



Modular Integrated Sustainable Datacenter

Containerized distributed edge datacenter

Deliverable D5.4

Deliverable:

Work-package: 5.4

Lead author: Eurofiber Innovations B.V.

Contributors: Team WP 5

Due Date: 30-06-2025

Submission Date: 27-06-2025

Pages:25

Version: 1.0

Type: Public/Internal Report:

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)



**MODULAR INTEGRATED
SUSTAINABLE DATACENTER**



Rijksdienst voor Ondernemend
Nederland

8ra CLOUD-EDGE
CONTINUUM

Table of Contents

Management summary	2
2. Characteristics of Distributed Edge	5
3. The Operation of Distributed Edge	6
3.1 Challenges and Considerations	8
4. Relevant Aspects	10
4.1 Sustainability	10
4.2 Data Sovereignty	11
4.3 Edge Network (Distributed Computing Architecture)	14
4.4 Security	16
4.5 Interoperability	18
5. Use Case Requirements	20
5.1 Use Cases	21
6. Conclusion and Recommendations	23

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

Management summary

The main drivers of the Western European colocation data center market are cloud computing and AI, as both private and public sectors increasingly adopt AI and cloud computing for their workloads. Around 71% of UK government agencies have transitioned from conventional data centers to cloud infrastructure, reflecting a significant shift in operational strategies.

As GenAI experiences explosive growth and can lead to massive power consumption to meet demand, AI colocation data centers are planning to develop a robust computing infrastructure both at the edge and the central data center. Equinix, Vantage Data Centers, Data4, and Virtus Data Centres are among the key colocation operators in Western Europe focusing on building data centers specifically designed for AI workloads. This addresses power grid issues and the high electricity demands inherent to AI workloads by retrofitting existing facilities to meet AI requirements. For example, Data4 plans to adapt its current campuses—including a 180 MW site in Frankfurt and two in Paris (120MW and 250MW)—to be optimally suited for AI needs.

In the Western European colocation market, the major hubs are Frankfurt, London, Amsterdam, Paris, and Dublin (FLAPD). London saw a substantial increase in data center proposals starting in 2023. Investments such as Reef Group's 600MW campus in London's Havering borough, Vantage's 48MW campus, and Global Switch's 3.1MW data center reflect this growth.

Power and land constraints pose major issues in Europe's growing data center ecosystem. Major hubs such as FLAPD have reached their capacity limits, leading to constraints on the power grid and land availability. Dublin, for instance, faces a de facto moratorium due to grid limitations, halting new data center developments. Additionally, land prices in FLAPD regions are rising; for example, property values in Slough have increased by over 40% since 2019, while some parts of London are projected to have no electrical capacity for development until 2035.

The rapid developments—especially in AI—are causing disruptive effects in the data center industry, impacting not just revenue but also its overall structure. On one hand, there will be a shift toward data centers located in regions with low electricity costs and low CO₂ emissions, such as Scandinavia and Spain. On the other hand, data security and ease of access remain critical. This is further driven by European regulations, data sovereignty, and local implementations. The development of distributed AI training is perhaps the strongest force fueling the growth of edge data centers.


Due to current uncertainty in data center colocation strategy—particularly around the rise of AI and mounting political and public pressure—edge data center development has lagged behind expectations. Nevertheless, (the often necessary) low latency of <10ms provided by such centers remains at the top of users' must-have lists.

Edge data centers require far lower investments in equipment, network infrastructure, and operational costs. However, they still consume substantial amounts of electricity for AI applications—electricity that is increasingly scarce or even unavailable.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)



The outlook is optimistic: it is estimated that around 20% of total workloads will end up in edge data centers. This is supported by recent private equity investments in new edge data center specialists.



In an era where data is being generated at an unprecedented pace, the need for efficient data processing and management has never been greater.

Distributed edge computing is a revolutionary approach that promises to transform how data is handled by bringing computing power closer to the source. This innovative method not only improves performance but also offers a scalable solution to meet the growing demands of modern applications.

This type of edge solution operates on a much smaller scale, which not only reduces dependency on the power grid but also allows for flexibility in choosing optimal locations. Because the energy consumption is significantly lower, processing can take place where power is available—and most importantly, where it is green. The heat output of new equipment is substantial and requires efficient liquid cooling techniques with the option to return waste heat to a local grid.

This flexibility lies in the ability to deploy a containerized solution at any location where green energy is available and where residual heat can be usefully applied. This makes it a plug-and-play solution, enabled by appropriate network connectivity.

