EUCLORA – European Cloud Computing Research Alliance

Beyond Latency – and Federation

Closing Europe's Cloud Efficiency Gap through Open, Coherent Infrastructure Software

Policy – Position Paper

EUCLORA Research Series

v2.0 · 5 January 2026

Table of contents

Authors

Camilla Ley Valentin

EUCLORA / Innomasters

ley@euclora.org

Niels Henrik Sodemann

EUCLORA / Innomasters

nhs@euclora.org

(January 2026)

This paper forms part of the EUCLORA Research Series on Cloud Efficiency and Sovereignty.

Preface note

This paper constitutes the first volume in the EUCLORA Founding Series. Subsequent volumes will define the legal statutes of the European Cloud Computing Research Alliance AISBL (EUCLORA), together with the InnoFabric technical architecture and the operational blueprint for the neutral pilot data-centre network scheduled for launch in 2026.

Publication notes

This document represents the formal policy position of EUCLORA at the time of publication. It sets out the Alliance's foundational technical and operational direction.

The paper is intended for policy, research, and institutional reference and may be cited accordingly.

Executive Summary

Europe's open-source and digital-sovereignty strategies have established a political foundation for technological independence; this paper defines the operational capability required to realise that ambition at infrastructure scale.

For the purposes of this paper, European cloud sovereignty exists when cloud infrastructure remains operable, governable, and evolvable within the European Union – independently of any individual vendor, platform, or external control. Sustaining sovereignty over time is therefore not primarily a question of ownership or market structure, but of operational performance.

Europe's cloud and edge infrastructure is physically mature but operationally fragmented. Hundreds of regional data centres and strong fibre networks represent immense collective investment, yet operational efficiency still lags substantially behind global hyperscalers that integrate compute, network, orchestration, and compliance within a coherent operational cloud fabric. The resulting efficiency–sovereignty gap is measured not in policy ambition, but in watts, utilisation, and cost – and therefore in lost competitiveness, higher energy demand, and unnecessary emissions.

Europe has pursued federation: connecting independent systems through shared rules and APIs. Federation can build trust, but it also multiplies operational overhead. Europe must therefore complement federation with pooling: a shared, open operational cloud fabric in which automation, telemetry, and policy evolve together, allowing improvements to propagate across providers.

EUCLORA – the European Cloud Computing Research Alliance – enables this pooling model. Its open-source operational fabric, InnoFabric, unifies identity, policy, orchestration integration, and telemetry in an architecture that any provider can adopt while retaining full sovereignty over infrastructure and data. Other open-source infrastructure components remain compatible and optional, but the sovereignty-critical capability is anchored in the shared operational fabric, shared test environments, and auditable metrics.

EUCLORA coordinates shared test data centres and measurable efficiency benchmarks, enabling procurement, investment, and ESG decisions to be evaluated against transparent operational metrics. Europe now needs an IMEC-style centre for cloud and infrastructure software where deep engineering is shared pre-competitively and efficiency gains compound across providers.

Key takeaways for policymakers

- Measurable sovereignty – auditable efficiency outcomes and benchmarks.
- Targets the root cause – fragmentation drives Europe's reinvestment disadvantage.
- Practical open framework – shared operational cloud fabric for telemetry, policy, and automation reuse.
- De-risked investment – testbeds and benchmarks support objective funding and tender criteria.

This paper does not propose industrial consolidation or a procurement-led exclusion strategy. It addresses a distinct structural constraint: the absence of shared, open, hyperscaler-grade operational software and test environments through which European operators can collectively develop, validate, and evolve infrastructure capability over time.

1 Introduction

This paper addresses European policymakers, national digital-infrastructure agencies, and cloud-service operators. It explains why latency and governance are no longer Europe's primary bottlenecks – and how shared, open-source operational cloud fabric software can translate existing public and private investment into measurable efficiency gains.

Over the past decade, initiatives such as Gaia-X, national Trusted-Cloud programmes, and the IPCEI Cloud & Edge framework have improved coordination but not efficiency. European providers still require substantially more hardware, energy, and personnel to deliver the same compute output as global hyperscalers. The consequence is structural: higher costs, slower scaling, and limited capacity to compound operational improvements across regions.

This gap manifests across three interconnected layers of performance:

- Infrastructure efficiency – how effectively data centres convert energy and hardware into usable compute.
- Economic efficiency – how that technical output translates into sustainable profitability and reinvestment capacity.
- Perceived efficiency – the value experienced by customers and developers, reflected in ecosystem adoption and developer gravity.

The imbalance is not a result of geography or talent but of software fragmentation. Hyperscalers operate coherent internal stacks in which automation, identity, policy, and telemetry function as a single operational cloud fabric. Europe operates a fragmented landscape of partially compatible stacks and operational practices. Each improvement – whether in orchestration logic, telemetry coverage, or energy scheduling – remains isolated within an individual provider.

Europe's cloud market also functions within an open-trade and regulatory framework. Under these rules, any provider – including non-European hyperscalers – may operate as an EU entity if it complies with European law. The paradox is that hyperscalers have become some of the most compliant actors in the market, largely because of continuous regulatory and political pressure from the EU itself. Over the past decade, they responded by automating compliance inside their operational cloud fabrics – turning legal requirements into software features. European providers, by contrast, often adapted through manual processes and isolated tools. The result is a structural asymmetry: hyperscalers convert compliance into efficiency, while domestic providers experience it as overhead. In a free market, sovereignty cannot rely on exclusion; it must rely on efficiency.

This dynamic reveals a deeper structural truth: without efficiency, no infrastructure sector can remain competitive in the long run. Cost advantages erode, energy consumption rises, and engineering talent becomes trapped in maintenance rather than innovation. Efficiency is not merely a technical metric – it is the compound engine of competitiveness, sovereignty, and sustainability. Protective policy can slow decline, but it cannot indefinitely sustain structurally inefficient operators without imposing escalating costs on customers, taxpayers, and energy systems. Europe must therefore treat software efficiency as a first-order policy objective, on par with data protection and energy security.

The traditional European approach has been federation – linking independent systems through governance frameworks or APIs. While effective for data exchange and trust management, federation also multiplies operational overhead. Five providers still mean five orchestration layers, five monitoring stacks, and five sources of truth. The result is predictable: more interfaces, more integration work, and less efficiency.

EUCLORA proposes to complement federation with pooling – a model in which providers share an open, deterministic operational cloud fabric. Improvements made anywhere in the system become immediately reusable everywhere. Pooling replaces duplication with compounding progress: each gain in automation density, telemetry resolution, or energy optimisation propagates across the network, turning individual innovation into shared efficiency.

At the centre of this architecture lies InnoFabric (see Appendix 11.2), an open-source operational cloud fabric (the combined substrate layer and control layer) that unifies identity, policy, orchestration integration, automation, and observability across cloud and edge domains. InnoFabric's resource naming model (XRN, see Appendix 11.2.2), telemetry schema, and policy layer allow participating providers to measure, compare, and continuously improve operational efficiency while retaining full sovereignty over infrastructure and data. Other open-source infrastructure components remain compatible and optional – including some whose governance or primary operational setup sits outside the EU – but the sovereignty-critical capability is anchored in the shared operational cloud fabric, shared test environments, and auditable metrics.

This logic of shared efficiency has precedent in Europe's own innovation ecosystem. It mirrors the structural model that made IMEC in Belgium a world-leading semiconductor R&D hub: pre-competitive pooling of engineering resources under shared governance, enabled by neutral test environments and shared measurement. Semiconductor R&D and infrastructure software share the same integration challenge: many specialised components must function flawlessly together, and progress depends on environments where changes can be validated at system scale.

EUCLORA applies the same principle to operational software: an open, continuously integrated operational cloud fabric environment in shared test data centres where orchestration logic, telemetry models, and automation frameworks from many contributors can be tested, benchmarked, and deployed under shared standards. In both cases, the goal is the same: to translate diversity of contributors into compounding technical progress rather than duplication.

Europe's strength has always been collaboration. EUCLORA turns that collaboration into code.

2 Context – Reading the 2025 Alliance Roadmap

The European Alliance for Industrial Data, Edge and Cloud issued its 2025 Roadmap to coordinate national initiatives and outline the investment priorities of the IPCEI Cloud and Edge Infrastructure and Services (CIS) programme. The document recognises the strategic need for sovereign infrastructure and sets three broad objectives: strengthening European supply chains, reducing dependency on foreign hyperscalers, and accelerating deployment of edge-to-cloud capabilities across industrial sectors.

In practice, however, the roadmap remains descriptive rather than operational. It catalogues project clusters but stops short of defining a quantitative framework for measuring efficiency or cross-provider interoperability. Many of its milestones refer to governance models or certification schemes rather than to shared automation, telemetry, or energy orchestration software. The result is an emphasis on coordination rather than compounding.

The roadmap's logic reflects Europe's traditional reliance on federation: linking national or sectoral systems through legal and contractual instruments. This approach can protect sovereignty but does not automatically produce efficiency. Without a coherent **operational cloud fabric**, each participant must still maintain its own orchestration, monitoring, and scaling stack – effectively re-building the same operational capabilities multiple times.

From an engineering perspective, the roadmap's three action pillars – data spaces, edge deployment, and sovereign cloud frameworks – correspond closely to the three layers of efficiency described earlier:

1. Infrastructure efficiency – affected by energy usage, hardware utilisation, and automation density at the facility level.
2. Economic efficiency – determined by operational cost per unit of compute and reinvestment potential.
3. Perceived efficiency – shaped by developer experience, interoperability, and market adoption.

Yet none of the roadmap's instruments explicitly measure or connect these layers. Investments risk remaining fragmented: improving infrastructure without addressing software cohesion, or expanding data-space governance without improving runtime automation.

Reading the roadmap through the EUCLORA lens reveals the missing mechanism: a common, open-source operational cloud fabric (substrate layer and control layer) that allows projects under the IPCEI CIS and the Alliance for Industrial Data, Edge and Cloud to share code, telemetry, and metrics rather than merely policy statements. In this sense, EUCLORA does not compete with the roadmap but completes it – supplying the operational means by which Europe's strategic intent can become measurable performance improvement.

3 Physical and Network Reality

Digital sovereignty cannot be asserted in policy alone; it must be expressed through the physical, network, and silicon layers that deliver compute, storage, and connectivity. Europe's cloud ecosystem is geographically distributed and operationally fragmented. Hundreds of regional data centres and fibre-network nodes have been built across the Union, representing an immense collective investment in physical capacity and engineering expertise.

Contrary to common perception, Europe's limitation is not the ability to build infrastructure. The continent has designed and constructed many of the world's most advanced facilities – often using the same engineering firms, contractors, and suppliers that delivered hyperscale campuses for AWS, Google, and Microsoft. Capital also exists within Europe's institutional and industrial base to fund large-scale deployments. What remains missing is the software and silicon substrate that allows these assets to operate with hyperscale efficiency.

The same asymmetry appears in silicon. Europe designs and packages servers but lacks leadership in the advanced data-centre chips that now define computational efficiency – AI accelerators, DPUs, and custom power-management silicon. Without access to these integrated components, European operators rely on imported architectures optimised for foreign cloud ecosystems. Software inefficiency is therefore amplified by a silicon-dependency loop that Europe does not yet control.

Hyperscalers achieve their efficiency not only through scale but through deep silicon–software co-design. Their custom processors, network-interface controllers, and data-processing units expose telemetry hooks and programmable power-management features directly to orchestration layers. This integration allows real-time tuning of voltage, frequency, and thread scheduling based on workload and platform conditions.

By contrast, most European operators rely on off-the-shelf CPUs and accelerators that provide limited visibility into such parameters, forcing orchestration decisions to operate one or more layers above the hardware. Bridging this instrumentation gap is therefore as crucial as increasing fabrication capacity: open telemetry standards must extend down to firmware, buses, and DPUs so that Europe's operational cloud fabric software can fully exploit available silicon capability, regardless of origin.

