Open-E & Fraunhofer SCAI

SCALABLE POWER FOR DATA-INTENSIVE RESEARCH:

FRAUNHOFER SCAI USES OPEN-E JOVIANDSS TO SUPPORT DATA-DRIVEN SCIENTIFIC COMPUTING

Summary

- Fraunhofer SCAI conducts data-intensive research and develops software solutions for complex industrial optimization processes
- Very heterogeneous system landscape with Windows, Linux and macOS clients, including many parallel computing environments
- Proproprietary storage appliance replaced with Open-E JovianDSS for increased flexibility and scalability
- High-availability storage system based on ZFS with snapshots, automatic failover, and integrated backup strategy
- Used for user directories, project data, scientific workloads and backup
- Scalable system architecture, already extended multiple times with additional licenses
- Active use of snapshots and self-service recovery by users

open-e

Fraunhofer SCAI

Requirements and Objectives

About Fraunhofer SCAI

The Fraunhofer Institute for Algorithms and Scientific Computing SCAI excels at combining expertise in mathematical and computational methods. It focuses on developing and implementing innovative algorithms in industrial practice, delivering substantial benefits to its customers and partners. One of the institute's key specialties is the development of professional software.

The ten business areas of Fraunhofer SCAI encompass the entire spectrum from research to customer-specific solutions: Bioinformatics, Optimization, Multiphysics, Fast Solvers, High-Performance Computing, Network Evaluation Technologies, Virtual Material Design, Numerical Data-Driven Prediction, Meshfree Multiscale Methods, and Computational Finance.

The institute maintains academic partnerships with the Institute for Numerical Simulation (INS), the Institute of Computer Science, and the Faculty of Medicine at the University of Bonn. Furthermore, Fraunhofer SCAI collaborates with the Bonn-Aachen International Center for Information Technology (b-it), the Hochschule Bonn-Rhein-Sieg, University of Applied Sciences, and the Hochschule Koblenz, University of Applied Sciences, RheinAhrCampus. Fraunhofer SCAI operates a branch lab at the University of Bonn, located within the INS.

How data storage contributes to research at Fraunhofer SCAI

In its daily research, Fraunhofer SCAI processes vast amounts of data from a wide range of application fields. The IT infrastructure must meet the highest demands in terms of performance, scalability, and platform independence – particularly due to the use of heterogeneous system environments (Windows, Linux, macOS) and highly parallel computing processes.

Data processing is carried out by both internal research teams and in the context of customer-specific projects. The storage and protection of sensitive – and often legally protected – data is essential for the quality of research and project results.

Customer Statement

We needed a system that could grow with our requirements, support different platforms, and still offer high availability – without being tied to expensive hardware.

Nils Spring, IT Staff, Fraunhofer SCAI

Key Challenges Before the Project Launch

Before switching to Open-E JovianDSS, Fraunhofer SCAI used a proprietary storage appliance that increasingly proved to be costly and inflexible – especially due to its dependency on proprietary hardware and limited expansion capabilities. As data volumes grew (e.g., due to machine learning workloads), the system reached its limits.

Another key challenge was the lack of flexibility in heterogeneous user environments: staff from research, development, and operations needed simultaneous access to shared project and user data – each with differing system requirements.

The project focused on building a flexible and high-performance storage system capable of meeting the demanding requirements of data-intensive research environments in the long term.

Objectives of the New Storage Solution:

- ✓ Platform-independent access to shared datasets
- ✓ High performance and availability – especially for data-driven research processes
- ✓ Easy scalability in line with increasing storage requirements
- ✓ Snapshot-based data protection for fast recovery
- ✓ Integration into existing backup infrastructure

Storage Strategy Over Time:

Before 2018

Proprietary storage appliance in use – increasingly costly and hardware-dependent

2018-2024

Migration to Open-E JovianDSS. Objective: more flexibility, scalability, and platform independence. Snapshots, self-service restore, and backup integration established.

Today

Stable, high-performance and highly available storage solution with approx. 800 TB of usable capacity. Managed by IT, self-service by users – with room to grow further.

The Implemented Solution: High-Availability, Scalable Storage Cluster

Fraunhofer SCAI deployed an Open-E JovianDSS Active-Active Cluster. The system consists of two high-performance bare-metal servers directly connected to JBODs (each with 44 drives). Each node is equipped with ample RAM (e.g., 500 GB) to fully leverage ZFS caching mechanisms – a critical requirement for data-intensive high-performance computing tasks.

