

Archives and Data Management Systems in the Big Data Era

Wednesday 26 February 2025 - Friday 28 February 2025

CNR Bologna

Book of Abstracts

Contents

Toward a Public MAGIC Gamma-Ray Telescope Legacy Data Portal	1
Towards a data platform providing a holistic support to AtLAST operations	1
The AGILE data archive and data management system c/o the ASI Space Science Data Center	1
Updates on the IVOA standards ecosystem	2
The IVOA ExecutionBroker, an overview	2
ALMA Archive towards the Wideband Sensitivity Upgrade Era	3
Web-based approach to Data Management	3
Astronomical Archive as a Service: A Microservices-Based Hyperconverged Infrastructure	3
The Central Role of Database Technology in Astronomical Archives	4
The QUBRICS database for machine learning: architecture and performance.	4
RUCIO Data Management System: a Simple Archive or just a Distributed Storage System?	5
ASTRI-Horn, ASTRI-Miniarray and CTA Observatory a new Archival Perspective Design	5
Overview Data Lakes Spoke 3 and IDL	6
Advanced archival solutions	6
Towards FAIRness of radio data in the SKA era	6
The Gaia mission: towards a whole-sky Legacy Data management and archive infrastructure at OPS4 INAF-ASI facility	7
SSDC	7
IA2 INAF Data Center	7
ASTRI Horn, ASTRI Mini-Array and CTA Observatory a new Archival Perspective Design	7
SKA SRC - an international effort	8

Overview Data Lakes Spoke 3 and IDL	8
The GAIA use case in Spoke 3 and Innovation Grants	8
QStar Tape as NAS: sicurezza e affidabilità per l'archiviazione a lungo termine di immagini satellitari e dati scientifici	8
Data Management: best practices	9
Archive as a Service: A Microservices-Based Hyperconverged Infrastructure	9
Web-based approach to Data Management	9
Workflow management Systems	9
The origin of RAP	9
Effortless Identity Management	9
Round table	10
INAF Open Science	10
DAL implementations	10
The IVOA ExecutionBroker, an overview	10
TBD	10
Round table	10
Round table	11
Round table	11
Round table	11
Round table	11
Discussion	11
Discussion	11
Open Data Long Term Data Preservation for HEP: the CNAF Experience	11
Workflow Management Systems	11
The Sardinia Radio Telescope data management, from observing to archiving	12
Mining Archives: needs for Machine Learning	12
Soluzioni Quantum per l'Archiviazioni a Lungo Termine	12
How to handle Multi PB Environments for years, easily - The Pure Storage Way	12
VSP ONE Object Storage	13
Updates on the IVOA standards ecosystem	13

Session 2 - International and National Projects Archives / 1**Toward a Public MAGIC Gamma-Ray Telescope Legacy Data Portal****Author:** Michele Doro¹**Co-authors:** Elisa Prandini²; Ilaria Viale³; Stefano Marchesi⁴¹ *University of Padova*² *DFA Padova*³ *University and INFN Padova*⁴ *Istituto Nazionale di Astrofisica (INAF)***Corresponding Authors:** ilaria.viale.1@phd.unipd.it, elisaprandini@gmail.com, stefano.marchesi@inaf.it, michele.doro@unipd.it

The MAGIC telescopes are one of the three major IACTs (Imaging Atmospheric Cherenkov Telescopes) for observation of gamma rays in the TeV regime currently operative. MAGIC operates since 2003, and has published data of more than 80 sources, mostly blazars, in several emission states. MAGIC already distributes astronomical .fits files with basic final scientific products such as spectral energy distributions, light curves and skymaps from published results.

We are working on a updated format of high-level data files that contains more information (for a complete legacy of results), in ascii format (for human eye readability) and compliant with VO requirements. The final goal is the legacy of all high-level products in MAGIC paper, possibly including multi-wavelength data specifically analyzed for the publications. This activity is also meant toward the new generation of IACT, the Cherenkov Telescope Array Observatory.

In this contribution we describe this project.