When implemented within a **distributed model across multiple locations**, it becomes possible to make choices about where, when, and under which conditions processing takes place. Linking multiple distributed edge containers and/or combining them with edge data centers and/or third parties is also feasible.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

This part of the IPCEI program focuses on the **development and validation of HPC use cases**, with (containerized) distributed edge computing at the core.

The program will examine definitions, features, operations, and challenges that lie ahead. The structure is designed to provide a holistic overview—both technical and strategic—with a focus on the benefits of distributed edge computing in a sustainable context.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

2. Characteristics of Distributed Edge

Distributed edge computing builds on the principles of traditional edge computing but adds a new dimension. By linking a network of edge nodes in different locations, it enables collaborative processing and resource sharing, paving the way for new opportunities in data-intensive industries.

What is distributed edge computing?

Distributed edge computing is an advanced evolution of traditional edge computing, designed to meet the demands of an increasingly data-driven world. Unlike conventional edge computing, where isolated edge nodes often operate independently, distributed edge computing leverages a network of interconnected edge nodes.

These nodes jointly handle data processing, storage, and transmission tasks closer to the data source or end user, reducing dependence on centralized data centers. This distributed system enables real-time data processing at the network edge and improves resource efficiency.

A strong example of distributed edge computing can be found in **smart cities**. Here, a network of edge devices and nodes supports applications like autonomous vehicles, where local data processing minimizes latency and ensures real-time decision-making. For example, edge nodes within a city collect raw traffic data, process it locally, and send only essential information to a central system. This reduces strain on cloud computing resources and bandwidth, while ensuring immediate processing for critical tasks such as traffic flow management or workplace safety at construction sites.

Distributed edge computing significantly improves on traditional edge computing by lowering transmission costs, enabling faster response times, and enhancing data storage capabilities. By processing data locally or nearer to the source, the system minimizes data transfers to remote servers or public clouds—a key advantage in scenarios with limited connectivity. This makes it well-suited to enterprise-generated data and IoT devices, which require consistent connectivity and low latency.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

3. The Operation of Distributed Edge

At the core of distributed edge computing lies **collaborative data processing through distributed networks**. Edge nodes form the backbone of this system, enabling real-time processing of data collected from connected devices directly at the edge.

Applications ranging from **predictive maintenance** in manufacturing to **quality control** in power grids benefit from local processing. For example, an edge computing system deployed near a healthcare facility can locally process patient data, detect anomalies using business intelligence tools, and prioritize critical information for centralized storage only when necessary.

This distributed computing model utilizes **fog computing principles** to enhance processing power and optimize computing resources in remote locations. By bringing data storage closer to the data source, edge computing reduces latency while improving system reliability. Whether performing tasks such as analyzing data at retail stores or supporting **augmented reality** applications, this technology is part of an ongoing evolution toward smarter, faster, and more efficient solutions for today's data challenges.

Distributed edge computing remains a crucial component of future technologies, adapting to various needs such as managing remote IoT devices, enhancing internet connectivity in rural areas, or enabling real-time data processing in industrial environments. By combining the strengths of cloud computing with localized computing power, it pushes the boundaries of what's possible in modern network systems.

Distributed edge computing expands the capabilities of standard edge computing by emphasizing **resource distribution** and **connectivity**. Traditional edge setups may rely on stand-alone devices or servers near specific endpoints. In contrast, distributed edge computing actively connects these resources into a unified network that enhances performance and reliability.

One key benefit of this approach is **load balancing**. Distributed systems manage workloads across multiple nodes, ensuring consistent performance even during periods of peak demand.

Additionally, the system's **fault tolerance** is improved by distributing tasks across various locations, allowing operations to continue smoothly even if certain nodes fail or experience disruptions.

Distributed edge computing offers many advantages, particularly in terms of **scalability and efficiency**. By distributing workloads across a network of edge nodes, it ensures that data processing occurs closer to the source, significantly reducing latency. This is especially useful for applications like the Internet of Things (IoT), where devices are widely dispersed and require immediate data processing.

A standout advantage is **optimized bandwidth usage**. By processing data locally, distributed edge computing minimizes the amount of raw data that must be transferred to centralized servers. This not only saves bandwidth but also improves system speed and responsiveness. In the healthcare

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)



sector, for instance, a distributed edge network can enable **real-time monitoring of vital signs** in remote clinics, reducing the reliance on centralized data centers for analytics.

