



MISD – WP2 design guidelines

Eurofiber Cloud Infra B.V.
Safariweg 25-31
3605 MA Maarssen

Owner	:Eurofiber
Reference	:20241209/MISD Design Guidelines
Date	:09 December 2024
Version	:Final

Index

1	Introduction	3
1.1	Background	3
1.2	Work Package 2 – Research approach	4
1.3	Work Package 2 – Starting points	4
2	Work Package 2 – Design Guidelines	5
2.1	Environmental Impact Assessment	5
2.2	Energy Efficiency Metrics	5
2.3	Renewable Energy Integration	6
2.4	Power level monitoring	6
2.5	Load Balancing Algorithms	7
2.6	Orchestration Layer	9
2.7	Architecture	10
2.8	Infrastructure Security	11
2.9	Information Privacy	12
2.10	Governance	13
3	Appendixes	14
3.1	Appendix 1 – Project definition	14

1 Introduction

1.1 Background

The ESG ambitions on energy and CO2 reduction is for Eurofiber an important multiyear project, and it is defined in the several stages. The one of the stages is the research on Green Distributed Data Center Initiative (hereafter: **GDDCI**) and how it can contribute on ESG ambitions.

In order to meet ESG criteria's, Eurofiber has defined research and development project which should cover the following capabilities of GDDCI:

1. Applied hardware that would be used for low power computing power and with demonstrable contribution to zero emissions;
2. Data Center edge architecture with physical containers;
3. Datacenter design for replicated storage at scale;
4. Datacenter design for (liquid)cooling system;
5. Datacenter heat exchangers (in the field and in the traditional datacenters);
6. Connectivity between physical containers and datacenters;
7. Power grid for physical containers;
8. Research on low and regulation for power consumption;
9. Research on low and regulation for placing containers in the fields;
10. Virtualization layer for compute power;
11. Orchestration layer/software for the integration between virtualization layer and the electricity providers;
12. The data points that serve as input to the orchestration layer; and
13. The data analytics platform that provides hindsight, insight and foresight.

1.2 Work Package 2 – Research approach

The GDDCI research is divided into five main work packages. The work package 2 (hereinafter: **WP2**) is defined for hardware- and software-layer together with orchestration software. The federated layer with continental scope is not part of this research.

The research approach for WP2 is divided into five main phases:

1. Design guidelines¹;
2. Building blocks definition²;
3. Preliminary research (creating the prototype on conventional hardware³);
4. Research on advanced technology⁴ (based on prototype result); and
5. Orchestration software development⁵.

This artifact provides design guidelines that can serve as input for the defined building blocks and building the prototype.

1.3 Work Package 2 – Starting points

To define the design guidelines, Eurofiber will use the following starting points:

- Environmental Impact Assessment⁶;
- Energy Efficiency Metrics⁷;
- Renewable Energy Integration⁸;
- Power Level Monitoring⁹;
- Load Balancing Algorithms¹⁰;
- Orchestration Layer¹¹;
- Architecture¹²;
- Infrastructure Security¹³;
- Information Privacy¹⁴; and
- Governance¹⁵.

¹ Design guidelines will be used as starting points for the research.

² Defining the fundamental building blocks is essential for a sustainable and environmentally responsible approach to information technology.

³ In order to safeguard and uncover potential limitations in the initial phases of development, preliminary research relies on well-established building blocks and technologies. This method effectively circumvents the challenges associated with adopting cutting-edge, advanced hardware solutions. Once the concept has proven successful, the research will advance to address ESG-requirements. This will involve the utilization of hardware and software designed for minimal power consumption, efficient heat exchange, and enhanced software capabilities.

⁴ Advanced research drives the adoption of open technology, reducing supplier dependency and increasing agility. This research includes advanced cooling and computing solutions for sustainable and efficient datacenter infrastructure.

⁵ It is the path to optimizing energy consumption.

⁶ Incorporate methodologies for assessing the environmental impact of your IT operations.

⁷ Definition of clear and comprehensive energy efficiency metrics.

⁸ Solar and wind power are crucial to supply clean energy for computational workloads.

⁹ The robust systems for monitoring power levels in real time is required and must serve orchestration layer.

¹⁰ Develop load balancing algorithms that prioritize stable power sources for computational tasks.

¹¹ The orchestration layer ensure that workloads are distributed based on power availability.

¹² The architecture is designed to be scalable, portable and resilience ensuring it can adapt to future computational demands.

¹³ Maintain the focus on Physical-, Network- and Hosts-security and privacy.

¹⁴ Maintain the focus on the GDPR and ENISA as cornerstones of EU data privacy law.