Session1 - International Project Archives status / 2**Towards a data platform providing a holistic support to AtLAST operations****Author:** Francisco Miguel Montenegro Montes¹¹ *Universidad Complutense de Madrid***Corresponding Author:** fmontene@ucm.es

AtLAST (Atacama Large Aperture Sub-millimeter Telescope) is a project that aims at building and operating the next large single-dish facility observing at sub-mm wavelengths in Chile. In addition to pursuing transformational science and expanding technological limits, we put a strong emphasis on sustainability aspects. Among the many outcomes from the recently completed EU-funded design study was a number of transformational science cases covering several different of astrophysics. In addition, we have created a comprehensive operations plan that envisions remote distributed operations, a modern user support model and the implementation of infrastructures enabling an easy and transparent access to the data. Here, we want to explore strategies for describing, storing and sharing the different types of data that AtLAST will produce including engineering, weather and science data at different stages of processing. We envision a platform where the different AtLAST stakeholders (science users, proposal reviewers, telescope astronomers and operators, engineers etc) will be able to access the technical and scientific data they need as well as the tools necessary to analyse these data.

Session 2 - International and National Projects Archives / 3

The AGILE data archive and data management system c/o the ASI Space Science Data Center

Author: Carlotta Pittori¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: carlotta.pittori@inaf.it

AGILE (Astrorivelatore Gamma ad Immagini LEggero) has been a unique and successful space mission of the scientific program of the Italian Space Agency (ASI) focused on high-energy astrophysics, built and operated with the programmatic and technical support of INAF and INFN. During almost 17 years of observations in orbit (from April 23, 2007, to January 18, 2024), AGILE contributed in fundamental ways to high-energy astrophysics, cosmic-ray physics, solar physics and to the study of terrestrial gamma-ray flashes. Its archives and catalogs, compliant with FAIR principles, are available to the community through the AGILE Data Center, which is part of the ASI multi-mission Space Science Data Center (SSDC, previously known as ASDC). In this presentation, we give an overview of the AGILE data archive and data management system.

Even if AGILE was a *small* space mission that could not be classified within the context of so-called “big data”, it also had a Guest Observer Programme and addressed many issues that can be a useful example for the future.

Session 9 - Open Science, FAIR and Interoperability / 4

Updates on the IVOA standards ecosystem

Author: Marco Molinaro¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: marco.molinaro@inaf.it

Big data and the dynamic software infrastructure landscape are challenges also for the interoperability ecosystem defined by the IVOA standards.

This contribution reports on the current status of the IVOA activities and on how the VO faces the change in software paradigms and data analysis challenges keeping a eye continuously focused on open standards, preservation and interoperability.

5

The IVOA ExecutionBroker, an overview

Authors: Dave Morris¹; Sara Bertocco²

¹ *Manchester University*

² *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Authors: dave.morris@manchester.ac.uk, sara.bertocco@inaf.it

IVOA provides data models and tools to build data access and transfer services.

Size and complexity of the datasets available in the Virtual Observatory are increasing, thanks to the last generation of astronomical tools, and requires new solutions. In the VO community, Science Platforms to bring together the datasets co-located with the compute resources needed to analyse them are deployed or under development. However, to date, the IVOA does not provide any APIs or services that enable science platforms to exchange the software used to analyse the data (i.e. “to bring the code to the data”).

The IVOA Execution Broker provides a step towards making this possible.

This talk is an overview of the IVOA ExecutionBroker specification and how to use it.

Session1 - International Project Archives status / 6

ALMA Archive towards the Wideband Sensitivity Upgrade Era

Author: Marcella Massardi¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: massardi@ira.inaf.it

ALMA is undergoing the Wideband Sensitivity Upgrade (WSU), that will result in an increase of the instantaneous spectral bandwidth, spectral scan speed, and sensitivity for all observations. However, this will also bring some technical and data management challenges. I will present the major consequences of the upgrade for the data flow and the ALMA Science Archive development. I will also present the approaches that in the European ARC, and in particular in the Italian node, we are considering to face these new challenges.

