

CS3 2026 - Cloud Storage Synchronization and Sharing



Tuesday, March 17, 2026 - Thursday, March 19, 2026

University of Oslo

Programme

The 2026 CS3 Conference will take place on 17-19 March 2026 at University of Oslo.

We look forward to meeting old and new colleagues within the CS3 community. Reconnect, inspire and get inspired, learn from each other and have some fun together, too.

The main session will start on Tuesday, 17 March and will be preceded on Monday 16 March afternoon by a co-located SIG-CISS session of the GEANT Association.

Keynotes

Digital Sovereignty or Digital Dependency: Europe's Tech Moment of Truth
Frank Karlitschek, CEO, Nextcloud

Amid intensifying geopolitical tensions, Europe is accelerating its drive toward digital sovereignty. Technology providers, enterprises, and public institutions are being reshaped by the continent's push for technological independence supported by an expanding regulatory framework that now spans the DMA, DSA, GDPR, NIS2, and upcoming European cybersecurity initiatives.

This keynote examines how these developments are redefining the landscape for collaboration platforms and File Sync & Share solutions technologies that sit at the core of modern productivity and information security.

We will explore what the next generation of collaboration software must deliver to thrive in this environment: decentralised and federated architectures, robust Open Source ecosystems, privacy respecting AI by design, transparent data flows, and open standards that empower organisations rather than lock them in.

Attendees will gain a behind-the-scenes perspective on how European tech policy is shaping the future of digital collaboration along with insights and lessons learned from more than 16 years of building Open Source collaboration technologies.

Frank Karlitschek is a long time open source contributor and former board member of the KDE e.V. He founded ownCloud in 2010 and the successor Nextcloud in 2016 to create a fully open source and decentralized alternative to big centralized cloud companies. Frank was an invited expert at the W3C to help to create the ActivityPub standard. Frank has spoken at MIT, CERN, Harvard and ETH and keynoted many conferences. Frank is the founder and CEO of Nextcloud GmbH. He is also a fellow of Open Forum Europe and an advisor to the United Nations regarding Open Source. Frank won the European SFS Free Software Award 2023 and the Acteurs du Libre European Award 2023.

The HPC-Cloud Pipeline: Data Provisioning for Destination Earth and the AI Revolution in Weather Forecasting
Dr. Tiago Quintino, Head of Development, ECMWF

The European Commission's Destination Earth (DestinE) initiative represents a paradigm shift in Earth system simulation, aiming to develop highly accurate digital twins of the Earth. At the core of this endeavour lies the Digital Twin Engine (DTE), an innovative software framework designed to connect the extreme data generation capabilities of High-Performance Computing (HPC) with the

interactive, user-centric flexibility of cloud environments.

This keynote discusses the architectural challenges and solutions involved in designing the DTE to enable seamless data provisioning and sharing. A key focus will be the convergence of Numerical Weather Prediction (NWP), Climate Information, and Machine Learning (ML). We will examine how the DTE is evolving to support "AI-ready" datasets, particularly addressing the extensive data handling requirements for ML training, including the upcoming ERA6 reanalysis—the successor to ERA5 and a vital component for future AI model training. Additionally, we will outline the data throughput challenges related to operationalising ECMWF's Artificial Intelligence Forecasting System (AIFS) and explain how we are developing scalable workflows to support the next generation of data-intensive prediction systems.

Dr. Tiago Quintino is the Head of Development at ECMWF, leading the delivery of high-throughput software infrastructure critical to the Centre's operational forecasting capabilities, from observation acquisition to product dissemination. His work supports the European Commission's Destination Earth initiative by developing cloud-native services that bridge High-Performance Computing with AI/ML applications and end-users. Spanning over 25 years in Aerospace and Numerical Weather Prediction (NWP), Dr. Quintino's career currently focuses on optimising massive NWP data workflows to enable the convergence of HPC and machine learning towards Exascale computing.

Clinical research and data in the AI era: Lessons and future directions from the AI-Mind project
Dr. Hatlestad-Hall, Oslo University Hospital

Concluded in February 2026, the ambitious AI-Mind project (Grant ID: 964220) was a pioneer in the new era of clinical AI research. During the five-year project period, the initiative successfully established one of the world's most comprehensive longitudinal datasets focused on Mild Cognitive Impairment (MCI) and dementia progression.

Collected through harmonised procedures across Norway, Finland, Italy, and Spain, the dataset covers over 1,000 individuals. It comprises more than 3,500 EEG recordings and digital cognitive assessments, 1,800 blood biomarker samples, and nearly 1,000 diagnostic conclusions. However, the existence of valuable data is only the first step; preparing such complex, sensitive information for clinical AI applications demands robust engineering combined with domain expertise.

In this talk, Dr. Hatlestad-Hall details how the AI-Mind project leveraged and collaborated with the Services for Sensitive Data (TSD) at the University of Oslo to build a comprehensive research platform for data curation, quality assurance, and AI model development. The talk will highlight critical lessons learned and identify essential focus areas for the future co-creation of research infrastructure that bridges the gap between clinical research and technology development.

...

AI and storage

The "AI and Storage" track explores the potential of artificial intelligence and machine learning to revolutionize data storage and collaborative services.