Additionally, distributed edge computing enhances system scalability across different regions. This is essential for IoT applications, where devices are often deployed in diverse and widespread locations. By leveraging a distributed network, organizations can ensure consistent performance and reliability regardless of geographic limitations.

Distributed edge computing is transforming industries by enabling **real-time decision-making** and **efficient data processing**. In IoT, it supports smart homes, factories, and agriculture by allowing sensors to analyze data locally while sharing key insights via connected nodes. This local processing capacity is vital for applications requiring immediate analysis and response.

Another area where distributed edge computing shines is **real-time analytics**. It is used in autonomous vehicles, drones, and robotics to support near-instant decision-making. For instance, connected edge nodes along highways can process vehicle-to-infrastructure (V2I) communications, improving traffic flow and safety. Additionally, distributed edge computing supports **content delivery networks (CDNs)** by caching content across various edge locations, reducing latency for end users and enabling efficient video streaming and online gaming. By bringing content closer to users, CDNs can deliver high-quality services with minimal delays.

It also supports **research and healthcare applications** by processing patient data locally, enabling real-time diagnostics and personalized treatment plans while ensuring data privacy and security.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)



3.1 Challenges and Considerations

Although **distributed edge computing** offers significant benefits, it also presents several challenges that organizations must address to ensure successful implementation.

One of the primary challenges is **network synchronization**. Maintaining consistency and synchronization between interconnected edge nodes can be complex—especially when processing large volumes of data at the edge. It is essential to keep all nodes updated and aligned to preserve system integrity and performance.

Security is another critical issue. The distributed model introduces additional endpoints, which can increase the potential attack surface if not properly secured. Implementing robust security measures—such as zero-trust architectures and edge-native security protocols—is essential to protect data and maintain user trust.

Infrastructure costs may also present a barrier to adopting distributed edge computing. The need for multiple edge nodes and interconnections can lead to higher initial deployment costs. However, these costs may be offset by long-term benefits, including enhanced performance and efficiency.

Centralized control platforms can help manage distributed edge nodes more effectively, reducing some of these challenges by providing streamlined management and oversight.

To successfully implement distributed edge computing, organizations should focus on **promoting collaboration between nodes**. Prioritizing systems that enable seamless communication between nodes can optimize processing and reduce redundancies. This interconnected approach ensures that resources are used efficiently and that data processing is as effective as possible.

Improving security protocols is another essential step. Implementing edge-native security measures tailored to the unique challenges of distributed systems can help defend against potential threats. By developing comprehensive security strategies, organizations can protect their data and preserve the integrity of their networks.

Planning for scalability is also crucial. Starting with a localized edge solution and gradually expanding to a distributed model based on workload requirements and geographic considerations can ensure a smooth transition. For example, by launching an edge computing pilot at a specific location, organizations can analyze performance and make necessary adjustments before scaling up to broader regions using distributed configurations.

Integrating **distributed edge computing with existing systems** requires a strategic approach to ensure compatibility and maximize benefits. A key first step is assessing current infrastructure to identify areas where distributed edge computing can improve performance. This includes evaluating the network edge and determining how edge nodes can be deployed to optimize processing and storage capabilities.

Organizations must also consider the role of **cloud computing** in their distributed edge strategy. By combining the strengths of cloud resources with the local processing power of edge nodes, businesses can create a **hybrid model** that offers the best of both worlds. This approach enables

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

efficient data transfer and storage, allowing essential data to be processed locally while using cloud resources for broader analytics and storage needs.

Another critical aspect is ensuring **network connectivity and bandwidth optimization**. While distributed edge computing can reduce network bandwidth usage through local processing, maintaining reliable internet connectivity is essential to support seamless data transfer between nodes and centralized systems. By optimizing network configurations, organizations can ensure that their distributed edge computing systems operate efficiently and effectively.

The future of **distributed edge computing** is poised for significant growth as organizations increasingly recognize its potential to transform data processing and management. As demand for real-time data processing and instant insights continues to rise, distributed edge computing will play a crucial role in meeting these needs by processing data locally and reducing dependence on centralized data centers.

A key trend shaping the future of distributed edge computing is integration with emerging technologies such as **fog computing** and **IoT devices**. By combining these technologies, organizations can create a more robust and flexible data processing ecosystem that supports a wide range of applications—from smart cities to autonomous vehicles. This integration will enable more efficient data handling and enhance the ability to perform tasks requiring immediate processing and response.