Session 5 - Workflow management systems and Data Mining / 7

Web-based approach to Data Management

Author: Luciano Nicastrò¹

¹ *INAF-OAS Bologna*

Corresponding Author: luciano.nicastrò@inaf.it

Web-based services make astronomical data, such as images and object catalogs, easily queryable and accessible via a browser or script.

Nowadays the size of the data underlying these services does not impose any restriction on their usability. Relational or schema-less database management systems can be employed to manage both data and metadata. Furthermore, web tools can also include interactive graphs that allow intuitive use even by inexperienced users. Finally, software maintenance is centralized and no installation is required.

Despite astronomical web tools becoming more and more common and effective, the INAF community is apparently little involved in this kind of activity.

In the hope of increasing interest, I will briefly show some examples of web-based data management as part of projects in which I and INAF colleagues are involved.

Session 7 - Data Management Systems - part 2 / 8

Astronomical Archive as a Service: A Microservices-Based Hyper-converged Infrastructure

Author: Federico Fiordoliva¹

Co-author: Stefano Gallozzi¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Authors: federico.fiordoliva@inaf.it, stefano.gallozzi@inaf.it

The explosion of astronomical data necessitates innovative solutions for archiving, accessing, and processing these vast datasets.

This abstract presents “Astronomical Archive as a Service” (AAaaS), a novel approach built upon a hyperconverged infrastructure (HCI) and leveraging microservices technology.

AAaaS aims to provide a scalable, resilient, and cost-effective platform for managing astronomical data. HCI simplifies infrastructure management by converging compute, storage, networking, and virtualization resources.

By adopting a microservices architecture, AAaaS decomposes the archive into independent, loosely coupled services, each responsible for a specific function (e.g., data ingestion, metadata management, query processing, data retrieval). This modular design enhances scalability, fault tolerance, and maintainability.

Session 7 - Data Management Systems - part 2 / 9

The Central Role of Database Technology in Astronomical Archives

Author: Stefano Gallozzi¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: stefano.gallozzi@inaf.it

The common modus operandi in astrophysical research is to recycle knowledge already acquired and shared within a single team. Too little attention is given to the discovery of new technologies to better fulfill project specifications and requirements unless a technological limit is reached during any testing use-case.

Although the Database Management System (DBMS) is the backbone on which astronomical archives are based, the choice of the best suited DB for scientists’ needs often comes up against the need to put into operation a testing prototypes, so the search for an improved technology is simply bypassed in favor of a solution already known and mastered.

In this talk we will try to make a general excursus on different kinds of DBMS, related tool and show possible approaches to consciously make the best technology choice for the scientific project, or in case, adopt a new database paradigm called “Polyglot Persistence”, where different DB solutions are joined together within a common archive ecosystem. The Polyglot Persistence focus on atomic services identified by data providers and data consumer and build the service adaptively to match the data management and access requirements.

Session 5 - Workflow management systems and Data Mining / 10

The QUBRICS database for machine learning: architecture and performance.

Author: Giorgio Calderone¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: giorgio.calderone@inaf.it

The QUBRICS project aims to identify new, bright and high-redshift QSOs in the Southern Hemisphere using photometric data from several astronomical surveys, and machine learning methods for classification and redshift regression.

I will briefly describe the architecture and performance of our internal database, and its role within the QUBRICS ecosystem.

Session 7 - Data Management Systems - part 2 / 11

RUCIO Data Management System: a Simple Archive or just a Distributed Storage System?

Author: Stefano Gallozzi¹

Co-author: Georgios Zacharis²

¹ *INAF/OAR*

² *INAF/OAR*

Corresponding Authors: georgios.zacharis@inaf.it, stefano.gallozzi@inaf.it

In this talk we will present the features of RUCIO DM.

This software, developed and maintained by CERN on which the storage of different national and international research projects is based, represents a technological solution capable of federating together different pools of storage distributed in various locations and, through the use of a centralized catalog, provides the most common access interfaces to this data as well as a series of management services to promote data redundancy and the possibility of connecting directly to Workload Management Systems for the automation of data processing.