¹⁵ Define the key metrics that Eurofiber will address.

2 Work Package 2 – Design Guidelines

The first three design guidelines relate to electric power and are used to achieve a low PUE.

2.1 Environmental Impact Assessment

The Eurofiber impact assessment is based on energy consumption. The scope of WP2 includes computing power, storage, and networking. At each step, the environmental impact assessment must be reviewed.

Guideline 1

ISO 50001 Energy Management and the practical experience of Eurofiber Datacenter architecture is the reference for research.

2.2 Energy Efficiency Metrics

The targeted efficiency in datacenters is divided into four (main) streams:

1. Optimizing workload capacity to influence PUE;
2. Enabling the consumption of sustainable energy;
3. The reduction of CO2 emissions; and
4. The supply of excess energy to heat exchangers in advanced stages.

Adaptive capacity is measured by assessing the consumption of sustainable energy, the reduction in CO2 emissions, and the energy supplied to heat exchangers.

Guideline 2

PUE is measured according to industry standards¹⁶.

Workload optimization is essential to positively impact PUE.

Adopting a sustainable energy and energy exchange strategy within a new architecture is crucial to reduce the consumption of conventional kW and decrease CO2 emissions.

The efficiency metrics relate to entire Eurofiber datacenter facility including the total output of ICT equipment's: racks, physical servers, storage and networking.

¹⁶ IT efficiency (W) = Total energy consumption / Total IT equipment consumption, with for example:
IT efficiency (W) = 1,300 W / 1,000 W = 1.3 PUE.

2.3 Renewable Energy Integration

Solar, wind energy and efficient cooling are crucial for the supply of clean energy for datacenters and computing workloads. The goal is to initially achieve PUE values of 1.3 where part of the energy used for distributed datacenters comes from solar and wind energy. This strategy will not reduce POE, but it will increase responsible energy use from sustainable sources.

Guideline 3

Integrate renewable energy grids to achieve PUE 1.3

2.4 Power level monitoring

Power consumption of compute nodes will be measured by using embedded management tools¹⁷. API commands retrieve data from selected tools for online and offline compute nodes, including previous and current power consumption but also statistics such as CPU utilization and memory usage. The REST API allows the interaction with the retrieving-appliance programmatically to retrieve data and perform management tasks¹⁸. These statistics will be collected to a central location and will be used to identify available nodes for that particular cluster and what the current power consumption is across all compute nodes.

Guideline 4

IT platform management tooling is central management component providing unified API interface and orchestration tool for managing and monitoring the infrastructure.

The hardware for enclosures, servers, storage and network are physical and logical entities that represent the compute managed by the platform management tools.

Orchestration management tooling provides management and orchestration for the compute platform. Optionally: it could also run integrated IT platform management software and acts as the management point

¹⁷ Every major vendor provides such tooling

¹⁸ An example of REST API Summary of the process:
Authenticate with the IT platform management appliance to obtain a session token;
Query the power devices or specific enclosures to get power consumption data;
Query the server hardware for CPU utilization data; and
Query the server hardware for memory utilization data

2.5 Load Balancing Algorithms

The primary function of this building block is to dynamically utilize available green energy from the connected grid to this concept and also to protect asymmetrical load.

In order for the solution to work efficiently, it is required that:

- Multiple containers are available;
- Load must be distributed to optimize efficiency regarding supply and demand;
- Load can be spread across containers; and
- The metrics are gathered using available agents.

The (rebalancing) algorithms serve as a comprehensive set of rules and processes that secure both data collection and workload orchestration. To assign and or move a workload to container, the metrics needs to be defined, along with the decision making process and underlying technology. This building block is divided in two layers: control- and execution-layer.

The control layer is depending on monitoring metrics. The technology architecture of monitoring metrics consists of:

- Collector agents: the agents used to gather data metrics. These agents are natively provided by embedded software, Out Of Band Management or they can be installed after initial setup;
- Data collection layer: it collects data from physical and logical components via automated collector;
- API layer: it provides RESTful interfaces for communication;
- Storage layer: it combines databases and caching systems to efficiently store and retrieve different types of data; and
- Datasets: it defines data structure within the system, including statistics, historical data and system configuration, ensuring consistent data processing and analysis.

The key components for execution layer are:

- Infrastructure services: platform including OS and workload space to serve the scheduling services; en
- Scheduling services: the platform with listener to receive trigger, complete control process and make decision. Then execute the migration by automated process and finally the verification where abort- or commit-status follows.