Network topology adds a further dimension. Europe's long-haul and metropolitan fibre infrastructure is already extensive and high-capacity – indeed, hyperscalers and European operators use the very same fibre routes and optical systems across the continent. The difference lies not in the glass but in the operational control and orchestration. Hyperscalers operate private backbones over the same fibre infrastructure used by European carriers, typically leasing dedicated wavelengths or long-term capacity from them. Software-defined routing and integrated telemetry enable these backbones to deliver deterministic performance and end-to-end visibility across their global regions. European operators could do the same; the capability already exists within national carriers and data-centre networks. The constraint lies not in the optical domain but in the *local compute infrastructure* beneath each cloud region. Within data centres, the local-area network (LAN) must behave as a near-zero-latency fabric where compute, storage, and acceleration units operate as one deterministic system. Only when this inner domain achieves lightning-fast coherence can inter-region networks perform effectively for replication and API transport.

At continental scale, these same principles extend naturally to the wide-area network (WAN). Across cities and borders, the WAN is already fast enough for many cloud workloads: metro-to-

metro round-trip times are typically in the tens of milliseconds across Europe (common for many major city pairs to fall in the ~10–25 ms range, and typically remaining below 100 ms even on longer cross-continent routes), enabling efficient replication, caching, and asynchronous workload distribution, provided applications connect to data stores within their local region. Hyperscalers follow precisely this logic – local compute for user-facing workloads, global replication for durability. Europe can do the same, because the physical backbone is sufficient; what is missing is shared orchestration intelligence that aligns compute placement and network transport across providers (see Appendix 11.4 for empirical latency data and methodological detail).

These realities define the foundation of Europe's cloud challenge. The continent does not suffer from a deficit of capital or engineering capability, but from a deficit of coherent code and controllable silicon. The physical infrastructure is already present; what is missing is an open operational cloud fabric that turns physical distribution into operational unity.

InnoFabric addresses this gap by providing a shared open-source layer for identity, policy, orchestration, and observability. It allows independently owned infrastructures to operate under common telemetry and control semantics, effectively transforming Europe's distributed data centres into a single measurable machine. Sovereignty, in this view, will not be achieved through new construction alone but through shared software and silicon that allow existing assets to act – and improve – as one.

4 Infrastructure Efficiency – Converting Energy and Hardware into Compute

EU-owned operators exhibit a persistent efficiency gap versus global hyperscalers driven by two primary technical factors, which then propagate into capex and opex outcomes.
First, facility overheads (PUE) remain higher: under the EU reporting baseline, the energy-weighted average PUE is ≈ 1.36 for EU-owned data centres, compared with ≈ 1.15 for hyperscale benchmarks. This directly increases electricity required per unit of IT-delivered energy.
Second, fleet productivity (effective utilisation, U) remains lower: hyperscalers use automation and fleet-wide workload placement to keep a larger share of installed capacity doing useful work, whereas EU-owned operators typically deliver less useful compute per installed server and per kWh of IT power.
Taken together (higher PUE and lower U), the uniform baseline model in Appendix 11.6 estimates energy per unit useful compute for EU-owned operators at ≈ 1.8×–3.0× relative to hyperscale (midpoint ≈ 2.1×). At EU-wide scale, this corresponds to ≈ 22–33 TWh/year of avoidable electricity consumption (midpoint ≈ 26 TWh/year), valued at ≈ € 3.5–6.3 bn/year (central estimate ≈ € 4.8 bn/year) at EU non-household electricity prices.
These technical differentials create downstream structural penalties. Lower utilisation implies more installed capacity is required to deliver equivalent useful compute (capex inefficiency), while higher overheads and weaker automation reduce operations leverage (opex inefficiency). The capex/opex implications are discussed further in Section 5, with full assumptions and method traceability provided in Appendix 11.6.
While individual inputs can be debated, the direction of the gap is robust: higher facility overheads, lower effective utilisation, and lower automation leverage necessarily increase energy, capex, and opex per unit useful compute.

4.1 Network Topology and Transport Architecture

Europe's physical network infrastructure is among the densest and most advanced in the world. Long-haul and metropolitan fibre routes interconnect every major city and data-centre cluster across the continent, owned or operated by carriers such as Orange, Deutsche Telekom, Telia Carrier, Colt, GTT, Lumen, and numerous national and regional providers. Hyperscalers and European operators alike rely on this same optical infrastructure: the same fibre routes and optical systems. The glass in the ground is already fast enough.

4.1.1 Optical Layer

At the physical layer, each fibre pair carries multiple wavelengths using dense wavelength-division multiplexing (DWDM). Each wavelength – or lambda – functions as an independent optical channel with capacity typically in the 100-800 Gb/s range, depending on modulation and distance. Hyperscalers typically secure dedicated wavelengths or indefeasible rights of use (IRUs) on carrier fibre, giving them deterministic bandwidth without owning the underlying cable. European operators can and often do the same; the technology and commercial model are identical.

4.1.2 Transport and Routing Layer

Above the optical layer, hyperscalers deploy private backbones built on standard technologies such as MPLS, Segment Routing, and Software-Defined Networking (SDN). These

frameworks allow deterministic routing, traffic engineering, and real-time telemetry collection across global backbones. The distinguishing factor is not the hardware but the tight coupling between network telemetry/traffic engineering and workload orchestration. In hyperscale environments, routing decisions are aware of workload placement and data-replication policies; the same telemetry informs both transport optimisation and service scheduling.

European providers possess all the technical components to replicate this model. What remains missing is a shared orchestration framework that can extend routing and telemetry semantics across ownership boundaries – so that multiple sovereign networks can operate as one logical fabric. This is a central design objective of the InnoFabric operational cloud fabric, which aims to unify network telemetry, workload placement, and policy enforcement across heterogeneous infrastructure.

4.1.3 LAN versus WAN Domains

Performance sensitivity differs sharply between the local and wide-area domains:

4. Local-area (LAN) domain – Inside each data centre or regional cluster, the compute fabric must operate at near-zero latency. CPUs, GPUs, and DPUs communicate over lossless, deterministic networks (RoCEv2, InfiniBand, CXL) where microsecond delays directly translate into wasted silicon cycles. Orchestration, storage, and telemetry must function as a single electrical system.

5. Wide-area (WAN) domain – Across cities and borders, latency budgets in the tens of milliseconds (commonly in the ~10–25 ms range for many major metro pairs in the Appendix 11.4 sample, and higher towards the geographic periphery) are acceptable for asynchronous replication, API transport, and content caching. The critical requirement is coherence: applications should connect to their local data store, while replication occurs transparently across regions. This mirrors the design pattern used by hyperscalers (for example, Amazon DynamoDB, Google Spanner, and Azure Cosmos DB) – local read/write performance with cross-region replication for durability.

4.1.4 Integration with EUCLORA and InnoFabric

InnoFabric's telemetry and orchestration interfaces are designed to accommodate both domains. Within each facility, it exposes real-time metrics for link utilisation, queue depth, and energy profile to the control layer; across facilities, it models aggregate latency and throughput as dynamic resources in the same XRN (eXtended Resource Name) space. This enables policy engines to place workloads intelligently: close to users, near data, and within sovereign jurisdictions – while using inter-region networks only for replication or API transport.

4.1.5 Summary

Europe already owns the fibre and optical capacity required for a sovereign, federated cloud. The bottleneck lies not in bandwidth but in the lack of shared orchestration and telemetry semantics across domains. By aligning network telemetry, optical routing, and workload orchestration through open standards, EUCLORA's architecture can transform Europe's fragmented connectivity into a coherent, measurable backbone for digital sovereignty.

4.2 Quantifying the Efficiency Gap

Part of Europe's efficiency gap stems from structural advantages that software alone cannot immediately offset. Global hyperscalers benefit from scale economics, vertically integrated supply chains, advanced silicon tuned to their operating models (including DPUs and AI accelerators), and power procurement and siting advantages. Yet these factors explain only part of the divergence.

Even under comparable hardware and power conditions, hyperscalers achieve materially higher fleet productivity through coherent operational cloud fabric software – unified telemetry, scheduling, and automation loops that increase effective utilisation and reduce operational overheads. Scale and custom silicon amplify these gains, but software coherence remains the lever Europe can apply immediately and collectively within its existing infrastructure footprint.

As summarised in the Section 4 headline metrics and quantified in Appendix 11.6, the gap is driven by two primary technical factors that propagate into capex and opex outcomes: facility overheads (PUE) and fleet productivity (effective utilisation, U). Higher overheads increase electricity per unit of IT-delivered energy, while lower utilisation increases the installed capacity required to deliver the same useful compute. Together, these differentials create downstream structural penalties: capex inefficiency from excess capacity, and opex inefficiency from higher energy intensity and lower automation leverage. The capex/opex implications are discussed further in Section 5, with full assumptions and method traceability provided in Appendix 11.6.

4.3 Operational cloud fabric Automation and Orchestration Density

At hyperscale, a defining operational efficiency indicator is automation density – typically measured as the number of servers managed per operations/SRE full-time equivalent (FTE). Published case studies and industry analyses report system-to-operator ratios in the thousands at leading hyperscalers; European providers commonly operate at materially lower ratios. Appendix 11.6 therefore uses transparent modelling bands to express this differential.

A central driver is software cohesion. When identity, naming, policy, scheduling, and telemetry share a unified data model and API surface, entire categories of operational work shrink or disappear: manual stitching, brittle integrations, per-service tooling, and duplicated error handling. EUCLORA's InnoFabric architecture targets this cohesion through the eXtended Resource Name (XRN) and a shared telemetry schema across modules. This alignment makes autoscaling, placement, and failover more deterministic, creating the conditions for European operators to improve automation density within existing staffing and infrastructure constraints.

4.4 Software and Telemetry Coherence

Fragmentation is a persistent structural condition in Europe's cloud ecosystem. Many research programmes and grants produce excellent components, but each reinvents its own identifiers, policy models, and telemetry conventions. This makes integration costly and brittle, preventing efficiency gains from compounding across the ecosystem.
InnoFabric addresses this by defining a coherent operational cloud fabric – a shared foundation for identity, policy, orchestration, and observability. It also creates a basis for open-source development to proceed along a coordinated track, rather than through competing projects and disconnected frameworks. Building one coherent system instead of fragmented pieces is imperative if Europe is to achieve hyperscaler-scale automation density.

Telemetry is the other half of coherence. A unified schema for metrics, traces, and events allows automation systems to reason over a consistent view of state, enabling closed-loop operation: scaling up or down, optimising for energy efficiency, and automatically rolling back on regression. This coherence turns cloud operation from a collection of manual tasks into a measurable, continuously improving system.

4.5 Silicon Utilisation and Compatibility

The cloud-efficiency frontier has moved downstream into silicon. Modern utilisation depends on precisely matching workloads to the right compute resources – CPU cores (x86, ARM, RISC-V), accelerators (GPU, NPU, TPU), and interconnect technologies (PCIe, CXL) – and keeping those resources busy with minimal orchestration overhead.

Europe's challenge is structural. The region lacks hyperscaler-scale semiconductor fabrication capacity and vertically integrated processor programmes, and European operators rarely have access to the same capital scale that enables hyperscalers to build end-to-end hardware stacks. As a result, most European providers rely on imported, general-purpose processors, while hyperscalers deploy increasing amounts of custom silicon co-designed for their workloads, achieving higher performance per watt and tighter coupling between hardware capability and orchestration logic. The result is a widening gap in both hardware sovereignty and system-level efficiency.

Addressing Europe's fabrication and industrial-scale capital constraints will require a long-term strategy and targeted EU investment beyond the scope of this paper. EUCLORA's near-term role is to ensure that, regardless of where silicon is designed or manufactured, Europe's operational cloud fabric uses it efficiently and coherently across providers.

InnoFabric therefore treats silicon as first-class metadata. The XRN (see Appendix 11.2.2 – eXtended Resource Name (XRN) Specification) and the operational cloud fabric describe capabilities such as instruction-set architecture, acceleration type, and memory tiers, enabling schedulers to place workloads intelligently. This approach can raise effective utilisation without depending on proprietary chips and helps Europe's compute base remain competitive as new architectures emerge.

4.6 Energy Efficiency and Sovereignty

Energy per workload is a sovereignty metric: Europe cannot be competitive at scale if watts per transaction remain high or opaque. Power efficiency is not purely a facility issue; it depends on software coordination – deciding when, where, and how workloads run relative to available power, cooling capacity, and grid conditions.

InnoFabric integrates energy-awareness directly into the operational cloud fabric. Telemetry ingests power and cooling signals; placement policies can respect energy constraints and carbon objectives; and automation can shift, scale, or defer workloads to minimise energy per unit useful compute while maintaining Service Level Objectives (SLOs) – quantitative performance and reliability goals (for example, response time, throughput, or availability) that define acceptable service quality. Unlike contractual SLAs, SLOs are internal technical benchmarks used by operators to ensure that efficiency gains do not come at the cost of performance.