We will show how it is possible to use RUCIO for a simple astronomical archive, but also how to use RUCIO directly as distributed storage, interfacing it to the capabilities of an external database to ensure greater performance in terms of reliability and availability.

Session1 - International Project Archives status / 12

ASTRI-Horn, ASTRI-Miniarray and CTA Observatory a new Archival Perspective Design

Author: Stefano Gallozzi¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: stefano.gallozzi@inaf.it

In the framework of the international Cherenkov Telescope Array gamma-ray Observatory (CTAO), the Italian National Institute for Astrophysics (INAF) has developed an end-to-end prototype of the CTA small-size class of telescopes in dual-mirror (SST-2M) Schwarzschild-Couder configuration. The prototype, named ASTRI-Horn, is located at the INAF “M.C. Fracastoro” observing station in Serra La Nave (Mt. Etna, Sicily). The ASTRI-Miniarray Project is the continuation of ASTRI-Horn and consists of a nine ASTRI telescopes currently in construction phase sited at the Teide Observatory in Tenerife Island of Canary Islands.

Off-site data archiving & handling management system for both ASTRI-Horn prototype and Miniarray are developed within INAF institutes and take care of the full data-chain produced by the different observatory scientific devices at different analysis steps. The Archive System provides the access to data at different user-levels and for different use-cases, each one with a customized data organization. The Astri Miniarray Archive System (AMAS) environment is used as technological testbed for archival, simulation, data-processing, quick look and quality check as well as scientific validation of intermediate and final data products during the workload.

In this context AMAS system applied to VHE data has to be considered as a preproduction prototype of the CTAO Archive System and may become a de-facto standard for multipurpose adaptive storage applications to be adopted by several heterogenous scientific communities as well as in European Projects and Collaborations.

Under the CTAO Collaboration the INAF has the leadership of several working package, in particular the Bulk Data Management System (BDMS) that is a part of the Data Processing and Preservation System (DPPS), which directly interfaces with Science Archive hosted in Science User Support System (SUSS). Due to the huge amount of data (produced in two different sites), to be processed and for which the scientific community needs to have access, all the most important operations has been moved off-site, distributing both computing and archival in 5 different datacenters, geographically distant and separated. One of these datacenters is hosted in Frascati in synergy between INAF and INFN.

Session 4 - Industrial contribution / 13

Overview Data Lakes Spoke 3 and IDL

Authors: Giacomo Coran¹; Massimo Costantini¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Authors: giacomo.coran@inaf.it, massimo.costantini@inaf.it

The following presentation will illustrate the Spoke 3 WP4 and IDL projects and the work done in creating a data lake for each project and in developing the infrastructure capable of managing data and metadata ingestion and retrieval, respecting interoperability and FAIR and Open Science principles.

Session 4 - Industrial contribution / 14

Advanced archival solutions

Author: Sandro De Santis¹

¹ *IBM*

Corresponding Author: sandro.de.santis2@ibm.com

IBM illustrerà le soluzioni di memorizzazione ed archiviazione per ambienti complessi e multi-attaforma: on premise, distribuiti e cloud oriented.

Partendo dalle soluzioni storiche (come le librerie robotizzate) arriveremo alle soluzioni software defined (come Storage Scale/GPFS, Ceph e Cloud Object Storage), che possono utilizzare sia appliance IBM che HW standard.

Big Internationa Projects;

* Gestione Dati in centri HPC, PNRR e Industria;

* Data Management Systems;

Session 9 - Open Science, FAIR and Interoperability / 15

Towards FAIRness of radio data in the SKA era

Author: Vincenzo Galluzzi¹

Co-authors: Alessandra Zanichelli¹; Cristina Knapic¹; Fabio Roberto Vitello; Giuseppe Tudisco¹; Marco Molinaro¹; Robert Butora¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Authors: marco.molinaro@inaf.it, giuseppe.tudisco@inaf.it, a.zanichelli@ira.inaf.it, fabio.vitello@inaf.it, cristina.knapic@inaf.it, vincenzo.galluzzi@inaf.it, robert.butora@inaf.it

The increasing importance of Science Archives and archive mining in defining the ultimate productivity of an observing facility motivated the Italian Centre for Astronomical Archives (IA2) to develop and maintain the INAF Radio Data Archive. Such a geographically-distributed archival facility flexibly handles different data models and formats aimed at data discovery/access through Virtual Observatory (VO).