The storage system serves as a central platform for user directories, project shares, and data-intensive workloads (e.g., machine learning). Snapshots taken at short intervals (every 5 minutes, hourly, daily, weekly) enable users to restore data independently. Backup operations are executed locally on the storage server.

Snapshot functionality is actively used by staff – for example, to recover from accidental deletions or configuration errors in project directories. Thanks to integration with Windows (Previous Versions) and direct access via Linux, users can independently roll back data to previous states without involving the IT department.

The system architecture was developed in close collaboration with Open-E's PreSales team, and tailored to the specific needs of data-driven scientific computing and storage-intensive research environments. Fraunhofer SCAI received comprehensive support during hardware selection and sizing. The hardware itself was procured independently by the customer – a clear example of the vendor-agnostic flexibility of Open-E JovianDSS.

Results and Benefits in Daily Use

- ✓ Automatic failover in under 15 seconds in case of hardware or memory failure
- ✓ Snapshot restore by users themselves – relieves the IT department
- ✓ Local backup processing on the storage system
- ✓ Continuously expanded storage environment to accommodate growing data volumes from research projects
- ✓ Seamless replication enabled by ZFS
- ✓ System monitoring via Checkmk, with email alerts for critical events

Customer Statement

With Open-E JovianDSS, we got exactly what we needed: a stable, high-performance system that we can scale at any time – and that simply runs without constant intervention. Even snapshots can be restored independently by our staff if necessary.

Nils Spring, IT Staff, Fraunhofer SCAI

System Architecture

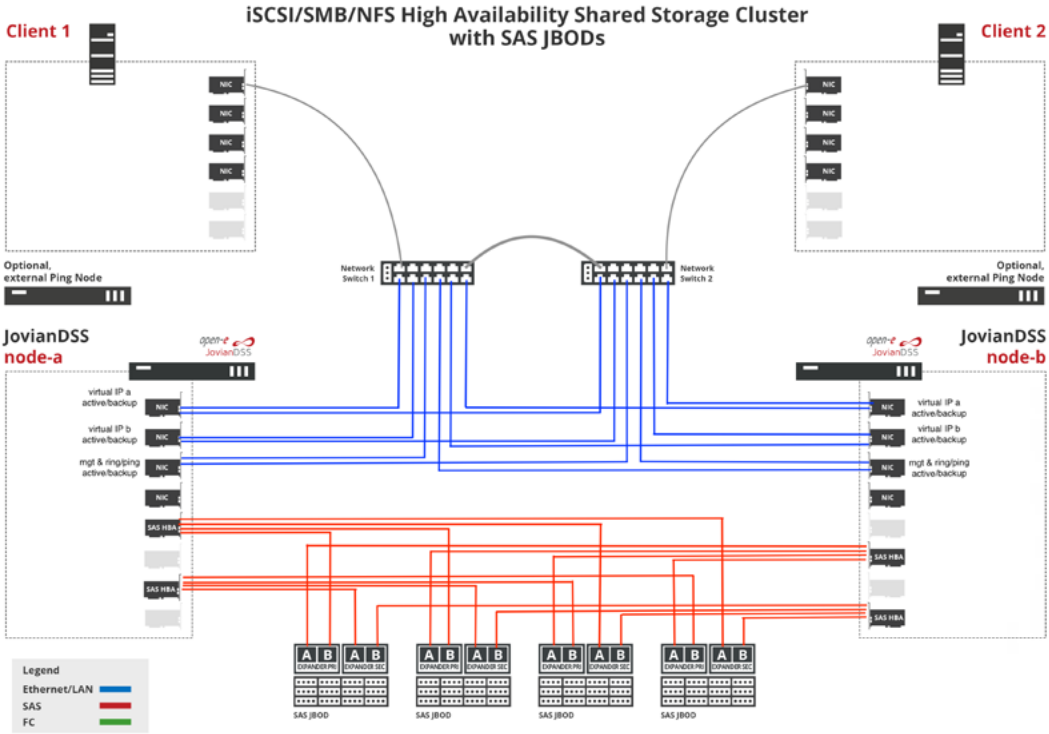

Redundant Cluster Architecture for Maximum Availability

Active-active setup with direct SAS connection and automatic failover – **ideal for uninterrupted storage operations in scientific computing environments.**

Snapshot Function and Data Access Optimized for Research

ZFS snapshots with self-service recovery and high-performance network structure – **ensuring maximum efficiency when working with machine learning and simulation data.**