Under EUCLORA, these efficiency and SLO metrics can be measured consistently across providers and, where appropriate, verified and published – supporting transparency and policy accountability.

4.7 Cross-Layer Inefficiencies

Optimising each layer in isolation leads to system-wide inefficiency. Networks optimise for throughput, storage for IOPS, and compute for utilisation – but without shared identity, policy, and telemetry semantics across providers, these optimisations can conflict: aggressive autoscaling can thrash storage caches, and security controls can disrupt routing or load balancing.

A coherent operational cloud fabric resolves this by unifying identity, policy, and telemetry across layers, making global optimisation both possible and safe. When every component reasons over the same view of state, performance tuning no longer creates downstream instability – it compounds efficiency across the entire system.

4.8 Quantitative Framing

Efficiency must be expressed in numbers. Without quantifiable metrics, Europe cannot measure progress, compare providers, or verify that public investment produces tangible results. Numeric indicators make efficiency auditable, repeatable, and improvable – turning policy objectives into engineering outcomes.

EUCLORA proposes the following initial Key Performance Indicators (KPIs):

- Automation density – servers (or VMs/pods) per operations/SRE FTE.
- Time to scale and time to recover under load or fault.
- Energy per workload – kWh per million requests, per training epoch, per GB processed, or per unit useful compute (where applicable).
- Placement optimality – share of workloads running on best-fit silicon, and the extent to which idle or underutilised capacity is effectively reclaimed through placement and consolidation.
- Policy fidelity – percentage of placements and resolutions that adhere to sovereignty and partition rules without manual override.

These KPIs are objectively auditable when all modules share a common eXtended Resource Name (XRN) and a unified telemetry schema. The InnoFabric operational cloud fabric will support automated KPI extraction per provider, producing standardised metrics that can be aggregated, compared, and published across the EUCLORA framework. Together, these capabilities establish the foundation for transparent, data-driven efficiency governance that Europe can measure, trust, and continuously improve.

4.9 Security, Compliance, and Open Governance

Sovereign infrastructure cannot rely on external compliance checklists alone; it must encode policy, identity, and trust directly in the operational cloud fabric. InnoFabric applies this principle as security and governance by design – embedding verification and transparency into the same operational cloud fabric that delivers efficiency.
Policy enforcement as code. All workload-placement, network, and data-access rules are declared and enforced through the shared operational cloud fabric. Each policy is a signed, versioned object linked to an XRN identifier and evaluated at runtime against live telemetry and system state. This ensures that sovereignty constraints – such as data location, encryption, or

residency – are enforced automatically and continuously rather than retrospectively through periodic audit.

4.9.1 Identity and Attestation

Every component – from container to API endpoint – is issued a cryptographically verifiable identity through Fabric IAM, and is required to present that identity for access and placement decisions. Runtime attestation can be used to confirm that workloads execute only on approved hardware and within authorised jurisdictions. Audit trails are tamper-evident and exportable, enabling independent certification bodies to validate compliance without requiring direct, privileged access to provider internals.

4.9.2 Continuous Verification Loop

Telemetry and policy state feed into a closed feedback loop: deviations trigger automatic remediation, alerts, or quarantine. Security becomes a continuous control function, measurable using the same quantitative discipline applied to performance and efficiency.

4.9.3 Alignment with EU frameworks

InnoFabric's compliance logic is designed to support alignment with EU frameworks including EUCS, NIS2, and CSRD. By embedding SLO and ESG reporting hooks directly into operational telemetry, providers can demonstrate conformity continuously while reducing manual compliance overhead.

4.9.4 Open governance and anti-capture safeguards

EUCLORA maintains the shared operational cloud fabric under an open RFC and review process. Technical standards, interface contracts, and code contributions follow transparent approval workflows with plural oversight from public, academic, and private members. No single vendor, state, or consortium can unilaterally alter governance or critical interface contracts: material changes require documented review and multi-party approval. EUCLORA statutes require open publication of interface definitions, reproducible build pipelines, and conflict-of-interest disclosures for maintainers.

Together these mechanisms ensure that Europe's efficiency fabric is also a trust fabric – secure by architecture, verifiable by data, and governed in the open. By fusing compliance, telemetry, and policy into one coherent operational cloud fabric, EUCLORA demonstrates that sovereignty and security are not trade-offs but properties of the same codebase.

5 Economic Efficiency – Capital Asymmetry and Reinvestment Capacity

Europe's cloud-efficiency gap is not limited to energy or automation metrics; it is reflected in the financial structure of the industry and the ability to reinvest at scale. In 2024, Amazon Web Services (AWS) reported approximately USD 107.6 bn in revenue and USD 39.8 bn in operating income. That operating surplus funds continuous reinvestment in R&D, custom silicon, and next-generation automation systems – compounding AWS's efficiency advantage year after year.

By contrast, OVHcloud – Europe's largest cloud provider – reported FY2024 revenue of approximately EUR 993 m, operating income (EBIT) of approximately EUR 25.7 m (around 2.6 %), and a small consolidated net loss. While OVHcloud and other European operators continue to invest, the scale of their financial headroom remains structurally smaller than that of the hyperscalers, limiting how quickly capital-intensive capabilities can be developed, validated, and deployed across large fleets.

The asymmetry is stark. In rough order-of-magnitude terms, AWS operates at around two orders of magnitude more revenue than OVHcloud, and AWS's annual operating income alone is approximately 40× OVHcloud's total revenue (using like-for-like year figures, before any exchange-rate nuance). High-margin operators accumulate financial and human capital that can be recycled into faster innovation cycles and self-funded expansion. Lower-margin operators, by contrast, face a tighter constraint: much of the organisation's capacity is absorbed by operational maintenance, integration work, and compliance overhead rather than by compounding platform improvements.

Capital and efficiency form a mutually reinforcing loop. Higher margins enable investment in automation leverage, silicon programmes, and deep software integration; those investments, in turn, raise utilisation, reduce overheads, and improve profitability. This compounding cycle has powered the hyperscalers for nearly two decades.

Europe must start the same loop from the efficiency side – using shared operational cloud fabric software and shared test environments to raise operational yield before capital follows. If sovereignty funding is allocated without measurable efficiency gains, it risks perpetuating the imbalance rather than correcting it. The role of EUCLORA is to make efficiency improvements measurable, comparable, and reusable across providers, so that operational progress compounds across the ecosystem rather than remaining isolated within individual operators.

Achieving durable convergence will require sustained, coordinated effort over multiple years, with EU-level support and broad industry participation – aligning reinvestment capacity with operational efficiency through measurable, auditable outcomes.

6 Perceived Efficiency – Developer Gravity, Adoption, and Ecosystem Compounding

Operational and economic efficiency ultimately express themselves as perceived efficiency – the value experienced by customers, developers, integrators, and partners. Hyperscalers translate technical performance and financial scale into trust and developer gravity; many European providers remain caught in a perception gap that reinforces the capital asymmetry described in Section 5. Perceived efficiency is not a branding problem. It is the cumulative outcome of predictable operational behaviour, low-friction integration, and measurable reliability, exposed through coherent tooling and auditable metrics.

A platform with high perceived efficiency becomes the default choice. Developers assume it is resilient, continuously improving, and supported by an ecosystem that "just works". That assumption drives network effects: more workloads, more partners, more tooling, higher utilisation, and more reinvestment. When perceived efficiency is low, the opposite dynamic dominates: adoption slows, integrations remain bespoke, and investment is diverted into maintenance and compatibility work rather than compounding product progress.

6.1 Fragmentation as a structural drag on developer gravity

A core driver of Europe's developer-gravity gap is ecosystem fragmentation. Building a solution that runs across EU providers is often cumbersome because providers expose different operational semantics: identity and naming conventions differ, telemetry and policy models are inconsistent, and operational behaviours (scaling, failover, networking, observability) vary by platform. Even when APIs appear similar, day-2 operations diverge – the part developers and integrators fear most.

This fragmentation imposes three penalties:

1. Integration friction – each additional provider requires bespoke adapters, deployment patterns, logging/metrics translation, and operational runbooks.
2. Risk premium – customers perceive higher operational risk because behaviour is less predictable across environments, and incidents are harder to diagnose without shared telemetry semantics.
3. Ecosystem dilution – ISVs and open-source projects cannot justify deep optimisation for a fragmented target; they optimise for the hyperscaler surface area that yields the largest addressable market.

In effect, Europe asks developers to target a landscape rather than a platform.

6.2 What developers actually adopt: semantics, not slogans

Developers and integrators do not adopt "sovereignty" as an abstraction. They adopt platforms with:

- Predictable operational semantics – stable behaviour under load, repeatable scaling, deterministic failover, and clear limits.

- Uniform tooling surface – coherent SDKs/CLIs, infrastructure-as-code patterns, reference architectures, and debugging workflows.
- Transparent reliability signals – measurable SLOs, consistent incident communication, and verifiable performance history.
- Low switching and deployment friction – portability that is practical across regions and providers, not an aspirational compliance claim.

Where these properties exist, trust accumulates. Where they do not, procurement and engineering teams default to hyperscalers because the hidden cost of uncertainty exceeds any nominal price difference.

6.3 EUCLORA's approach: make trust measurable and portability practical

EUCLORA addresses perceived efficiency by turning Europe's landscape into a platform-like surface through a shared operational cloud fabric. InnoFabric does not require consolidation of ownership or uniformity of underlying stacks. It provides a common layer of semantics and evidence:

- Common resource identity – XRN provides stable, provider-neutral naming and reference semantics across domains.
- Common policy semantics – sovereignty and partition rules become declarative objects evaluated consistently, rather than bespoke contractual interpretations implemented per provider.
- Common telemetry semantics – metrics, traces, and events are expressed through a shared schema so that performance, reliability, and sustainability can be compared and automated across providers.

This creates a single developer-relevant truth: when a workload is deployed under the operational cloud fabric, its operational behaviour becomes more predictable, its observability becomes more portable, and its compliance posture becomes more verifiable.

6.4 From metrics to market confidence

Perceived efficiency improves when trust is converted into data. EUCLORA therefore treats transparency as an engineering output: providers can expose auditable, comparable indicators such as time-to-scale, time-to-recover, automation density proxies, and energy-per-workload indicators, expressed under a common telemetry schema. The result is not "marketing claims", but reproducible evidence that procurement teams, regulators, and developers can evaluate.

This enables two practical mechanisms:

1. Comparable provider profiles – a standardised set of operational and sustainability metrics that allows buyers to compare providers without bespoke benchmarking exercises.
2. Evidence-driven procurement – tender criteria can refer to measurable SLO performance, recovery behaviour, and sustainability metrics rather than to brand or proprietary certification alone.

Over time, this reduces the risk premium attached to EU providers and turns sovereignty from an abstract requirement into an operationally testable property.

6.5 Compounding ecosystems: ISVs, open source, and the integration dividend

Developer gravity compounds when partners can build once and reuse everywhere. With shared operational semantics, ISVs and open-source projects can provide:

- portable integrations for identity, policy, and observability;
- reference architectures that work across providers;
- validated deployment patterns with predictable failure modes; and
- shared runbooks and operational tooling.

This is the integration dividend: reduced duplication transforms effort previously spent on bespoke compatibility into effort spent on capability. It also reduces switching costs created by proprietary tooling loops and shifts ecosystem compounding back towards open interfaces.

6.6 Test environments as adoption infrastructure

Perceived efficiency is strengthened by environments where claims can be validated. EUCLORA's shared test data centres provide a neutral place to:

- run conformance and interoperability tests against the operational cloud fabric;
- benchmark operational KPIs under repeatable load and fault scenarios; and
- validate telemetry, policy behaviour, and recovery semantics before production rollout.

These testbeds turn "works in one provider" into "works across providers", and allow improvements to be verified and propagated rather than re-implemented in parallel.

6.7 Closing the loop

Perceived efficiency is the bridge between engineering output and market adoption. Operational cloud fabric coherence reduces integration friction; shared telemetry and comparable KPIs reduce perceived risk; and shared test environments make interoperability real. Together, these mechanisms convert Europe's distributed infrastructure footprint into a platform developers can target with confidence.