The activity related to the INAF Radio Data Archive led us to join the IVOA Radio Interest Group, whose goal is the interoperability, and more generally the FAIRness, of radio data. In this talk we will present our contribution, mainly focused on the definition of requirements and use cases for the representation of radio astronomy data in the VO. This implied the identification of metadata concepts needed by the radio domain that are not currently supported by the VO, hence the definition of a radio-specific data model that has been published as an IVOA Proposed Recommendation.

Complementary to data discovery, archival models for data ingestion and operations are required by large projects like the SKA and its SKA Regional Centre Network. In this respect we will present our activity within the Orange and Azure SRCNet Teams, mainly focused on the definition of observatory and advanced data products, their relation with science use cases and the definition of requirements for visualisation tools. Also, a fundamental contribution is being given on the mapping of SKA metadata onto existing and widely used data models, like the CADC Common Archive Observation Model and the IVOA ObsCore itself.

Session1 - International Project Archives status / 16

The Gaia mission: towards a whole-sky Legacy Data management and archive infrastructure at OPS4 INAF-ASI facility

Corresponding Author: busonero@oato.inaf.it

Session 2 - International and National Projects Archives / 17

SSDC

Corresponding Author: matteo.perri@inaf.it

Session 2 - International and National Projects Archives / 18

IA2 INAF Data Center

Corresponding Author: cristina.knapic@inaf.it

Session 2 - International and National Projects Archives / 19

ASTRI Horn, ASTRI Mini-Array and CTA Observatory a new Archival Perspective Design

Corresponding Author: stefano.gallozzi@inaf.it

Session 2 - International and National Projects Archives / 20**SKA SRC - an international effort**

Corresponding Author: gheller@ira.inaf.it

Session 3 - Archival activities in HPC centers / 21**Overview Data Lakes Spoke 3 and IDL**

Corresponding Authors: giacomo.coran@inaf.it, massimo.costantini@inaf.it

Session 3 - Archival activities in HPC centers / 22**The GAIA use case in Spoke 3 and Innovation Grants**

Corresponding Authors: sara.gelsumini@inaf.it, enrico.licata@inaf.it

Session 4 - Industrial contribution / 23**QStar Tape as NAS: sicurezza e affidabilità per l'archiviazione a lungo termine di immagini satellitari e dati scientifici**

Nel panorama della gestione dei dati, la crescita esponenziale delle informazioni impone soluzioni di archiviazione che siano non solo scalabili, ma anche sicure e sostenibili.

QStar Tape as NAS risponde a queste esigenze offrendo un'alternativa strategica ai tradizionali sistemi di storage su disco, particolarmente adatta per la conservazione di immagini satellitari, dati provenienti da ricerche scientifiche e processi di High-Performance Computing (HPC).

Protezione e sicurezza dei dati

L'archiviazione di dati scientifici e immagini satellitari richiede elevati standard di sicurezza e affidabilità nel tempo. I sistemi a nastro garantiscono una protezione avanzata contro le minacce informatiche: a differenza dei dischi, le cartucce LTO possono essere conservate offline, eliminando il rischio di attacchi ransomware o di corruzione accidentale dei dati. Inoltre, la cifratura hardware garantisce la protezione delle informazioni sensibili, fondamentali per le agenzie spaziali, i centri di ricerca e le istituzioni governative.

Archiviazione a lungo termine e sostenibilità

L'enorme quantità di dati generati dai satelliti e dai progetti scientifici richiede un sistema capace di garantire la conservazione per decenni senza degradazione. Le tecnologie a nastro offrono una durata superiore rispetto ai dischi magnetici, con una stabilità comprovata per oltre 30 anni. Inoltre, il consumo energetico è drasticamente inferiore: a differenza degli hard disk attivi 24/7, le cartucce LTO non necessitano di alimentazione continua, riducendo i costi operativi e l'impatto ambientale.