Another trend is the growing focus on edge computing for **small businesses and enterprise applications**. As companies of all sizes seek to harness the power of data, distributed edge computing offers a scalable and cost-effective solution that can be tailored to specific needs. By strategically deploying edge nodes, businesses can optimize data processing and storage, enabling them to make better data-driven decisions.

As distributed edge computing continues to evolve, it will also drive advancements in **network connectivity** and **bandwidth optimization**. By reducing the need to transfer data to centralized locations, distributed edge computing can help minimize network congestion and improve overall system performance. This will be particularly important as the number of connected devices and data sources continues to grow, placing increasing demands on existing network infrastructure.

The future of distributed edge computing is promising, with the potential to revolutionize how organizations process and manage data. It represents a significant shift in how computing power is deployed, bringing it closer to end-users and devices, and offering unmatched benefits in **latency, reliability, scalability, and efficiency**. While challenges remain, the advantages and future potential of distributed edge computing make it a cornerstone of next-generation IT infrastructure.

As we continue to embrace digital transformation, the role of distributed edge computing will only grow, paving the way for a faster, more connected, and intelligent world.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

4. Relevant Aspects

The context in which a **distributed edge compute** environment must operate places high demands on various aspects, each of which will be highlighted individually below.

4.1 Sustainability

Network congestion on the power infrastructure, political considerations, and public pressure make it nearly impossible to build large, central data centers in or near major cities. Modifying existing data centers, on the other hand, is a very costly operation that few parties can afford. Necessary provisions must be thoroughly adapted to future demands, including power capacities >125kW per rack, liquid cooling solutions, returning heat to heating networks, reinforcement of the structure, etc.—all while keeping the data center operational.

This has a **positive effect** on the choice for distributed edge computing at locations where green energy is available and where residual heat can be meaningfully utilized via a heating grid.

Energy consumption and efficiency: Although current data centers deliver seven times more computing power using the same electrical power as five years ago, energy consumption remains enormous and is still increasing. The biggest energy consumers in a data center are cooling systems (about 40%) and the high-bandwidth network connections. Distributed edge computing setups are compact and modular, contributing to energy-efficient solutions. They have a **small footprint** and make efficient use of space, meaning lower cooling demands. They're easily scalable, flexible, and often hot-swappable.

On the other hand, distributed edge computing is often used for **AI/ML/GPU inference**, where resource usage is inherently lower, while AI training happens at a larger scale within edge or regional data centers.

Cooling:

Liquid cooling with the ability to return waste heat to a heating grid is an ideal solution for contributing to sustainability. Many vendors have adopted this method in recent years due to the immense heat generated by the new generation of GPUs.

Green energy and renewable sources:

The role of renewable energy in the sustainable deployment of distributed edge computing is of major strategic importance. This includes wind and solar energy, but also alternative forms.

Reducing latency and network infrastructure:

Minimizing the need for long and/or costly network connections—thereby indirectly reducing the ecological footprint—also contributes to a sustainable solution.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

4.2 Data Sovereignty

Data sovereignty is the principle that organizations retain **full control over their data**, including where it is stored, how it is processed, and who has access to it.

This concept has become increasingly crucial as businesses rely more heavily on digital technologies and external services for managing and sharing data.

A **sovereign cloud** includes one or more of the following six capabilities, depending on geographic, regional, or other legal requirements:

1. **Access restrictions** limiting use of the cloud to specific users, software, systems, and services from a company and its trusted partners, clients, and vendors—potentially restricted to certain geographies or individuals with security clearance.
2. **Organizational control over data location**, such as a specific country, region, or in the provider's or client's own data center.
3. **Compliance with specific government, regulatory, or industry requirements**, including technical specifications and legal, contractual, or operational practices.
4. **Operational support** by the cloud provider that complies with strict expectations regarding employee nationality, security clearances, and residency.
5. **Dedicated network capacity**, ranging from secure VPNs to air-gapped environments fully isolated from the public internet and other customers.
6. **Advanced encryption**, allowing the customer to bring and control their own encryption keys—keys the cloud provider can never access or view.