In this sense, developer gravity is not won through slogans or procurement mandates alone. It is won by making Europe's infrastructure predictable to operate, easy to integrate, and auditable to trust – so that sovereignty becomes the by-product of a platform developers actively choose.

7 Lessons from Prior Initiatives

Europe has not been idle in its pursuit of digital sovereignty.

Over the past five years, Gaia–X, IPCEI Cloud & Edge, and numerous national programmes have worked to build federated and sovereign infrastructure layers.

These initiatives succeeded in mobilising political will, building trust among providers, and establishing shared vocabularies such as data spaces and federation services.

However, none has yet produced an operational, continuously evolving software base.

The lesson is simple: governance alone does not create efficiency – software coherence does.

Gaia–X demonstrated the power of branding and convening, and it released valuable open-source components and reference implementations.

Yet it never matured into a broad, continuously maintained operational platform capable of addressing Europe's underlying data–centre efficiency challenge.

Its output remains primarily specifications, schemas, and pilot code rather than a unified operational cloud fabric deployed at scale across providers.

IPCEI Cloud & Edge assembled impressive consortia, but each participant implemented its own stack, resulting in parallel systems that do not interoperate.

Likewise, the EU research framework has funded hundreds of cloud-related projects, yet their outputs rarely persist beyond the grant period.

The recurring challenge is fragmentation: many deliverables, little integration.

7.1 What Federation Is – and Is Not

Federation addresses business and governance logic, not operational logic.

It defines how independent systems exchange trust, policy, or data, but it does not determine how those systems actually run.

Done well, federation prevents Europe from rebuilding the same service logic repeatedly – for example, shared identity frameworks, catalogues, or certification registries.

It can streamline compliance and coordination across sectors, reducing administrative duplication.

Yet federation alone cannot close the efficiency gap that defines Europe's structural disadvantage.

Every federated service still runs inside a physical data centre with its own orchestration, scaling, and monitoring stack.

The underlying operational cloud fabric fragmentation – where automation, telemetry, and energy management remain separate per provider – persists unchanged.

Federation aligns intent; pooling aligns operations.

Table 1: Dimensions

Dimension	Federation	Pooling
Domain	Business and governance logic	Operational cloud fabric logic
Purpose	Coordination and interoperability	Efficiency and automation
Typical output	APIs, schemas, catalogues	Unified telemetry, orchestration, identity models
Efficiency impact	Avoids duplication of services	Eliminates duplication of operations

Federation remains valuable where services are generic and cross-sectoral – such as identity and access management, data-exchange catalogues, sustainability-reporting APIs, or procurement registries.

Federation has given Europe a coherent governance layer: common rules for trust, identity, and compliance across sectors.

What it cannot by itself deliver is operational parity with global hyperscalers.

Pooling complements and completes federation by adding the shared operational cloud fabric software where automation, telemetry, and policy logic evolve together across providers.

It transforms the success of Europe's federated governance model into measurable efficiency – turning coordination into compounding performance.

Efficiency at the hardware, software, and energy layers demands pooling: shared automation frameworks, unified telemetry, and common control logic.

Europe should *federate services but pool systems.*

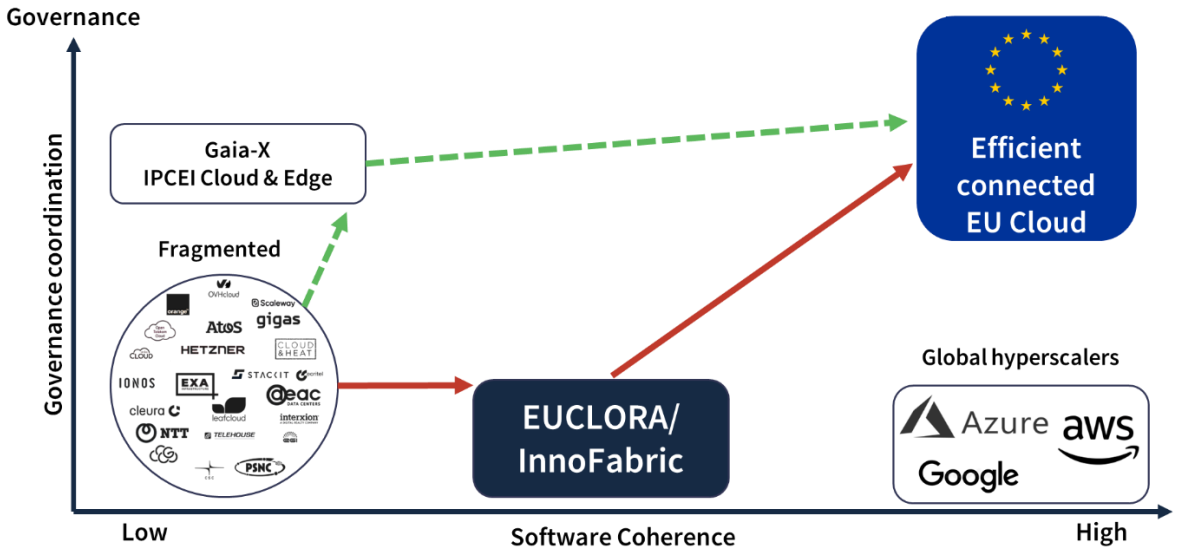

Figure 1 Two European Pathways Toward a Coherent Cloud: Europe's cloud ecosystem can evolve along two complementary routes. Governance-driven initiatives such as Gaia-X and IPCEI Cloud & Edge strengthen coordination, while engineering-driven pooling through EUCLORA / InnoFabric builds shared operational efficiency. Together, they lead to an efficient, connected EU cloud.

EUCLORA draws three direct lessons:

- Efficiency must be measurable – and demonstrated in running code.
- Open governance must accompany continuous integration, not replace it.
- Funding should follow proven reuse and interoperability, not isolated prototypes.

Together, these experiences reveal a deeper structural issue: Europe has tried to integrate before it automates.

Federation was a rational first step to align governance and trust, but it cannot on its own deliver operational parity with hyperscalers.

Pooling complements federation by providing the shared operational fabric that makes federation efficient – a common control layer through which identity, telemetry, and policy logic evolve together across providers.

EUCLORA's substrate architecture requires strong central integration of code and interface contracts.

For complex operational cloud fabric software to function, telemetry models, orchestration APIs, and data schemas must remain consistent across all implementations.

In this sense, the substrate is technically centralised in logic, with one canonical codebase and integration contract.

Yet the infrastructure that runs on top of it remains fully decentralised: each operator deploys the same substrate locally, under its own policy and regulatory jurisdiction.

Governance of the shared code follows an open-foundation model: transparent RFC processes, community review, and plural oversight prevent vendor or state capture.

The result is a system centralised in logic but distributed in control – a necessary balance between efficiency and sovereignty.

Hyperscalers win not by owning data centres but by owning the software substrate that makes all data centres behave as one machine.

Europe keeps trying to design federations of data centres without first building that substrate.

That is why EUCLORA – through its open InnoFabric architecture and shared test data centres – focuses on the operational cloud fabric layer: the true source of efficiency, sovereignty, and composability.

7.2 EuroStack and the Operational Gap

Recent EuroStack proposals (EuroStack, 2025) further reinforce the direction of travel: Europe needs a coherent digital stack that is governable, auditable, and resilient under European jurisdiction. EUCLORA is compatible with this ambition but focuses on a specific missing layer: the operational cloud fabric through which infrastructure efficiency is created, measured, and compounded across providers.

Without a shared operational fabric, "stack" initiatives risk repeating the pattern of prior programmes – strong governance and specification, but limited integration and limited efficiency compounding. Pooling provides the operational mechanism that turns EuroStack-like intent into measurable performance improvement.

8 Conclusion – Building Efficiency Through Collaboration

Europe's digital sovereignty will not be secured through governance or regulation alone. It requires shared, measurable software efficiency – a common operational cloud fabric in which identity, policy, telemetry, and automation evolve together across providers.

The analysis in this report identifies three interlinked efficiency gaps – infrastructure, capital, and perceived value – each rooted not in a lack of talent or resources but in fragmentation: in code, in coordination, and in investment logic. Without efficiency, Europe's cost to deliver compute remains structurally higher than that of global hyperscalers. Transactions, models, and workloads consume more power, require more operational effort, and return less capacity to reinvest.

Federation alone cannot close that gap – it multiplies overhead across parallel stacks and operational silos. Europe cannot compete in the long run from an inefficient base. Efficiency is not a secondary concern but the foundation of sovereignty and competitiveness; without it, governance frameworks have little to sustain.

The consequences are not only economic but environmental. Small percentage losses in efficiency translate into material wasted electricity and avoidable CO_2 emissions at continental scale. Power is the new capital of the digital economy – and it is becoming scarce. In the coming decade, sovereignty will depend as much on conserving and optimising energy as on producing it.

Hyperscalers achieved dominance not by owning data centres, but by operating coherent software that makes large, distributed fleets behave as one system. Europe can apply the same lesson in an open, sovereign way: by building a shared operational cloud fabric that unifies and automates existing infrastructure across providers.

EUCLORA provides that path. By aligning open engineering with quantitative accountability, and by operating shared test data centres where orchestration and telemetry models can be validated, EUCLORA can unify Europe's cloud and edge ecosystem around a single measurable objective: turning public investment into compounding operational efficiency.

The next step is collaborative. European providers, research institutions, and policymakers must now join under the EUCLORA framework to define the technical and governance instruments that will make measurable sovereignty a reality.

Europe's strength has always been collaboration. EUCLORA turns that collaboration into code.

9 References

1. Akamai Technologies. *State of the Internet Report: Performance 2015*. Cambridge, MA: Akamai, 2015. https://www.akamai.com/resources/state-of-the-internet
2. Akamai Technologies. *What Is a CDN (Content Delivery Network)?* Cambridge, MA: Akamai, n.d. https://www.akamai.com/glossary/what-is-a-cdn
3. Amazon Web Services. *Amazon CloudFront Developer Guide*. AWS Documentation, 2024. https://docs.aws.amazon.com/AmazonCloudFront/latest/DeveloperGuide/Introduction.html
4. Amazon Web Services. *Annual Results 2024*. In Amazon.com, Inc. *Form 10-K 2024*. Seattle, WA: Amazon.com, Inc., 2025. https://s2.q4cdn.com/299287126/files/doc_financials/2024/q4/AMZN-Q4-2024-Earnings-Release.pdf
5. Amazon Web Services. *AWS European Sovereign Cloud – Official Information Portal*. Seattle, WA / Luxembourg: Amazon Web Services, 2025. https://aws.eu/
6. Amazon Web Services. *Sustainability Report 2024 – AWS Summary*. Seattle, WA: Amazon Web Services, 2024. https://sustainability.aboutamazon.com/2024-amazon-sustainability-report-aws-summary.pdf
7. Amazon Web Services. *AWS re:Invent (2019–2023)*. Conference presentations and public metrics discussing operational scale. Las Vegas, NV: Amazon Web Services, 2019–2023.
8. AMS-IX. *Realtime Statistics – SLA KPIs (Delay / Delay Variation / Frame Loss)*. Amsterdam: AMS-IX, n.d. https://stats.ams-ix.net/rt-stats.html
9. Cloudflare, Inc. *0-RTT Connection Resumption – Developer Documentation*. San Francisco: Cloudflare, n.d. https://developers.cloudflare.com/speed/optimization/protocol/0-rtt-connection-resumption/
10. Cloudflare, Inc. *HTTP/3 (QUIC) – Developer Documentation*. San Francisco: Cloudflare, n.d. https://developers.cloudflare.com/speed/optimization/protocol/http3/
11. Cloudflare, Inc. *The Cloudflare Global Network*. San Francisco: Cloudflare, 2024. https://www.cloudflare.com/network/
12. Datacenter Knowledge. *Hyperscalers Will Command 60 Percent of Global Data Center Capacity by 2030 – Report*. Datacenter Knowledge, 2025. https://www.datacenterknowledge.com/hyperscalers/hyperscalers-will-command-60-of-global-data-center-capacity-by-2030-report
13. DE-CIX. *Annual Report 2024*. Frankfurt: DE-CIX, 2024. https://www.de-cix.net/en/about-de-cix/annual-report
14. DE-CIX. *DE-CIX Service Levels for DE-CIX Locations (RTT / Jitter Tables)*. Frankfurt: DE-CIX, 19 September 2023 (Version 6.0). (PDF). https://www.de-cix.net/_Resources/Persistent/d/7/6/d/d76d0116dcb2e0287da4f2faa62e41296360dd68/DE-CIX%20Service%20Levels.pdf
15. Deutsche Telekom. *Corporate Responsibility Report 2024*. Bonn: Deutsche Telekom, 2024. https://report.telekom.com/cr-report/2024/environment/energy.html

16. European Alliance for Industrial Data, Edge & Cloud. *Open-Source Way to EU Digital Sovereignty and Competitiveness*. Brussels: European Commission, 2025. https://ec.europa.eu/newsroom/dae/redirection/document/117980

17. European Commission. *IPCEI Cloud & Edge Computing – Project Overview and Objectives*. Brussels: European Commission, 2024. https://digital-strategy.ec.europa.eu/en/library/ipcei-cloud-edge-computing

18. European Commission, Directorate-General for Energy (DG ENER). *Assessment of the Energy Performance and Sustainability of Data Centres in EU: First Technical Report.* Luxembourg: Publications Office of the European Union, 2025. DOI: 10.2833/3168794. ISBN: 978-92-68-29508-3.