Scalabilità ed efficienza per i Big Data

I progetti di osservazione terrestre, le simulazioni scientifiche e le applicazioni HPC generano petabyte di dati, spesso con la necessità di un accesso rapido e continuo. QStar Tape as NAS consente di utilizzare i nastri come un filesystem NAS tradizionale, offrendo una soluzione scalabile e conveniente per la gestione dei Big Data. I costi complessivi dello storage risultano notevolmente ridotti rispetto ai dischi, senza compromessi in termini di accessibilità e prestazioni. Conclusione QStar Tape as NAS rappresenta la scelta ideale per enti spaziali, istituti di ricerca e aziende che necessitano di un archivio sicuro, scalabile e sostenibile. Grazie alla combinazione di affidabilità,

protezione e ottimizzazione dei costi, questa tecnologia consente di gestire e preservare enormi volumi di dati, garantendo l'accessibilità e l'integrità delle informazioni nel lungo periodo.

Session 6 - Data Management Systems - part 1 / 24

Data Management: best practices

Corresponding Author: stefano.gallozzi@inaf.it

How to build an on-demand system to meet the challenges of large projects and big data

Session 6 - Data Management Systems - part 1 / 25

Archive as a Service: A Microservices-Based Hyperconverged Infrastructure

Corresponding Author: federico.fiordoliva@inaf.it

Session 6 - Data Management Systems - part 1 / 26

Web-based approach to Data Management

Corresponding Author: luciano.nicastro@inaf.it

Session 7 - Data Management Systems - part 2 / 27

Workflow management Systems

Corresponding Author: andrea.bignamini@inaf.it

Session 8 - A&A / 28

The origin of RAP

Corresponding Author: f.tinarelli@ira.inaf.it

Session 8 - A&A / 29

Effortless Identity Management

Author: Massimo Costantini¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: massimo.costantini@inaf.it

In this presentation, we will explore how to quickly and efficiently implement authentication, authorization, and Single Sign-On (SSO) using Keycloak, an open-source identity and access management solution. Through a video demo, we will show how easy it is to set up a secure authentication system, manage user permissions, and enable login across multiple applications.

Session 8 - A&A / 30

Round table

Session 5 - Workflow management systems and Data Mining / 31

INAF Open Science

Corresponding Author: riccardo.smareglia@inaf.it

Session 10 - Interoperability and Virtual Observatory / 32

DAL implementations

Corresponding Author: robert.butora@inaf.it

Session 10 - Interoperability and Virtual Observatory / 33

The IVOA ExecutionBroker, an overview

Corresponding Author: sara.bertocco@inaf.it

Session 10 - Interoperability and Virtual Observatory / 34

TBD

Corresponding Author: martina.vicinanza@inaf.it

Session 3 - Archival activities in HPC centers / 35

Round table

Session 6 - Data Management Systems - part 1 / 36**Round table****Session 7 - Data Management Systems - part 2 / 37****Round table****Session 5 - Workflow management systems and Data Mining / 38****Round table****Session 9 - Open Science, FAIR and Interoperability / 39****Discussion****Session 10 - Interoperability and Virtual Observatory / 40****Discussion****Session 9 - Open Science, FAIR and Interoperability / 41****Open Data Long Term Data Preservation for HEP: the CNAF Experience**

Author: Luca Dell'Angello¹

Co-author: Stefano Dal Pra¹

¹ *INFN-CNAF*

Corresponding Authors: luca.dellagnello@cnafe.infn.it, stefano.dalpra@cnafe.infn.it

In this talk, we quickly remind the FAIR data principles, the importance and need for Long Term Data Preservation (LTDP) solutions in High Energy Physics (HEP). We present a case study of the CDF experiment. We showcase how CNAF addressed the preservation challenges highlighting both technical solutions and broader implications for open science and data fairness. We finally share the insights gained through this experience.