The figure below shows process steps from data collection to migration execution:




Figure 1. The control and orchestration overview with process flow

Guideline 5

Decision-making algorithm considers three primary factors:

- Availability (**A**): Binary state (1/0) indicating container operational status;
- Preference Rating (**PR**): Scoring range of 1.0-3.0 based on container efficiency; and
- Power Supply (**P**): Available power capacity from renewable sources (measured in kW)

The software (algorithms) should be dynamic and adaptable to changing energy conditions.

Load balancing weight calculation formula = $A * PR * P$

(**A**)vailability * (**PR**)eference * (**P**)ower supply

The load balancing system operates as a dynamic workload distribution mechanism optimized for green energy availability.

Load balancing weight calculation example with formula:

$$A * PR * P = 1/0 * 1.0 - 3.0 * 1kw$$

Values	Output
C #1: 1.987	(1 st in order)
C #2: 1.654	(2 nd in order)
C #3: 1.321	(3 rd in order)
C #4: 0.0	(not available)

2.6 Orchestration Layer

Note: the execution layer, part of Loadbalancing algorithm, as described in the previous chapter 2.5, is integral part or orchestration layer.

Typically, a cloud-like environment has a hierarchical layout of racks, data centers (DCs), and regions. The software stack ensures that workloads are spread across racks within a DC. Within the same DC, 3-way synchronous replication at the storage level is commonly applied, and a form of asynchronous replication is used to replicate data to a different region for disaster recovery and global availability.

If the workload allows, Eurofiber consider using per uDC redundancy and use VM-motion like features to move workloads. However, whenever this is applicable, it will depend very much on the type of tenant(s).

Within the context of WP2, a single physical container can be viewed as a rack, which we will refer to as uDCs. To achieve the same level of availability and durability, a workload **MUST** be replicated to three different micro DCs or physical containers. This replication level is also required for workloads (not storage) between uDCs. Due to the implicit need for workload mobility, initially, the replication level cannot be adjusted.

The orchestration layer needs to know what uDCs can be used to provision a workload. A group of 3 uDCs is equivalent to a traditional DC. Users should be able to select one more DC when they provision workloads. Compared to public clouds, this additional feature allows workload/data placement closer to the end users of said data.




Figure 2. The overview of 3 Unified Data centers is East.DC

As within the public cloud, object storage should be made available to all DCs. However, the resolution of DCs that provide object storage services can be much lower since object storage workloads are less latency sensitive. This also means that the object storage services do not have to be hosted using the uDC approach but could be centralized instead.

Every uDC should have a DNS name following the standard RFC 1035. In the above diagram, a DC called north consists of 3 uDCs in three separate locations. Async replication can be configured to replicate to, for example, "south.dc" which is further broken down into multiple uDCs.

Tenants should be able to create resources via the API and integrate via SAML with various identity providers. API calls can be made to one or more DCs to create VMs. The tenant is responsible for setting up any form of async replication, back-up, and recovery.

As the API request is received, scheduling logic based on whatever heuristic can determine the uDC the workload will initially be scheduled too. Scheduling constraints among others are load and power (usage and availability).

Guideline 6

Using uDCs or the “traditional” Eurofiber Datacenter should be fully transparent to tenants.

Storage and compute must be segregated as the storage services need to be operational without compute.

The uDCs need to be placed in latency domains since latency will dictate storage performance.

2.7 Architecture

The objective is to build architecture as a generic solution for broad consumers (industry and public sector), without penalty on performance compared to the current Eurofiber Private Cloud.

To achieve scalable architecture, it is important to define the building blocks that are part of the distributed ecosystem and limit these to a minimum to keep distribution as efficient as possible.

The software stack consists out of the following components:

1. Instances: these are KVM based virtual machines that have a fixed set of configurations. There are no further constraints on the VM itself in terms of filesystems used;
2. Windows: Windows instances will require netkvm.sys and viostor.sys;
3. Disks, snapshots, and Images:
 - this provides the ability to upload images and deploy them from said images;
 - once deployed the disks can be snapshotted;
 - depending on the storage solution, a retention policy needs to be specified;
4. Object storage target to be used by tenants to backup and restore disks and images;
5. Virtual private network as defined by RFC1918:
 - Multiple VPCs can be created via the API;
 - Firewall rules that can be attached to VPCs;
 - Floating IPs which can be assigned from preconfigured IP pools;
 - Load balancing for public IPs assigned to instances; and
6. Tenants¹⁹ should be able to create multiple projects within their own tenant and delegate access to these projects, and per project resource allocation and alerting.

¹⁹ Overtime additional hosted services can be provided to tenants. For example, services like, K8s, PostgreSQL, MongoDB, KeyDB. The roadmap of these services will be provided based on customer demand.

Guideline 7

The architecture is scalable, portable and resilience.