19. European Commission. *Trusted Cloud Label – Framework for Secure and Sovereign Cloud Services in Europe*. Brussels: European Commission, 2023. https://digital-strategy.ec.europa.eu

20. European Commission Joint Research Centre (JRC). *Energy Efficiency of EU Data Centres*. Luxembourg: Publications Office of the European Union, 2023. DOI: 10.2760/297157

21. European Commission, Joint Research Centre (JRC). Kamiya, G., and Bertoldi, P. *Energy Consumption in Data Centres and Broadband Communication Networks in the EU*. Luxembourg: Publications Office of the European Union, 2024. EUR 31841 EN (JRC135926). DOI: 10.2760/706491. ISBN 978-92-68-12554-0. https://publications.jrc.ec.europa.eu/repository/bitstream/JRC135926/JRC135926_01.pdf

22. European Data Centre Association (EUDCA). *Energy Efficiency Trends in European Data Centres 2024*. Brussels: EUDCA, 2024. https://www.eudca.org

23. EuroStack. EuroStack White Paper (final, 19 May 2025). EuroStack Initiative, 2025. https://eurostack.eu/wp-content/uploads/2025/08/eurostack-white-paper-final-19-05-25-3.pdf

24. Eurostat. Electricity price statistics. Statistics Explained. Luxembourg: Eurostat, 2024. https://ec.europa.eu/eurostat/statistics-explained/index.php?title=Electricity_price_statistics (dataset: nrg_pc_205).

25. Forsgren, N., Humble, J., and Kim, G. *Accelerate: The Science of Lean Software and DevOps*. Portland, OR: IT Revolution Press, 2018.

26. France-IX. *QoS (RTT / Jitter / Loss) – Paris POP*. Paris: France-IX, n.d. https://tools.franceix.net/qos/par/rtt/sla3.nce1-lyn/sla.dat1-par

27. France-IX. *Statistics*. Paris: France-IX, n.d. https://www.franceix.net/en/infrastructure/statistics/

28. Gaia-X AISBL. *Gaia-X*. Brussels: Gaia-X AISBL, n.d. https://gaia-x.eu

29. GHG Protocol. *ICT Sector Guidance: Built on the Product Life Cycle Accounting and Reporting Standard*. Washington, DC: World Resources Institute (WRI) / World Business Council for Sustainable Development (WBCSD), n.d. https://ghgprotocol.org/ICT-sector-guidance

30. Google LLC. *Environmental Report 2024*. Mountain View, CA: Google, 2024. https://sustainability.google/reports/environmental-report-2024/

31. Google LLC. *Site Reliability Engineering (SRE) Book*. Sebastopol, CA: O'Reilly Media, 2022.

32. Google LLC. *The Site Reliability Workbook: Practical Ways to Implement SRE*. Sebastopol, CA: O'Reilly Media, 2022.

33. Hamilton, J. *On Designing and Deploying Internet-Scale Services*. Seattle, WA: Microsoft, 2007. https://www.mvdirona.com/jrh/TalksAndPapers/JamesRH_LISA2007.pdf

34. Hetzner Online GmbH. *Environmental and Energy Statement 2023*. Gunzenhausen: Hetzner, 2023. https://www.hetzner.com/company/environment/

35. Huang, R., Masanet, E., and Kurnik, C. *Data Center IT Efficiency Measures Evaluation Protocol (Uniform Methods Project – Chapter 20)*. Golden, CO: National Renewable Energy Laboratory (NREL), 2017. NREL/SR-7A40-68576. https://docs.nrel.gov/docs/fy17osti/68576.pdf

36. IDC. *Data Center Operations Efficiency Study 2023*. IDC Market Analysis Report (subscription). Framingham, MA: IDC, 2023.

37. International Energy Agency (IEA). *Electricity Market Report Update 2023*. Paris: IEA, 2023. https://www.iea.org/reports/electricity-market-report-update-2023

38. International Energy Agency (IEA) and Eurostat. *Electricity Information 2023 and Energy Statistics (nrg_cb_e, nrg_d_hhq)*. Paris / Brussels: IEA / Eurostat, 2023–2024.

39. ISO/IEC. *ISO/IEC 30134-2: Information technology – Data centres – Key performance indicators – Part 2: Power Usage Effectiveness (PUE)*. Geneva: ISO, n.d. https://www.iso.org/standard/63485.html

40. Kaffes, K., Sbirlea, D., Lin, Y., Lo, D., and Kozyrakis, C. *Leveraging Application Classes to Save Power in Highly-Utilized Data Centers*. In Proceedings of the 11th ACM Symposium on Cloud Computing (SoCC '20). New York, NY: ACM, 2020. https://dl.acm.org/doi/10.1145/3419111.3421274

41. Knorr, E. *Microsoft exec: We "get" the cloud*. InfoWorld, 2010. https://www.infoworld.com/article/2627010/microsoft-exec--we--get--the-cloud.html

42. LINX. *Traffic Visualisation (SNMP)*. London: LINX, n.d. https://portal.linx.net/services/lans-snmp

43. McKinsey & Company. *Riding the Hyperscaler Wave: The Investment Opportunity in Cloud Ecosystems*. McKinsey, 2024. https://www.mckinsey.com/industries/private-capital/our-insights/riding-the-hyperscaler-wave-the-investment-opportunity-in-cloud-ecosystems

44. McKinsey & Company. *The role of power in unlocking the European AI revolution*. Article, 24 October 2024. https://www.mckinsey.com/industries/electric-power-and-natural-gas/our-insights/the-role-of-power-in-unlocking-the-european-ai-revolution

45. McKinsey & Company. *The State of Cloud Computing in Europe: Increasing Adoption, Low Returns, Huge Potential*. McKinsey, April 2024. https://www.mckinsey.com/capabilities/mckinsey-digital/our-insights/the-state-of-cloud-computing-in-europe-increasing-adoption-low-returns-huge-potential

46. Microsoft Corporation. *Environmental Sustainability Report 2024*. Redmond, WA: Microsoft, 2024. https://www.microsoft.com/en-us/sustainability/emissions-impact-dashboard

47. Netflix, Inc. *Netflix Open Connect: Delivering Content at Scale*. Los Gatos, CA: Netflix, 2023. https://openconnect.netflix.com/en/

48. Netflix, Inc. *Open Connect Briefing Paper – A Cooperative Approach to Content Delivery*. Los Gatos, CA: Netflix, n.d. (PDF). https://openconnect.netflix.com/Open-Connect-Briefing-Paper.pdf

49. Netnod. *IX Statistics*. Stockholm: Netnod, n.d. https://www.netnod.se/ix/statistics

50. OVHcloud. *Annual Report 2024 – Consolidated Financial Statements*. Paris: OVHcloud, 2025. https://investors.ovhcloud.com

51. OVHcloud. *ESG Report 2023*. Roubaix: OVHcloud, 2023. https://corporate.ovhcloud.com/en/sustainability/environment/

52. OVHcloud. *OVHcloud presents its strategic plan, Shaping the Future, and its new financial targets for FY2026 (Investor Day press release)*. Roubaix: OVHcloud, 17 January 2024. https://corporate.ovhcloud.com/en/newsroom/news/investor-day-shaping-future/

53. Queue-it. *Operational case studies and technical blogs on global automation density and uptime practices*. Copenhagen: Queue-it, n.d.

54. RFC-0001 – XRN (eXtended Resource Name) Specification. Innomasters / InnoFabric RFCs, n.d. https://github.com/innofabric/rfcs/blob/main/rfc/rfc-0001-xrn-spec.md

55. RIPE NCC. *RIPE Atlas – Platform Portal*. Amsterdam: RIPE NCC, n.d. https://atlas.ripe.net/

56. RIPE NCC. *RIPE Atlas Anchors – Directory*. Amsterdam: RIPE NCC, n.d. https://atlas.ripe.net/anchors/

57. RIPE NCC. *RIPE Atlas Documentation – Anchors API*. Amsterdam: RIPE NCC, n.d. https://atlas.ripe.net/docs/apis/rest-api-reference/anchors/

58. RIPE NCC. *RIPE Atlas Documentation – Measurement Result Format*. Amsterdam: RIPE NCC, n.d. https://atlas.ripe.net/docs/apis/measurement-result-format/

59. RIPE NCC. *RIPE Atlas Documentation – Measurements API (Ping)*. Amsterdam: RIPE NCC, n.d. https://atlas.ripe.net/docs/apis/rest-api-reference/measurements/

60. RIPE NCC. *RIPE Atlas Statistics – Coverage*. Amsterdam: RIPE NCC, n.d. https://atlas.ripe.net/statistics/coverage

61. Scaleway. *Impact Report 2024*. Paris: Scaleway, 2024. https://www.scaleway.com/en/impact-report/

62. Scaleway. *Understanding Network Latency (Propagation Delay in Fibre)*. Paris: Scaleway, n.d. https://www.scaleway.com/en/blog/understanding-network-latency/

63. Synergy Research Group. *European Cloud Market Forecast 2024*. Market report (subscription). Reno, NV: Synergy Research Group, 2024.

64. Uptime Institute. *Global Data Center Survey 2024*. London: Uptime Institute, 2024. https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2024.GlobalDataCenterSurvey.Report.pdf

65. Verma, A., Pedrosa, L., Korupolu, M., Oppenheimer, D., Tune, E., and Wilkes, J. *Large-scale Cluster Management at Google with Borg*. Proceedings of the European Conference on Computer Systems (EuroSys '15). New York, NY: ACM, 2015. https://research.google/pubs/pub43438/

10 Glossary

Common computing terms (e.g. VM, container, pod) are used in their standard industry sense.

Automation density – Measure of operational automation within a provider, often expressed as the number of servers managed per engineer.

Cloud-sovereignty – Legal and operational control of cloud infrastructure by entities subject to European jurisdiction and policy.

Control layer – The upper layer of the operational cloud fabric that provides federation and sovereignty orchestration – unifying policy, placement objectives, compliance constraints, and lifecycle automation across domains, and applying these controls across the substrate and the underlying cloud platform(s).

Data centre – Physical facility housing compute, storage, and network systems, designed for continuous operation and efficient cooling, power, and security.

Data-space – Federated architecture for secure data sharing across organisations, often defined by sectoral or geographic boundaries.

DPU (Data-Processing Unit) – Programmable accelerator that offloads networking, storage, and security workloads from CPUs.

Edge–to–cloud – Computing model spanning from edge devices near users to centralised cloud data centres, enabling low-latency and distributed processing.

EU-level investment – Funding mechanisms coordinated across European institutions and member states to support collective digital infrastructure initiatives.

Hyperscaler – Large-scale cloud provider operating global data-centre fleets with extreme efficiency and automation (e.g. AWS, Microsoft Azure, Google Cloud).

Open-source – Software released under licences that allow inspection, modification, and redistribution of the source code.

Operational cloud fabric – The combined substrate layer and control layer: a cloud-scale software fabric that delivers efficiency, observability, and governance across distributed compute, storage, and network resources.

PUE (Power Usage Effectiveness) – Industry metric for data-centre efficiency, defined as total facility energy divided by IT equipment energy; lower values indicate higher efficiency.