Hardware details

Open-E JovianDSS Node (Core Components per Cluster Node)

Component	Model	Quantity
Chassis	Supermicro 826BE1C4-R1K23LPB	2
Motherboard	Supermicro H12SSL-NT	2
Network	Supermicro AOC-S40G-12Q	3
CPU	AMD EPYC 7282	2
ECC RAM	8 × 64 GB Samsung DDR4-3200 CL22 ECC (512 GB)	2
Boot-SSD	Toshiba SSD XG6 NVMe 256 GB	2
Cache-SSD	Samsung SSD PM983 NVMe 3.84 TB	2
Write-Log-SSD	Intel Optane DC P4801X	2
HBA	Broadcom HBA 9400-16e Tri-Mode	2

JBOD

Model	Capacity per HDD	HDD per JBOD	Total
Supermicro 847E2C-R1K23	10 TB	44	4 JBODs
			≈ 1.7 PB gross capacity

*total raw capacity of the cluster

Switch

Model	Purpose	Quantity
Cisco Nexus 9504	Supermicro 826BE1C4-R1K23LPB	2
Cluster-Netz	Supermicro H12SSL-NT	2

UPS

Component	Model / Type	Assigned to
USV	Data Center UPS	Node A
USV	Data Center UPS	Node B


Image credit: Fraunhofer SCAI

Data Storage Licenses

The licensing of the system with Open-E JovianDSS was implemented in several stages, aligned with the increasing storage demands of new projects and research data. Both cluster nodes were gradually expanded – first to 320 TB, and later to 640 TB per node.

License Development – Open-E JovianDSS per Node

Node A

Year	License Type	Capacity After Expansion
2017	Base license	64 TB
2017	256 TB capacity upgrade	320 TB
2017	HA Cluster Feature Pack	
2018	256 TB capacity upgrade	576 TB
2018	64 TB capacity upgrade	640 TB

Node B

Year	License Type	Capacity After Expansion
2017	Base license	64 TB
2017	256 TB capacity upgrade	320 TB
2017	HA Cluster Feature Pack	
2018	256 TB capacity upgrade	576 TB
2018	64 TB capacity upgrade	640 TB

Key Details for Fraunhofer SCAI

To support its data-intensive research activities, Fraunhofer SCAI implemented a high-availability, scalable storage system based on Open-E JovianDSS. The solution is specifically designed for data-driven scientific computing and fast data access.

- **Active-active cluster with 2 bare-metal servers and direct JBOD connectivity**
- **4 JBODs with a total of 176 × 10 TB HDDs – approx. 1.7 PB gross capacity**
- **Snapshot-based backup and restore concept with IBM Spectrum Protect**
- **Self-service restore by users via Windows and Linux**
- **Network redundancy through 2× Cisco Nexus + IPMI management**
- **Expanded to 640 TB per node via licensing (HA Cluster + add-ons)**

The Open-E Pre-Sales team supported Fraunhofer SCAI in designing the system – individually and hardware-independent, tailored to computing environments with high requirements for speed and availability.

Are you looking for a similar solution? Contact us for a free consultation.

Open-E, Inc. | +1 (678) 666 2880 (US) | +49 (89) 800777 0 (Europe) | info@open-e.com

About Fraunhofer SCAI

The Fraunhofer Institute for Algorithms and Scientific Computing SCAI excels at combining expertise in mathematical and computational methods. It focuses on developing and implementing innovative algorithms in industrial practice, delivering substantial benefits to its customers and partners. One of the institute's key specialties is the development of professional software.

Fraunhofer SCAI is located in Sankt Augustin near Bonn, Germany.

About Open-E

Founded in 1998, Open-E is a leading developer of IP-based storage management software. Its flagship product, Open-E JovianDSS, is a robust storage application known for its outstanding compatibility with industry standards and its ease of use and management. It is also one of the most stable solutions on the market, offering excellent price-performance value.

Thanks to its reputation, experience, and reliability, Open-E is a trusted technology partner for leading IT companies. With over 40,000 installations worldwide, Open-E has earned numerous industry awards.

For more information about Open-E, its products, and partners, visit: **www.open-e.com**.

This case study is published by Open-E GmbH. Contact person: Paweł Brzeżek, Marketing Director, info@open-e.com

More information:

- **Fraunhofer SCAI** → +49 2241 144300 / info@scai.fraunhofer.de
- **Open-E GmbH** → +49 898007770 / info@open-e.com