Session 5 - Workflow management systems and Data Mining / 42

Workflow Management Systems

Corresponding Author: andrea.bignamini@inaf.it

Session 3 - Archival activities in HPC centers / 43

The Sardinia Radio Telescope data management, from observing to archiving

Corresponding Author: antonietta.fara@inaf.it

Session 5 - Workflow management systems and Data Mining / 44

Mining Archives: needs for Machine Learning

Author: Stefano Cavuoti¹

Co-author: Giuseppe Riccio²

¹ INAF - Astronomical Observatory of Capodimonte Napoli

² Istituto Nazionale di Astrofisica (INAF)

Corresponding Authors: stefano.cavuoti@inaf.it, giuseppe.riccio@inaf.it

Session 4 - Industrial contribution / 45

Soluzioni Quantum per l'Archiviazioni a Lungo Termine

Author: Stefano Chiappella^{None}

Co-author: Emanuele Bascape

Corresponding Authors: emanuele.bascape@quantum.com, stefano.chiappella@quantum.com

Le organizzazioni necessitano sempre più di soluzioni moderne e facili da usare che le aiutino a memorizzare, gestire, proteggere, archiviare e analizzare enormi quantità di dati.

Quantum ActiveScale Cold Storage è una soluzione di archiviazione cloud ibrida progettata per offrire uno storage a lungo termine altamente scalabile, sicuro e conveniente per dati ad alta e bassa frequenza di accesso. Basata su un'architettura Object Storage distribuita e supportando le API S3 e S3 GLACIER, combina tecnologie di erasure coding e data durability per garantire affidabilità ed efficienza. Ideale per backup, archiviazione di dati scientifici, contenuti multimediali e conformità normativa, ActiveScale Cold Storage riduce i costi operativi grazie a un modello di accesso ottimizzato per il retrieval su richiesta. Integrandosi con servizi cloud e on-premise, offre un'alternativa economicamente vantaggiosa ai tradizionali archivi a nastro, con tempi di accesso ai dati più rapidi e gestione semplificata.

Session 4 - Industrial contribution / 46

How to handle Multi PB Environments for years, easily - The Pure Storage Way

Author: Umberto Galtarossa¹

¹ *Pure Storage*

Corresponding Author: ugaltarossa@purestorage.com

Pure Storage offers robust archiving solutions, prominently featuring the Evergreen architecture, high density, and lower costs. The Evergreen architecture allows Pure Storage to deliver continuous, non-disruptive upgrades, ensuring that the storage infrastructure remains modern and efficient over a 10+ year lifespan without the need for disruptive migrations or re-buys. This architecture guarantees that customers can minimize downtime and future-proof their storage needs.

Our Platforms provides efficient, high-throughput, scale-out storage, optimized for large-scale environments with both file and object storage capabilities. The storage solutions are designed to be energy efficient, consuming significantly less power and generating less heat compared to traditional systems. This focus on efficiency not only reduces operational costs but also aligns with modern sustainability goals by minimizing the environmental impact.

Additionally, Pure Storage's solutions are cost-effective. Consolidating storage silos into a unified platform reduces the complexity and costs associated with managing multiple systems. The predictable pricing model and the option for as-a-service offerings through Evergreen//One provide financial flexibility and help organizations better manage their storage expenses while ensuring robust data protection and rapid recovery capabilities via features like SafeMode snapshots.

These combined advantages make Pure Storage's archiving solutions an excellent choice for organizations looking to achieve high performance, low costs, and reliable, future-proof storage infrastructure.

Session 4 - Industrial contribution / 47

VSP ONE Object Storage

Corresponding Author: cristian.magro@hitachivantara.com

Session 10 - Interoperability and Virtual Observatory / 48

Updates on the IVOA standards ecosystem

Author: Marco Molinaro¹

¹ *Istituto Nazionale di Astrofisica (INAF)*

Corresponding Author: marco.molinaro@inaf.it

Big data and the dynamic software infrastructure landscape are challenges also for the interoperability ecosystem defined by the IVOA standards.

This contribution reports on the current status of the IVOA activities and on how the VO faces the change in software paradigms and data analysis challenges keeping a eye continuously focused on open standards, preservation and interoperability.