Pre-competitive collaboration – Co-development between firms prior to market competition, typically on shared infrastructure or standards.

Sovereign cloud – Cloud infrastructure that ensures national or regional control over data governance, access, and compliance.

Substrate layer – The lower layer of the operational cloud fabric that provides the primary efficiency and observability mechanisms – data-path optimisation and offload, storage and network integration, and platform services (e.g., HodogenDB integration) that expose hardware capabilities and reduce overheads for higher-layer orchestration.

Telemetry – Automated measurement and reporting of system metrics (e.g. performance, energy, or carbon data) used for optimisation and governance.

11 Appendices

11.1 EUCLORA – Purpose and Structure

EUCLORA – European Cloud Computing Research Alliance (Established December 2025).

EUCLORA serves as a pan-European, member-based research and coordination alliance dedicated to advancing open, efficient, and sovereign cloud-infrastructure software. Its mission is to unite academic institutions, public providers, private contributors, and EU-level research programmes around a shared, measurable framework for software efficiency and interoperability across Europe's digital infrastructure – and to operate EU-funded test data centres that verify real-world interoperability and performance of open components.

EUCLORA functions as a non-profit research alliance, coordinating standards, reference implementations, and benchmarking frameworks that enable European cloud and edge providers to reach hyperscaler-level efficiency through open and coherent software. EUCLORA's establishment and initial pilot operations are co-funded under EU digital-infrastructure programmes (including Horizon Europe and CEF Digital). The Alliance also oversees the InnoFabric RFC Series – a transparent, community-driven process for defining and validating technical and governance standards in areas such as identity, telemetry, orchestration, and automation.

By linking engineering transparency with policy accountability, EUCLORA aims to make measurable digital sovereignty a practical and verifiable goal across the European cloud ecosystem. EU-supported test data centres host reference workloads under real operating conditions, allowing members to participate both as providers and tenants. These facilities validate telemetry consistency, benchmark operational efficiency, and feed verified performance data back into the shared InnoFabric codebase. This continuous validation loop ensures that improvements to orchestration, telemetry, or automation logic are empirically tested at production scale and quantitatively reflected in EUCLORA's public efficiency metrics and annual benchmark reports.

In its operational capacity, EUCLORA acts as the steward of the InnoFabric ecosystem. The Alliance maintains the InnoFabric roadmap, coordinates the RFC process, and contributes to selected development streams where cross-provider functionality or neutral reference implementation is required. Beyond direct engineering activity, EUCLORA facilitates collaboration among academic, public, and private contributors, ensuring that progress in individual modules compounds into measurable, system-wide efficiency gains.

To ensure independence and accountability, EUCLORA's governance follows a transparent, multi-stakeholder model designed to prevent capture and guarantee public oversight. EUCLORA is governed by a General Assembly representing public research institutions, national and regional cloud providers, and independent experts appointed through open selection. An Executive Board oversees operational execution, with separate Technical and Policy Councils responsible for validating architecture changes, interoperability standards, and compliance metrics. All technical specifications and performance results are published under open-access terms, and decision procedures follow documented RFC and voting processes to guarantee balanced representation of public and private contributors. Financial reporting, project selection, and test-centre results are subject to independent audit and annual publication to maintain trust and neutrality across the European cloud ecosystem.

11.1.1 Pilot Application — DTU InnoFabric Testbed

The initial EUCLORA pilot will establish a controlled single-site testbed at the Technical University of Denmark (DTU), serving as the first full-stack implementation of the *InnoFabric* substrate (see Section 11.2). The facility will include both wet and dry cooling environments, representing liquid-cooled and air-cooled system configurations within the same data-centre footprint, allowing direct comparison of thermal efficiency and telemetry accuracy across cooling methods.

The pilot will deploy the core substrate components described in Section 11.2.1, including the *Registry* for XRN-based resource identification, *InnoDNS* for authoritative and alias-record handling, and the *Telemetry Plane* (see Section 11.3). The telemetry implementation will follow the standard *InnoFabric* model, combining operational instrumentation for workload and automation feedback with energy-domain telemetry for continuous measurement of power draw from connected compute, storage, and cooling systems.

Following each measurement period, verification will compare accumulated power consumption derived from telemetry with actual grid-level readings from certified meters, with subsequent independent audit. This post-event process validates telemetry accuracy and proportionality, ensuring a traceable correlation between reported and measured power efficiency for ESG reporting. Verified results will be incorporated into EUCLORA's public efficiency metrics and annual benchmark reports, forming part of EUCLORA's continuous transparency and accountability framework.

All components will be connected to the public Internet, allowing the DNS layer to serve live queries and validate end-to-end operational coherence. The DTU testbed will provide the first empirical validation of EUCLORA's efficiency-measurement model under real operating conditions. Results will inform subsequent multi-site pilots and contribute to EUCLORA's first *Common Efficiency Baseline*.

11.2 InnoFabric: Open architecture for Hyperscale Efficiency

This appendix defines the open InnoFabric architecture that underpins the EUCLORA operational cloud fabric concept.

InnoFabric is an open, modular software framework developed under the *EUCLORA* to help European cloud and edge providers close their efficiency gap to hyperscalers.

Its mission is to ensure that the logic executing Europe's business logic – the orchestration, placement, and automation software beneath every workload – operates with hyperscale-level efficiency.

Rather than replacing existing platforms, InnoFabric repackages and integrates proven open-source components such as Kubernetes, OpenStack, OpenNebula, and PostgreSQL behind a common set of operational semantics. It introduces shared operational cloud fabric services and interfaces (spanning substrate and control capabilities) that unify automation, telemetry, identity, and sovereignty policy across providers. This serves two purposes: (1) compounding – improvements made by any project or provider become reusable building blocks within a shared framework; and (2) developer gravity – a more uniform operational environment reduces integration friction, allowing developers and ISVs to build once and deploy across participating EU providers with more predictable behaviour and comparable observability.

These substrate services form Europe's open equivalent to the internal control infrastructure used by global hyperscalers – enabling the same degree of operational efficiency, but under transparent, federated governance.

The result is a shared framework through which each provider can achieve measurable, hyperscaler-grade performance, without losing independence or data sovereignty.

11.2.1 InnoFabric Architectural Overview

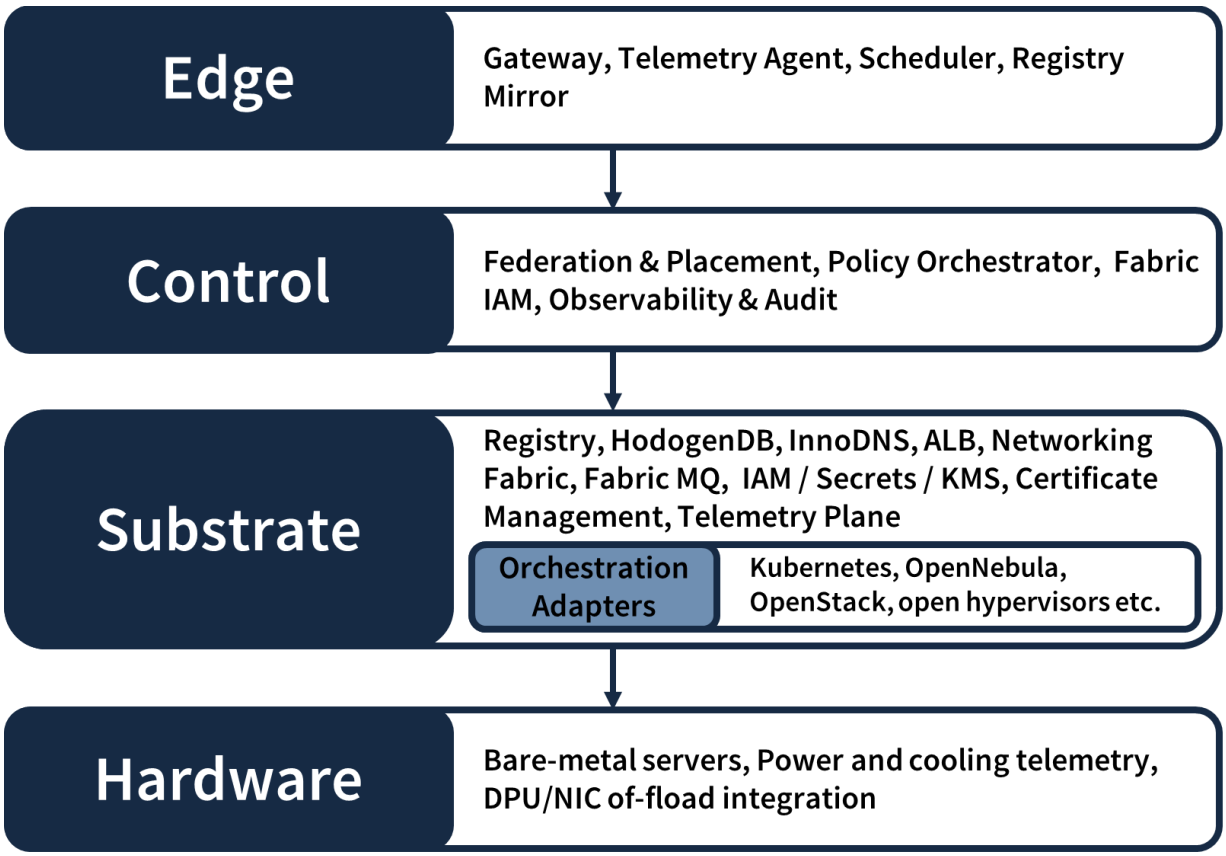

Figure 2: InnoFabric Reference Architecture: Logical layering under EUCLORA governance

The component mapping below reflects current architectural plans for the InnoFabric.

All components, priorities, and design details will be co-developed, validated, and governed collaboratively by EUCLORA members as the initiative progresses.

Table 2: InnoFabric Reference Architecture

Layer	Function	Core Capabilities
Edge Layer	Extends InnoFabric's automation and telemetry closer to users and devices, enabling	Edge Gateway (local control and caching) – Edge Telemetry Agent (active health and energy data collection) – Edge Scheduler (local placement and data-affinity logic)

	distributed, low-latency efficiency.	– Edge Registry Mirror (for offline or intermittent connectivity)
Control Layer	Provides policy, placement, and orchestration intelligence across the federated environment.	Federation & Placement (sovereign placement, data-export controls, latency-based rules, audit trail) – Policy Orchestrator (network/LB policies – weights, canary, failover, stickiness) – Fabric IAM (orgs, projects, roles, API keys) – Observability & Audit (control layer logs, tamper-evident trails)
Substrate Layer (including Orchestration)	Executes and automates workloads on existing open-source infrastructure. Houses the unified data, automation, and security substrate that enables hyperscale-level efficiency.	Registry (XRN objects, versioning, audit trail, API CRUD, event bus) – HodogenDB (working name for managed relational + NoSQL database) – InnoDNS (authoritative DNS, weighted/geo/sovereignty policies, ACME) – ALB (TCP/UDP data path, health checks, failover, autoscaling) – Networking Fabric (VPCs, subnets, routes, SGs, NACLs) – Fabric MQ (NATS-based pub/sub) – IAM / Secrets / KMS (SoftHSM/HSM integration) – Certificate Management (ACME support) – Telemetry Plane (active health checks, endpoint status) – Orchestration adapters for Kubernetes, OpenStack, OpenNebula, and open hypervisors (KVM, Firecracker); optional interoperability with VMware environments for migration and legacy integration; bare-metal lifecycle integration via MAAS (Metal as a Service)
Hardware Layer	Provides the physical compute, storage, and networking foundation managed by the substrate. Includes facility-level energy, cooling, and interconnect	Bare-metal servers – Storage arrays and JBOD / JBOF systems – Top-of-rack and spine switches – Optical and edge routers with BGP / EVPN integration – DPU / SmartNIC offload hardware – Hardware security modules (HSM) and

	systems that enable measurable efficiency across the infrastructure.	trusted platform modules (TPM) – Power and cooling telemetry sensors – Facility monitoring and environmental control interfaces

11.2.2 eXtended Resource Name (XRN) Specification

The eXtended Resource Name (XRN) defines a globally unique, structured, and machine-parseable identifier for all resources participating in the InnoFabric operational cloud fabric.

It provides deterministic identity, traceability, and reversibility across multiple providers, while maintaining strict separation between physical and logical resource domains. XRNs combine human readability with formal parsing stability, serving as the canonical reference for infrastructure, operational cloud fabric, and policy entities within the InnoFabric ecosystem.

