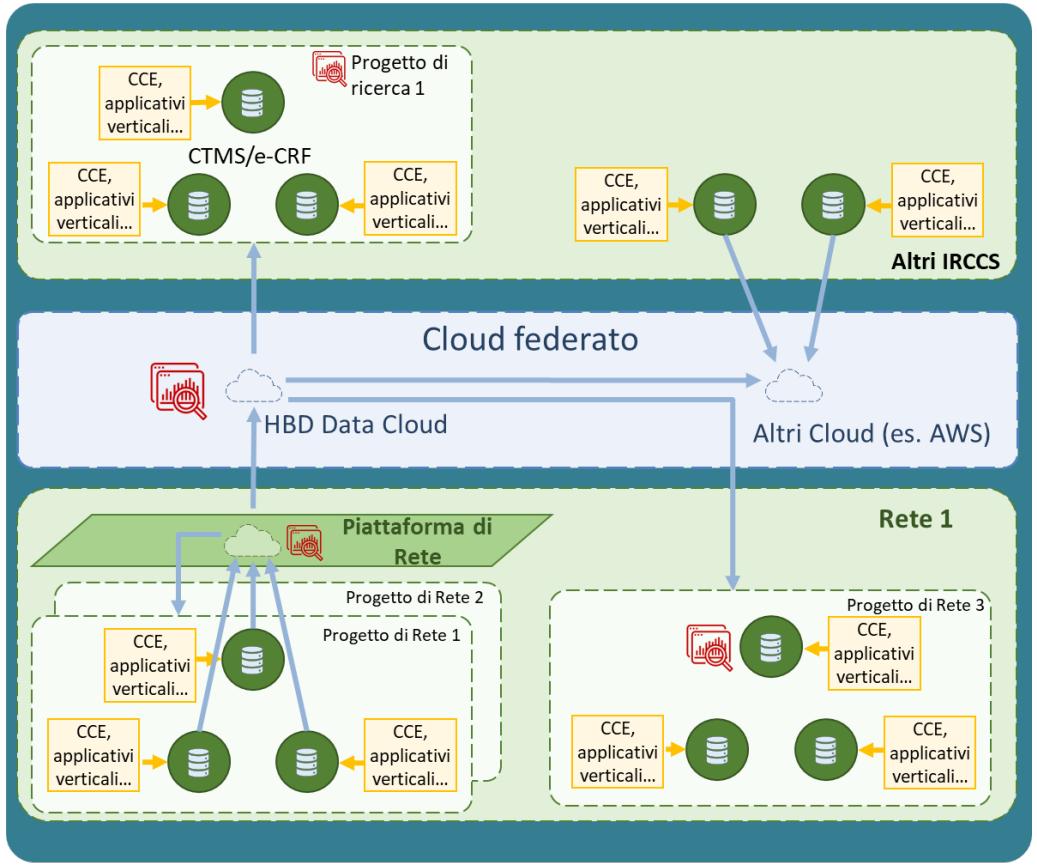

INFN-IRCCS Sant'Orsola Collaboration

Joint research agreement with the following objectives

- secure applications for genomic data
- GPU -based solutions for genomic analysis methods
- federated and integrated cloud platforms for homics data
- adaptation of genomic pipelines to cloud and data lake architectures based on microservices
- Integration of homics data and other clinical data like Electronic Medical Records (EMR)