Drivers Behind Data Sovereignty:

- Growing reliance on digital technologies
- Geopolitical tensions and regulation
- Competition with non-European tech companies
- Protection against cyberattacks and data breaches

Dimensions of Sovereignty:

- **Data Sovereignty:**
 - Control over data storage and processing
 - Compliance with local laws

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

- Protection against unauthorized access
- **Infrastructure Sovereignty:**
 - Control over physical and virtual infrastructure
 - Reduction of reliance on imported technology
 - Increased resilience of critical infrastructure
- **Technological Sovereignty:**
 - Development and control of proprietary technology
 - Encouraging local innovation
 - Autonomy in technological decision-making

Technological Foundation

- Implement **International Data Spaces (IDS)** standards
- Use **connectors** for secure data sharing
- Provide **brokers and catalogs** for business registration

The **IDS architecture** consists of three layers:

1. **Governance layer:** Responsible for policy, certification, and standards governing the entire data space
2. **Software connectors:** The technical link between governance and participants
3. **Participants:** End-users of the system, including data providers and consumers

Version Development

- Continuous improvement of the reference architecture
- **Version 3.0** released in 2019
- **Version 4.0** followed in 2023

European Context

IDS plays a key role in the European Data Space by:

- Developing standards for secure data exchange

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

- Giving full control to participating organizations over their data
- Providing a structured framework for automated negotiations and transactions
- Integrating with broader EU initiatives like **Gaia-X**

This implementation enables European organizations to **collaborate securely and responsibly** on data exchange, while guaranteeing **data sovereignty** at all times.

Practical Implementation

The **European implementation of Sovereignty as a Service** is a revolutionary concept currently being developed under the **ECOFED initiative** (European Cloud services in an Open Federated ecosystem) <https://www.techzine.nl/blogs/infrastructure/559310/ecofed-nederlands-cloudproject-voor-europese-digitale-soevereiniteit>.

This project aims to offer a fundamentally different approach to cloud computing, giving European organizations full control over their **digital assets and infrastructure**.

The implementation is built on three essential pillars of **digital sovereignty**:

1. **Data sovereignty**
2. **Infrastructure sovereignty**
3. **Technological sovereignty**

(Source: [SoftwareONE](#))

Future Outlook

The future of **data sovereignty** is promising, with developments such as:

- **App-store-like platforms** for data access
- **User-friendly infrastructure** for data sharing
- **Network effects** through increased adoption

Political debate is intensifying around the increasing reliance of governments on Microsoft, Google, and AWS for workplace services, email, file sharing, and software development.

This raises growing concerns: shouldn't **Western Europe** possess its **own sovereign cloud and network environment**—a secure system immune to U.S. government surveillance?

Political developments in the U.S. following the election of a new president are cause for concern, increasing the urgency for European action. **Legislation and political pressure** are expected to become major factors driving decisions in this area.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

4.3 Edge Network (Distributed Computing Architecture)

Edge networks have evolved into powerful solutions within the concept of distributed computing, driven by modern application demands for **fast response times**, **low latency**, and **efficient data processing**.

Designing and implementing a **distributed network** that intelligently routes traffic to available resources—based on factors such as location, power cost, or availability—is essential.

An **edge network** is a distributed computing architecture where network, storage, and computing resources are located closer to the data source or end user, rather than depending on a centralized data center. By relocating computing tasks to the edge, the burden on centralized infrastructure is reduced.

Key Benefits of Edge Networks:

- **Low latency:** Enables faster content delivery and minimal data transfer delays—crucial for real-time applications like online gaming or autonomous vehicles.
- **Reduced bandwidth usage:** Less data needs to traverse the core network to central data centers, saving on bandwidth and cost.
- **Reliability:** Edge nodes can continue serving users even if central servers fail.
- **Scalability:** Distributing workloads over multiple nodes reduces the load on any single server and improves distribution.
- **Security:** Sensitive data is less exposed during transmission, thanks to localized processing.


Components of an Edge Network:

- **Edge nodes (servers):**
These are real or virtual servers located close to users or devices. They perform local computing, caching, and data processing tasks—minimizing reliance on centralized data centers. These nodes may be VMs, physical servers, or containers, often deployed in specialized colocation facilities or cloud platforms.
 - **Edge caching:**
This involves storing data—both dynamic and static—at edge nodes to reduce bandwidth usage and delay. Data can be delivered directly from the edge, avoiding repeated fetches from the origin server.
-

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

Regional Load Balancing.

The process of distributing incoming traffic across geographically diverse edge nodes. The goal is to **balance load** and ensure **high availability** by routing users to the closest or most responsive node. This job distribution can also be based on policies like **electricity prices**, **availability**, or **geographic proximity**.