Highlighted topics include:

AI-driven optimization of storage systems

Innovative data management strategies
The role of machine learning in enhancing storage performance and efficiency
AI-enhanced file synchronization
Intelligent sharing protocols
Optimization of data accessibility across multiple devices
Automated data organization
Security measures powered by AI
Real-time collaboration tools that improve user experience.

FAIR and Open Data, Research Data Lifecycle, Data Science Environments & Preservation

EFSS services play a role in the research data lifecycle workflow, from data acquisition, collection, processing, analysis to publication, preservation, archiving and Open Data.

This track is intended to discuss integration between EFSS and other research services to cover the full data lifecycle: interactive notebooks, metadata aggregators, data repositories, preservation services, specialised data viewers, collaboration tools, documentation and more.

How can our services help in implementing and encouraging FAIR practices in real life? What should be the role of human interfaces and machine APIs?

This track is also a place to discuss novel applications and user scenarios which are enabled by the CS3 services with innovative data access and sharing functionality.

Keywords: Repositories, JupyterLab & Notebooks, FAIR, ORCID, OpenAIRE, GPUs, Spark, Analytics, DTN, FTS, Grid.

Applications handling sensitive data, data classification, data security and privacy

There is a growing number of use-cases where handling of sensitive data is key.

This track is intended to discuss research use-cases and applications handling sensitive data, data classification, data security and privacy and the role of EFSS services in this context.

File Sync & Share Solutions and Requirement from the Community

This is the track for software companies developing Enterprise File Sync and Share products (EFSS): evolution and latest releases, planned new features and development roadmap.

Past speakers included: Dropbox, Nextcloud, Owncloud, Powerfolder, Pydio, Seafile, Syncany

We expect to discuss community requirements in connection with the different solutions, such as scalability, technology evolution, interoperability, operational costs and the like.

We invite contributions from software companies and users.

Scalable Storage Backends and Integration with Data Processing

Effective storage solutions are a key factor for successful integration of CS3 services with processing farms for HPC, Machine Learning and Global Science.

In this track we provide a forum for providers and integrators of innovative storage solutions across all storage tiers: low-latency storage for data analysis and machine learning, online and near-line storage for data products and archival storage.

Solutions from vendors and experience from the sites will be discussed.

Collaborative Applications

This track focuses on collaborative applications and techniques to enhance sharing and user experience for Office, Groupware and Productivity.

While applications are typically provisioned fully on-premise some CS3 sites are exploiting hybrid models in which certain applications (e.g. office) are provisioned by external clouds while the storage stays on-premise.

We encourage input from application developers as well as site managers and policy owners.

EFSS and storage federations for EOSC and eResearch infrastructures.

European Open Science Cloud (EOSC) is the European Commission's Digital Single Market strategic initiative to remove virtual borders, boost digital connectivity, and make it easier for researchers, innovators, companies and citizens to access cross-border online content across the European Union.

In addition, there is a large overlap between CS3 sites and other eResearch Infrastructures (such as WLCG, EGI, SKAO, EU-DAT...) and services resulting from science cluster projects (such as ESCAPE, SSHOC, EOSC-Life,...).

In this session we will discuss the EOSC roadmap, available services and opportunities, as well as: the role of EFSS platforms and EFSS federations such as ScienceMesh in EOSC. co-evolution of EFSS services and EFSS federations and eResearch Infrastructures.

Contributions are encouraged from all stakeholders.

Interoperability: protocols, APIs, OpenCloudMesh (OCM)

The CS3 community has been the laboratory where prototyping and building federations of EFSS systems and efforts to ensure interoperability between EFSS vendors and systems were started and progressed.

These activities have revolved around technologies such as OCM, CS3APIs and Reva. The feasibility of these efforts hinges on the robustness and long-term stability and maintenance of these protocols and APIs.

Recent developments such as the progress of OCM in the IETF and the recognition of OCM as a foundation of the EOSC federation are good examples of progress started at CS3, which this conference wishes to support through this track; is is therefore dedicated to contributions that deal with the current status of development, standardization and testing and future of these protocols and initiatives, from a technology and policy perspective.

Technology & Research

The traditional CS3 track hosts technical presentations of the building blocks of CS3 services, as:

- Algorithms and protocols for file sync and sharing
- Sharing and metadata semantics
- Service reliability and data integrity
- Business Continuity and Disaster Recovery
- Innovative desktop and mobile integration
- Monitoring and performance analysis
- New user interfaces
- APIs and command-line tools.

CS3 Community Site Reports

There is a growing number of services deployed and operated in the CS3 community. This session is an opportunity to share the current status and plans, user feedback as well as operational experience. This session will provide a sort of family-photograph and a competence map of all CS3 services.

In particular we encourage newcomers to the community to introduce themselves and to present the status of their system and plans.

SIG-CISS meeting of GEANT Association

Co-located event: 17th SIG-CISS meeting

The SIG brings together those who are building/operating R&E clouds and are willing to share strategy, design, deployment, performance optimization, application integration, interoperability, security and other related information, knowledge and best practices as well as participate in joint efforts aimed at addressing needs of academic environment related to building, operating, brokering cloud services and infrastructure and ensuring their interoperability.

More information:

<https://wiki.geant.org/display/CISS/17th+SIG-CISS+meeting>
<https://wiki.geant.org/display/CISS/SIG-CISS+Home>

Lightning talks