All uDCS will run the same software, where each uDC registers itself with the provisioning system.

Generic customer architecture (building block) is:

- Virtual Machines (VM) Linux & Windows
- File system within VM
- Database for VM
- Internet access
- Direct connectivity
- Storage tiering & Backup
- Load balancer
- Firewall

2.8 Infrastructure Security

Infrastructure security is addressed in two areas: physical and logical security, according to best practices.

Guideline 8

Eurofiber's Private Cloud must rely on following security domains:

- Secure access
- Datacenter Security
- Hardware Security
- Host Security
- Network Security

Detailed services of Eurofiber's Private Cloud Infrastructure Security:

Secure access	SD-WAN with encryption
	Zero Trust Network Access
	Multi-factor authentication

Datacenter Security	Secure access controls, environmental monitoring, and physical barriers to prevent unauthorized access to facilities
	Optional with Biometric access control systems
	UPS and backup power systems
	Disaster recovery facilities

Hardware Security	Secure boot processes;
	Hardware-based encryption (e.g., Intel SGX or AMD SEV);
	Out-of-band management security (e.g., iLO, iDRAC);
	Secure hardware disposal process; and
	Optional Tamper-evident seals.

Network Security	Threat prevention to block unauthorized access and prevent attacks
	Network traffic monitoring for suspicious activity and raise alerts
	Optional AI-based threat detection
	Network Segmentation
	Micro-segmentation for containers
	VLAN isolation
	Zero Trust architecture
	VPN to securely connect remote users or offices to the private cloud

Host Security	Configuration hardening
	Minimal installation principle
	Privileged access management
	Regular security patching
	Antivirus and Anti-Malware management
	Host-based Intrusion Detection Systems

Container Security	Image scanning
	Runtime protection
	Pod security policies
	Container isolation

2.9 Information Privacy

The information privacy will be addressed by current ISO and ISAE certifications from Eurofiber, and in addition by processes and tools available on Private Cloud.

Guideline 9

Eurofiber's information privacy include:

- Data security measures
- Data minimization policy
- Storage minimization

Data Security	Implement appropriate technical and organizational measures to protect personal data against unauthorized access, alteration, disclosure, or destruction
---------------	--

Data minimization	Collect only the personal data necessary for the specific purpose of processing
-------------------	---

Storage minimization	Store personal data only for as long as necessary for the purpose for which it was collected
----------------------	--

2.10 Governance

The Eurofiber solution must have a traceable report via a scorecard, based on the regulations, to ensure compliance or our goals.

Guideline 10

Eurofiber's key compliance metrics include:

- Energy consumption
- Access control
- Risks assessments
- Policy enforcement
- Vulnerability management
- Incident response procedures
- Audit logging
- Security awareness training

Detailed description of Eurofiber's governance guidelines:

Energy consumption	Continuously tracks and validates the load balancing algorithms, energy consumption patterns and workload migrations to meet both green energy commitments and security requirements
Access control	Implement and assess the access controls and identity privilege management
Risks assessments	Identify, resolve, and periodically analyze emerging cybersecurity threats
Policy enforcement	Ensures continuous monitoring and automated remediation of security controls through defined frameworks, tools, and processes
Vulnerability management	Monitor and address vulnerabilities in the infrastructure
Incident response	Provide supportive processes to incident response teams
Audit Logging	Automatically captures, logs and stores all security relevant events
Awareness training	discuss and assess the effectiveness of current security awareness training programs with employees

3 Appendixes

3.1 Appendix 1 – Project definition

Project definition starts with the stakeholder engagement and project team formation.

The project is structured around five fixed (research)deliverables and to them corresponding deadlines, scheduled between January-2024, and January-2029.

Over the 60-month duration, milestones will be clearly defined, each associated with the required efforts and necessary knowledge profiles to ensure successful achievement.

The five fixed (research) deliveries are:

- 1. Design guidelines (starting points for the research);
- 2. Building blocks definition;
- 3. Preliminary research (creating the prototype on conventional hardware);
- 4. Research on advanced technology (based on prototype result); and
- 5. Orchestration layer research and development

24 - Q1	24 - Q2	24 - Q3	24 - Q4	25 - Q1	25 - Q2	25 - Q3	25 - Q4	26 - Q1	26 - Q2	26 - Q3	26 - Q4	27 - Q1	27 - Q2	27 - Q3	27 - Q4	28 - Q1	28 - Q2	28 - Q3	28 - Q4	29 - Q1	29 - Q2	29 - Q3	29 - Q4
Design Guidelines																							
Building blocks definition																							
Preliminary research																							
Research on advanced technology																							
Orchestration layer research and development																							