**Edge Data Processing:**

Allows **local computation** and **data handling at the network edge** instead of in centralized data centers. This reduces latency and enables **real-time decision-making**.

Data Synchronization:

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

Ensures **data consistency** across various edge locations. It replicates databases across regions so all edge nodes have access to up-to-date data with minimal latency.

Types of data synchronization:

- **Master-slave replication:** Data is written to a master and replicated to read-only replicas.
- **Master-master replication:** Allows updates at multiple locations, syncing changes across all nodes.
- **Eventual consistency:** Offers less strict consistency but better performance for globally distributed systems.

Edge networks, in combination with **distributed edge computing**, decentralize data processing and bring it closer to the user. With benefits like low latency, improved scalability, reduced bandwidth consumption, and enhanced security, **edge computing is becoming a vital part of modern digital infrastructure**. By leveraging edge nodes, caching, regional load balancing, and real-time synchronization, companies can deliver faster and more reliable services.

As technology advances, **edge networks will continue to shape the future of cloud computing**, enabling smarter, more responsive applications across all sectors.

4.4 Security

Security is a no-brainer—but nonetheless an extremely serious subject.

Edge security is critical for protecting the edge network against a wide range of threats and vulnerabilities. Edge nodes require strong security measures because they are more exposed to the internet and located closer to end users.

Monitoring and observability are essential for resilient edge network maintenance. They ensure that administrators can monitor the health of edge nodes, track performance indicators, and quickly respond to failures or traffic spikes.

Quantum Key Distribution (QKD)

QKD is a **secure communication method** that uses quantum mechanics to generate a shared, random secret key between two parties. This key can then be used to encrypt and decrypt messages.

Unlike traditional cryptographic methods, QKD makes it possible to detect any third party attempting to gain knowledge of the key—thanks to a fundamental quantum principle: **measuring a quantum system inevitably alters it**.

Key principles:

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

- **Eavesdropping detection:** QKD can detect eavesdropping because any attempt to measure the quantum states used in key distribution introduces detectable anomalies.
- **Quantum superposition and entanglement:** QKD uses these principles to encode information in quantum states, allowing any interception attempt to be discovered.
- **Provable security:** QKD security is based on the laws of quantum mechanics and information theory. It offers **forward secrecy** and **provable guarantees**.

Secure Data Processing in Hardware Cores

Another important area is **secure data processing inside the hardware cores** themselves.

The current breakdown of core technologies includes:

1. Quantum Processing Units (QPUs)

- Use quantum states for core processing
- Provide inherent security via quantum mechanics
- Currently limited by noise and error correction challenges
- Most promising for specialized cryptographic operations

2. Trusted Execution Environments (TEEs)

- Hardware-based secure zones for isolated processing
- Total separation from OS and apps
- Examples include Intel SGX and ARM TrustZone
- Already implemented in many modern processors

3. Homomorphic Encryption

- Enables computation on encrypted data without needing to decrypt it
- Preserves data privacy during processing
- Still faces major performance challenges
- Already applied in cloud computing scenarios

4. Secure Multi-Party Computation (SMPC)

- Allows multiple parties to compute jointly without revealing private input
- Complex coordination required

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

- Strongly developing with growing adoption

Practical Implementation Strategy

Based on current maturity levels and feasibility:

Short-term:

- Implement TEEs for immediate security benefits
- Use hardware-based security modules
- Deploy **secure boot mechanisms**

Medium-term:

- Monitor advances in homomorphic encryption
- Evaluate SMPC solutions
- Consider **hybrid solutions** combining multiple technologies

Long-term:

- Invest in QPUs as the most promising technology
- Develop **flexible systems** that can evolve with future innovations

4.5 Interoperability

Interoperability in IT refers to the ability of different systems to operate together and exchange information seamlessly.

This is essential for modern information systems because it ensures that diverse software applications, hardware components, and services can communicate and cooperate effectively.

Types of Interoperability:

- **Syntactic or Technical Interoperability**
Focuses on whether systems can practically communicate and function together. This includes the ability to **replace hardware or software with alternatives** from other vendors while maintaining functionality.
- **Semantic Interoperability**
Focuses on the **meaning** of exchanged data. It aims to prevent misunderstandings when different organizations use different terms for the same concepts.
Example: one organization uses "creator", another uses "author".

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

- **Cross-domain or Organizational Interoperability**

Addresses **societal and organizational aspects**. It involves how institutions such as libraries, archives, or public agencies share data across different domains.