XRNs are comparable in purpose to cloud-native resource identifiers such as AWS ARNs or Azure Resource IDs, but they are designed explicitly for shared operational semantics across independent EU providers. This enables third parties to build once – tooling, automation, policy objects, observability integrations, and compliance evidence pipelines – and apply them consistently across participating infrastructures under transparent European governance.

XRNs enable open interoperability across independent providers, research institutions, and national infrastructures, avoiding vendor lock-in while supporting alignment with EU frameworks such as Gaia–X, EUCS, and emerging digital-infrastructure sovereignty initiatives.

Each XRN provides:

- Global uniqueness – every resource has one deterministic identifier derived from its provider, domain, partition, and type.
- Interoperability – common normalisation and encoding rules allow consistent parsing across implementations.
- Reversibility – an XRN can be mapped back to the provider-native identifier through the registry.
- Cross-provider portability – XRN-stable identifiers allow tooling and policy logic to operate predictably across providers within the shared operational cloud fabric.

Source: RFC-0001 – XRN (eXtended Resource Name) Specification, InnoFabric RFC Series – Part of the InnoFabric Standards Track.

11.2.3 Design Principles

- Close the gap – Bring hyperscaler-grade efficiency to every provider through automation, telemetry, and measurable operational semantics.
- • Wrap, don't replace – Build on existing open-source components while adding shared operational cloud fabric services across the substrate layer and control layer.

- • Automate everything – Placement, scaling, routing, and remediation driven by real-time telemetry and health signals.
- • Pool sovereignty – Shared logic and open governance, with independent infrastructure ownership and local jurisdictional control.
- • Measure to improve – Unified telemetry quantifies efficiency, energy use, and latency from edge to hardware, enabling closed-loop optimisation.
- • Open and transparent – 100 % open-source, governed under EUCLORA with auditable interfaces and reproducible builds.

Detailed InnoFabric architecture, RFCs, and source code will be published through EUCLORA as development milestones are completed.

11.2.4 Objective

The InnoFabric stack will demonstrate that European providers can reach hyperscale efficiency using open, federated, and measurable software.

By automating placement, scaling, and telemetry from edge to hardware, InnoFabric establishes the technical baseline for a sovereign, energy-efficient European cloud substrate.

11.3 InnoFabric Telemetry Schema

Telemetry is the unifying data layer of InnoFabric's operational cloud fabric. It serves three tightly connected purposes that together enable hyperscale-level efficiency, reliability, and transparency across providers and tenants:

1. Operational automation and health – Continuous telemetry drives automated placement, scaling, and fault remediation across the operational cloud fabric. It provides the health signals and performance metrics that feed InnoFabric's orchestration and control logic for both provider and tenant workloads.
2. Technical monitoring and observability – Tenants and providers use telemetry for system-level insight: performance metrics, error rates, latency, and resource utilisation. This data supports real-time analytics, troubleshooting, and capacity planning across multi-provider deployments.
3. Efficiency and ESG reporting – The same telemetry pipeline underpins sustainability and compliance reporting. By linking operational telemetry with energy, cooling, and emissions data, InnoFabric enables verifiable efficiency metrics for both providers (facility and fleet footprint) and tenants (workload-level energy use, CO_2 intensity, and renewable ratios).

These functions are intentionally unified under a common telemetry schema. InnoFabric does not separate technical monitoring from sustainability reporting – the same auditable data drives both automation and accountability. This ensures that energy, performance, and reliability signals are measured once, reported consistently, and verifiable end-to-end.

A complete InnoFabric telemetry schema and associated data model will be published through EUCLORA as development milestones are completed.

11.3.1 Telemetry Schema for Efficiency and ESG Reporting

The *InnoFabric Telemetry Schema (ITS)* provides a unified data model for energy, utilisation, and sustainability metrics across all participating providers.

It allows data–centre operators to aggregate machine-level telemetry and correlate it with actual consumption (power, cooling, and network utilisation) for both operational optimisation and ESG reporting.

Under *EUCLORA*, telemetry collection and aggregation are auditable processes aligned with the EU Corporate Sustainability Reporting Directive (CSRD) and GHG Protocol (Scope 2) standards.

Each provider exposes a verified telemetry feed covering power, utilisation, and emissions.

Tenants can query a subset of this feed – limited to their workloads – to support their own ESG and sustainability reporting.

Table 3: Data Model Overview

Entity	Scope	Description
Provider	Aggregated	Reports total and per-region efficiency metrics: power draw, cooling energy, network throughput, CO_2 emissions.
Tenant	Logical / workload	Subset of above, exposing workload-level energy use, CO_2e, and efficiency per compute-hour.
Cluster	Regional / facility	Aggregates server, storage, and network telemetry; feeds facility-level ESG data.
Node	Machine-level	Reports instantaneous metrics: power (W), temperature (°C), utilisation (%), and energy source composition.

This ensures that environmental efficiency becomes a measurable, verifiable part of cloud operations rather than a marketing claim.

11.3.2 Example – Provider-Level Telemetry Record


```
{
 "xrn": "xrn:nimbus:infra:fr-paris-01",
 "start_time": "2025-10-21T14:00:00Z",
 "end_time": "2025-10-21T15:00:00Z",
 "power_kw": 18650.4,
 "pue": 1.48,
 "hardware_utilization_pct": 62.3,
 "energy_mix": {
 "renewable_pct": 82.5,
 "grid_pct": 17.5
```

```
  },
  "carbon_intensity_gco2_per_kwh": 68.9,
  "tenant_efficiency_aggregate": {
  "active_vms": 8231,
  "avg_vm_power_w": 42.7,
  "avg_vm_co2_g_per_hour": 215.6
  },
  "verification": {
  "source": "Facility Sensor Network v3.2",
  "audited_by": "EUCLORA-CERT-ESG-2026"
  }
 }
```

11.3.3 Example – Tenant-Level Telemetry Record

```
 {
  "xrn": "xrn:nimbus:tenant:org1234:workload:12",
  "start_time": "2025-10-21T14:00:00Z",
  "end_time": "2025-10-21T15:00:00Z",
  "energy_kwh": 0.58,
  "co2_g": 39.8,
  "renewable_ratio_pct": 78.3,
  "cpu_utilization_pct": 74.2,
  "network_bytes": 182000000,
  "scope": "tenant",
  "verified_by_provider": true
 }
```

11.4 European Network Latency and Topology Data

Physics sets a hard lower bound on network latency through propagation delay (Table 4) [N1]. In optical fibre, signals propagate at approximately 200,000 km/s; as a practical rule of thumb, each additional 1,000 km adds roughly 5 ms of one-way delay (≈ 10 ms round-trip time (RTT), the time for a packet to travel from sender to receiver and back), before accounting for routing stretch, switching, queuing, and access-network effects [N1]. In practice, end-to-end latency between European metropolitan areas is therefore determined by a combination of physical distance and routing policy, peering topology, and congestion conditions [N1–N3].

Table 4 reports end-to-end IPv4 RTTs derived from RIPE Atlas anchor-mesh ping results over 10–17 Dec 2025, using RIPE Atlas measurement semantics and result formats [N2], with aggregation as defined in Table 5 [N11]. Across the sampled routes, route-level median RTTs range from 8 ms (Amsterdam ↔ London) to 67 ms (Lisbon ↔ Helsinki), and the “Typical Range” is shown as the interquartile (P25–P75) range of pairwise medians [N2, N11]. For many Europe-wide service deployments, these metro-to-metro RTTs are consistent with latency budgets in which application-layer processing, request fan-out, and back-end dependencies can dominate the user-perceived critical path.

Content delivery networks (CDNs) such as Akamai, Amazon CloudFront, Cloudflare, and Netflix Open Connect can reduce user-perceived latency by serving cacheable responses from nearby points of presence, reusing established connections, and using modern transport protocols (for example HTTP/3 over QUIC, including 0-RTT resumption where applicable) to reduce connection-establishment and loss-recovery penalties [N4–N7]. Accordingly, for workloads with a high proportion of cacheable or edge-terminable interactions, “distance-to-compute” within Europe

may be a secondary determinant of perceived performance compared with application and operational cloud fabric behaviour (for example, request fan-out, repeated authentication and policy checks, and cold-start effects).

The remaining bottlenecks are therefore often logical rather than purely geographic: operational cloud fabric chattiness, repeated authentication and policy checks across systems, non-deterministic placement between compute and data, and slow autoscaling or cold-start behaviour. These are primarily software and orchestration constraints rather than propagation constraints. As a result, deploying numerous micro-data-centres at the "edge" without a unified control layer can increase cost and operational complexity, while delivering only limited marginal improvement in user-perceived performance for workloads that are already cache-friendly or otherwise served from nearby points of presence.

Published interconnection service levels and KPI reporting provide independent evidence that major interconnection fabrics can sustain low delay, delay variation (jitter), and loss under normal operating conditions [N3]. Observed variation in end-to-end RTT is therefore frequently consistent with topology, routing policy, and congestion effects, rather than geography alone [N1–N3].

Table 4: Measured Round-Trip Times between Major European Metropolitan Areas

Route	Distance (km)	Median RTT (ms)	Typical Range (ms)
Paris ↔ Frankfurt	≈ 480	10	9–14
Amsterdam ↔ London	≈ 360	8	8–9
Frankfurt ↔ Warsaw	≈ 900	20	18–22
Paris ↔ Warsaw	≈ 1,350	30	27–33
Madrid ↔ Paris	≈ 1,050	21	19–27
Lisbon ↔ Frankfurt	≈ 1,875	40	37–42
Helsinki ↔ Frankfurt	≈ 1,540	28	26–29
Lisbon ↔ Helsinki	≈ 3,370	67	65–68

Distances are great-circle (WGS-84) between city centres (author calculation); RTTs are end-to-end IPv4 round-trip times from RIPE Atlas ping measurements. RTT values are derived from RIPE Atlas anchor-mesh ping results over 10–17 Dec 2025. For each anchor pair, we compute the median RTT across the window using the per-measurement minimum RTT. Route-level median RTT

is the median of the pairwise medians; Typical Range is the P25–P75 range of the pairwise medians. Pair counts are 9 (3×3) unless noted; Lisbon routes use 6 (2×3) due to anchor availability. [N2, N11]

11.4.1 Appendix Reference Mapping (Network)

This Appendix distinguishes between two levels of referencing to ensure both readability and traceability:

1. References section (Section 9**)** – contains the complete bibliographic entries for all cited sources, including document titles and URLs.
2. Appendix reference mapping (this section) – provides a compact mapping from each quantitative statement, benchmark, or dataset used in the Network chapter to the specific source(s) that substantiate it.

Local identifiers [N#] are used throughout Section 11.4 and the associated tables to cross-reference the relevant mapping entries in this Appendix. Where a claim is based on empirical measurements, the mapping identifies both (a) the measurement platform or published KPI source and (b) the methodological definition required to interpret the figures (for example: RTT definition, measurement result format, time window, and aggregation approach).