Benefits of Interoperability:

- Improved efficiency through seamless data exchange
 - Better coordination between stakeholders
 - Lower costs through standardized methods
 - Increased innovation potential
-

Key organizations advancing interoperability:

- European Interoperability Framework (European Commission)
- U.S. General Services Administration initiatives
- International Alliance for Interoperability (buildingSMART)

Understanding and applying effective interoperability is **crucial in today's interconnected world**, where systems must seamlessly work together to achieve shared goals while maintaining security, efficiency, and reliability.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

5. Use Case Requirements

In chapter 4, the key values for the structure and operation of the **HPC AI propositions** within the program were outlined—these represent the minimum requirements use cases should meet.

Sustainability, sovereignty, edge networking, security, and interoperability are the cornerstones of the concept within this program.

Technological innovation is progressing at such a rapid pace that we must pause and reflect carefully on the future. The trend initiated by **NVIDIA**, featuring a single rack with **direct-on-chip liquid cooling** consuming roughly **135 kWh**, is rapidly becoming the **de facto standard**. Many manufacturers (such as Supermicro, Gigabyte, and others) are incorporating this into their offerings.

From a European perspective, there is significant pressure to implement this technology **in the most sustainable way possible**, prioritizing green energy and reusing residual heat from data centers.

Simultaneously, rapidly changing **geopolitical developments**, combined with **regulatory** and **social pressures**, demand that we find a solution for **data sovereignty** even sooner than expected.

The **edge network** is essential to deliver a **distributed solution** that meets **latency requirements** in practice.

The emergence of **quantum computing** will increase the demands placed on secure networks and secure data processing.

Interoperability must enable **resource sharing**, allowing for the formation of **larger clusters**.

During the development of use cases, partners in the program will first be consulted to see if they can contribute to one or more of these aspects. If not, external solutions will be sought—preferably **open source**.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

5.1 Use Cases

Medical and academic research institutions generate and process vast volumes of sensitive data, including genomic sequences, imaging scans, and clinical records. A **Distributed Edge Datacenter Architecture** enables localized compute, fast access to AI/ML tools, and real-time data processing — all while ensuring compliance with regulatory frameworks like HIPAA, GDPR, and academic data sharing protocols.

Many of the of the workload type that are specified before can be found within these academic environment.

There is an increasing use of especially extremely powerful HPC/AI and/or GPU by researchers, while the models they use are becoming more complex and a lot of data is being used for training purposes.

The decreasing availability or larger, central environments (which are somewhat outdated) has also a negative influence.

The raising concerns about the public cloud in combination with the rapid development of new hardware, which needs high investments, is not positive too.

Sharing resources with other institutes in a distributed way make the economics of a solution much better

And last but not least the sentiment of having your own environment is still predominantly there.

A solution for them would be a distributed, containerized edge data center

Use Case TU Delft

Focus is on the ability to develop its own **HPC GPU cluster**, in addition to the capabilities already offered by **SURF**. Due to rising demand for GPU power from research projects, current resources are far from sufficient—both now and in the future.

Setup has already been shared as a separate appendix.

Use Case Bytesnet D'Root Groningen

To enable on-premise use of **NVIDIA equipment consuming 135 kWh per rack**, in combination with **direct-on-chip cooling**, several design elements must be re-engineered.

Connection to NVIDIA's **DCU**, with **four feeds**, is required.

Use Case Oost NL

Not yet developed.

Use Case UMCG (University Medical Center Groningen)

An initial exploratory meeting took place with the **Genetics** and **Pathology** departments, along with the manager of the innovative **DASH** (Data Science Center in Health) and the AI portfolio manager.

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

- **Genetics (HPC compute)** and **Pathology (HPC GPU)** are interested in participating. Follow-up discussions are underway.
 - Critical to them is a **super-fast storage solution** capable of running on **bare metal** alongside **OpenStack/Ansible**.
 - **PhD students** could also use the HPC environment for their research.
-

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

6. Conclusion and Recommendations

Distributed edge computing presents a **fundamental shift** in how data is processed, stored, and managed. It empowers organizations to respond to rising **AI demands**, **energy constraints**, and **geopolitical challenges**—while promoting **sustainability**, **security**, and **sovereignty**.

By strategically combining **edge infrastructure**, **containerization**, and **federated governance models**, Europe can build a resilient and autonomous digital ecosystem. Investing now in **scalable**, **secure**, and **sustainable** edge solutions is vital to maintaining **technological leadership** and **data independence** in the years to come.

The emergence of containerized distributed edge datacenters represents a revolutionary transformation in how research centers approach data processing and management. This architectural shift fundamentally reimagines the relationship between data generation, processing, and analysis, creating a more responsive, efficient, and adaptable infrastructure for scientific advancement.

The Strategic Imperative

Traditional centralized data center architectures have reached a critical limitation in supporting modern research demands. As research centers increasingly rely on real-time data processing, artificial intelligence, and IoT devices, the latency and bandwidth constraints of centralized systems have become significant bottlenecks.

Containerized distributed edge datacenter address these challenges by bringing computational power closer to where data is generated and analysed, fundamentally changing how research centers operate.

The strategic importance becomes particularly evident in research environments where immediate data analysis is crucial. For instance, in high-energy physics experiments or medical imaging studies, researchers require instant access to processed data rather than waiting for transmission to distant central facilities. By distributing computing resources across multiple locations, research centers can now process data in real-time, enabling faster discovery cycles and more efficient resource utilization.

Operational Transformation

The distributed nature of containerized edge datacenters creates a resilient infrastructure that mirrors the collaborative nature of modern research. Unlike traditional centralized facilities that represent single points of failure, these distributed systems operate in a decentralized manner where no single location holds complete control or data.

This architecture enhances overall system reliability and ensures continuous operation even if one node experiences technical difficulties.

Research centers benefit significantly from this approach as it enables dynamic workload distribution across multiple nodes. When conducting large-scale experiments or simulations,

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

researchers can now allocate computing resources precisely where needed, maintaining high availability and optimal performance levels. This flexibility becomes especially valuable during peak research periods or when collaborating with remote teams.

Integration with Emerging Technologies

One of the most compelling aspects of containerized distributed edge datacenters is their seamless integration with cutting-edge research technologies. They naturally align with the growing adoption of artificial intelligence, machine learning, and advanced analytics in research environments. By processing data locally, researchers can implement real-time analytics and decision-making capabilities that would be impossible with traditional centralized architectures.

This integration extends beyond pure computing capabilities. The modular nature of containerized solutions allows research centers to deploy specialized computing resources exactly where needed – whether in laboratories, observatories, or experimental facilities. This proximity to research activities enables more efficient data collection, processing, and analysis cycles, ultimately accelerating the discovery process.

Cost Efficiency and Resource Optimization

From a strategic perspective, containerized distributed edge datacenters offer compelling economic advantages for research centers. Unlike traditional facilities that require significant upfront investment in infrastructure, these solutions provide a more manageable cost structure. The modular design allows research centers to scale their computing resources incrementally, matching infrastructure growth with actual research demands rather than investing in oversized facilities.

Moreover, the reduced infrastructure requirements of containerized solutions mean research centers can redirect resources previously allocated to facility maintenance toward actual research activities. This strategic reallocation of funds directly impacts research productivity and institutional competitiveness in securing grants and attracting talent.

Future-Ready Infrastructure

Perhaps most importantly, containerized distributed edge datacenters position research centers for future success in an increasingly data-intensive environment. As research methodologies continue to evolve and incorporate more sophisticated technologies, these distributed systems provide the necessary flexibility to adapt to changing requirements.

Whether supporting emerging technologies like quantum computing or expanding IoT sensor networks, the modular architecture ensures research centers can maintain pace with technological advancement without complete infrastructure overhauls.

The strategic significance extends beyond technical capabilities to institutional positioning. Research centers that adopt this architecture demonstrate forward-thinking leadership and commitment to innovation, making them more attractive to top researchers, funding agencies, and collaborative partners. This competitive advantage becomes increasingly important in today's fast-

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)

paced research environment where rapid adaptation to new methodologies and technologies often determines institutional success.

The transition to containerized distributed edge datacenters represents more than a technical upgrade – it embodies a fundamental shift in how research centers approach data management and processing. By distributing computing resources closer to research activities, these institutions can accelerate discovery processes, enhance collaboration, and maintain competitive advantage in an increasingly data-driven research landscape. This architectural transformation isn't merely tactical; it's a strategic imperative that positions research centers for sustained excellence in the decades ahead

This activity is conducted in the scope of Modular Integrated Sustainable Data Center (MISD) project which received funding from the Netherlands Enterprise Agency (RVO) under the 8ra initiative (IPCEI-CIS)