Table 5: Reference Mapping

Ref.	Data or Context Supported	Source
N1	Physics-bound propagation ceiling for intra-European latency (speed of light in fibre; order-of-magnitude "ms per 1,000 km" rule-of-thumb).	Scaleway. *Understanding Network Latency (Propagation Delay)*
N2	Empirical RTT measurement basis (how RTT is measured/represented, and how to reproduce/inspect ping results used for Table 4 aggregates).	RIPE NCC. *RIPE Atlas Documentation – Measurements API (Ping)* RIPE NCC. *RIPE Atlas Documentation – Measurement Result Format*
N3	Published interconnection latency and KPI benchmarks (IXP service-level RTT/jitter between locations; fabric delay, delay variation, and frame-loss KPIs). Used as independent corroboration alongside RIPE Atlas measurements.	DE-CIX. *DE-CIX Service Levels for DE-CIX Locations* (RTT / Jitter tables). AMS-IX. *Realtime Statistics – SLA KPIs (Delay / Delay Variation / Frame Loss)*
N4	CDN mechanism: caching and serving from nearby nodes reduces user-perceived latency (edge delivery principle).	Akamai Technologies. *What Is a CDN (Content Delivery Network)?*

N5	CloudFront mechanisms: edge caching can reduce latency; persistent connections and connection reuse reduce repeated TCP/TLS handshakes.	Amazon Web Services. *Amazon CloudFront Developer Guide*
N6	Protocol-level latency reduction at the edge: HTTP/3 over QUIC; 0-RTT resumption to reduce connection-establishment latency for returning clients.	Cloudflare, Inc. *HTTP/3 (QUIC) – Developer Documentation.* Cloudflare, Inc. *0-RTT Connection Resumption – Developer Documentation*
N7	Netflix Open Connect architecture: OCAs and localisation model bringing content close to users/ISPs (reduced long-haul delivery and improved performance).	Netflix, Inc. *Open Connect Briefing Paper – A Cooperative Approach to Content Delivery*. Netflix, Inc. *Netflix Open Connect: Delivering Content at Scale*
N8	Anchor discovery and selection basis: confirming which anchors exist for a given metro; anchor identifiers and FQDNs used as "from" and "to" endpoints.	RIPE NCC. *RIPE Atlas – Anchors (Directory / UI listing)*
N9	Anchor metadata for "actual from/to city" fields in your summary exports (city/country from RIPE Atlas anchor records), and reproducible lookup via API.	RIPE NCC. *RIPE Atlas Documentation – Anchors API (anchor lookup / search by FQDN)*
N10	Anchor availability / coverage justification (why some routes are 2×3 rather than 3×3, and why certain metros have limited anchors).	RIPE NCC. *RIPE Atlas Statistics – Coverage.* Amsterdam: RIPE NCC
N11	Table 4 aggregation method (median of pairwise medians; Typical Range = P25–P75 of pairwise medians; per-measurement minimum RTT used before aggregation)	RIPE NCC. *RIPE Atlas Documentation – Measurement Result Format* (fields and semantics used for aggregation)

11.5 Network Topology and Transport Architecture

Europe's physical network substrate is among the densest and most advanced in the world. Long-haul and metropolitan fibre routes interconnect every major city and data-centre cluster across the continent, owned or operated by carriers such as Orange, Deutsche Telekom, Telia Carrier,

Colt, GTT, Lumen, and numerous national and regional providers. Hyperscalers and European operators alike rely on this same optical infrastructure: the same ducts, amplifiers, and dense-wavelength equipment. The glass in the ground is already fast enough.

11.5.1 Optical Layer

At the physical layer, each fibre pair carries multiple wavelengths using dense wavelength-division multiplexing (DWDM). Each wavelength – or lambda – functions as an independent optical channel with capacity between 100 and 800 Gb/s, depending on modulation and distance. Hyperscalers typically secure dedicated wavelengths or indefeasible rights of use (IRUs) on carrier fibre, giving them deterministic bandwidth without owning the underlying cable. European operators can and often do the same; the technology and commercial model are identical.

11.5.2 Transport and Routing Layer

Above the optical layer, hyperscalers deploy private backbones built on standard technologies such as MPLS, Segment Routing, and Software-Defined Networking (SDN). These frameworks allow deterministic routing, traffic engineering, and real-time telemetry collection across global backbones. The distinguishing factor is not the hardware but the tight coupling between network control and compute orchestration. In hyperscale environments, routing decisions are aware of workload placement and data-replication policies; the same telemetry informs both transport optimisation and service scheduling.

European providers possess all the technical components to replicate this model. What remains missing is a shared orchestration framework that can extend routing and telemetry semantics across ownership boundaries – so that multiple sovereign networks can operate as one logical operational fabric. This is a central design objective of the InnoFabric control layer, which aims to unify network telemetry, workload placement, and policy enforcement across heterogeneous infrastructure.

11.5.3 LAN versus WAN Domains

Performance sensitivity differs sharply between the local and wide-area domains:

1. Local-area (LAN) domain – Inside each data centre or regional cluster, the compute fabric must operate at near-zero latency. CPUs, GPUs, and DPUs communicate over lossless, deterministic networks (RoCEv2, InfiniBand, CXL) where microsecond delays directly translate into wasted silicon cycles. Orchestration, storage, and telemetry must function as a single, tightly coupled system.
2. Wide-area (WAN) domain – Across cities and borders, latency budgets of 15–25 milliseconds are acceptable for asynchronous replication, API transport, and content caching. The critical requirement is coherence: applications should connect to their local data store, while replication occurs transparently across regions. This is the same design pattern used by hyperscalers such as AWS Dynamo, Google Spanner, and Azure Cosmos DB – local consistency with cross-region durability.

11.5.4 Integration with EUCLORA and InnoFabric

InnoFabric's telemetry and orchestration interfaces are designed to accommodate both domains. Within each facility, it exposes real-time metrics for link utilisation, queue depth, and energy

profile to the control layer; across facilities, it models aggregate latency and throughput as dynamic resources in the same XRN (eXtended Resource Name) space. This enables policy engines to place workloads intelligently: close to users, near data, and within sovereign jurisdictions – while using inter-region networks only for replication or API transport.

11.5.5 Summary

Europe already owns the fibre and optical capacity required for a sovereign, federated cloud. The bottleneck lies not in bandwidth but in the lack of shared operational cloud fabric. By aligning network telemetry, optical routing, and workload orchestration through open standards, EUCLORA's architecture can transform Europe's fragmented connectivity into a coherent, measurable backbone for digital sovereignty.

11.6 Quantifying the Efficiency Gap

This Appendix quantifies the efficiency gap between leading hyperscalers and EU-owned cloud operators across three dimensions: Power Usage Effectiveness (PUE), automation density (servers per operations engineer, SRE/ops), and effective compute utilisation (the share of server capacity doing useful work rather than sitting idle). The objective is to provide a transparent, data-driven baseline for Europe's cloud-efficiency strategy using publicly available ESG disclosures, industry surveys, and peer-reviewed operational literature. All inputs and derived figures are traceable through the local Reference Mapping table (Table 9), which links each assumption to its underlying publication(s) in the main References section (Section 9).

Power Usage Effectiveness (PUE) is the industry-standard metric for data-centre facility efficiency. It is defined as the ratio of total facility power consumption (including cooling, power distribution losses, lighting, and other overheads) to the power used directly by IT equipment (servers, storage, and networking):

$$PUE = \frac{\textbf{Total Facility Power}}{\textbf{IT Equipment Power}}$$

A PUE of 1.0 represents a theoretical optimum in which every watt drawn by the facility is delivered to IT equipment. In practice, modern data centres typically operate above 1.0, with achieved values depending on scale, design, cooling approach, and utilisation. Lower PUE values indicate higher facility efficiency. PUE is defined in industry standards and commonly reported in sustainability and ESG frameworks, enabling portfolio-level benchmarking across operators.

Table 6: Efficiency Gap

Operator	PUE	Automation density (servers per SRE/ops)	Effective compute utilisation, U (useful work vs idle; CPU proxy)	Indicative energy per unit useful compute vs AWS [A13]
AWS (benchmark)	1.15 (published [A1])	3,000–5,000 (author-defined modelling extension beyond the published ~2,500:1 observation; see [A2])	55% (conservative midpoint assumption for effective compute utilisation; see [A6] for differential framing and [A13] for the derived-metric definition).	1.00×
T-Systems (DE)	1.53 (published [A4])	150 (modelled; EU band anchored by [A5])	25% (midpoint model input; anchored by [A6])	2.93×
Hetzner (DE)	1.13 (published [A7])	200 (modelled; EU band anchored by [A5])	35% (midpoint model input; anchored by [A6])	1.54×
OVHcloud (FR)	1.26 (published [A8])	350–400 (modelled; EU band anchored by [A5])	40% (midpoint model input; anchored by [A6])	1.51×
Scaleway (FR)	1.37 (published [A9])	200 (modelled; EU band anchored by [A5])	35% (midpoint model input; anchored by [A6])	1.87×
EU-owned average	1.36 (energy-weighted baseline [A11])	150–250 (modelled; EU band anchored by [A5])	35% (midpoint model input; anchored by [A6])	1.86×

Notes

1. Effective compute utilisation (U) values are model inputs anchored by the utilisation differential described in [A6], because operator-specific fleet-average compute utilisation is generally not publicly disclosed.

2. Indicative energy per unit useful compute vs AWS is computed as (PUE / U) ÷ (AWS PUE / AWS U) using the midpoint utilisation assumptions listed above, as defined in [A13].

11.6.1 Energy Efficiency Gap

Across EU-owned data-centre operators, the energy-weighted average Power Usage Effectiveness (PUE) reported under the EU reporting framework is approximately 1.36, compared with 1.15 for AWS's fleet-average benchmark [A1], [A11]. This facility-efficiency differential implies that, for the same IT load, EU operator portfolios draw materially more total energy at the meter due to higher cooling and power-distribution overheads.

Beyond facility overheads, hyperscalers also benefit from two independent operational levers: higher automation density (modelled here as 3,000–5,000 servers per SRE/ops engineer versus 150–400 for non-hyperscale operators) [A2], [A5], and higher effective compute utilisation enabled by automated placement and workload mixing, which reduces idle capacity [A6]. Because operator-specific fleet-average compute utilisation is rarely disclosed, utilisation values in Table 6 are treated as model inputs, anchored by the utilisation differential described in [A6] and applied transparently through the derived metric defined in [A13].

To express the combined impact in a single, auditable indicator, Table 6 reports "indicative energy per unit useful compute vs AWS", computed as (PUE / U) ÷ (AWS PUE / AWS U) using the stated midpoint assumptions [A13]. Under these assumptions, the EU-owned average remains materially above the AWS benchmark, indicating substantial headroom for efficiency gains through improved facility performance, greater operational automation, and higher effective utilisation.

11.6.2 Methodology Overview

All quantitative estimates in this report are derived from a uniform baseline model comparing the operational efficiency of EU-owned cloud and data-centre infrastructure with that of leading hyperscalers. The model integrates three primary dimensions: facility energy efficiency (PUE), effective compute utilisation (useful work versus idle capacity), and automation density (servers per operations engineer, SRE/ops). Assumptions are drawn from aggregated industry data, public operator disclosures, independent research, and energy benchmarks, with each input traceable via the Reference Mapping in Table 9.

Sensitivity note – Quantitative results in this Appendix carry moderate uncertainty inherent in industry-aggregated and partially disclosed operational data. EUCLORA performed an internal one-at-a-time sensitivity check of ±10% on the three primary inputs (PUE, effective utilisation U, and automation density AD). Across these perturbations, the direction of the relative efficiency gap (EU-owned operators above the hyperscaler benchmark on the derived energy-per-useful-compute indicator) remains unchanged. (Author analysis; detailed sweep not shown.)

Silicon-level design advantages are not modelled as a separate term. Where relevant, their effects are treated as part of the observed operational outcomes captured by the model — in particular higher effective utilisation and reduced overhead through automation and workload placement. Custom accelerators and DPUs can offload network, security, and specialised compute tasks, potentially improving energy proportionality; future EUCLORA iterations may incorporate explicit silicon-efficiency terms where comparable measurement data enables quantification across hardware types.

Calculations are normalised to an estimated 2.0 million servers deployed across EU-owned data centres, with an assumed average electrical draw of 0.35 kW per server and 8,000 operating hours per year. Comparative benchmarks for hyperscalers use PUE = 1.15 [A1] and utilisation assumptions consistent with the utilisation differential described in [A6] and applied transparently through the derived metric defined in [A13]. Electricity costs are monetised using 2024 EU weighted-average non-household electricity prices in the € 0.16–0.19/kWh range (depending on tax treatment and contracting assumptions), per Eurostat's non-household price statistics (dataset nrg_pc_205) [A14].

Table 7: Model Estimates

Parameter	EU-owned operators	Hyperscale benchmarks	Unit / assumption	Source reference
Average PUE	1.36 (Energy-weighted baseline)	1.15	Ratio (fleet average)	Table 6; [A1], [A11]
Effective compute utilisation (U)	25–40% (midpoint used in model: 35%)	55% (midpoint)	Share of capacity running initiated workloads (fleet average)	Table 6; utilisation differential framing [A6]; midpoint assumptions applied transparently via derived indicator definition [A13].
Automation density	150–400 servers/FTE (midpoint: 250)	3,000–5,000 servers/FTE	Servers per ops/SRE FTE	Table 6; [A2], [A5]
Normalisation: server count	≈ 2.0 million	–	Servers	Model assumption
Normalisation: average IT draw per server	0.35	–	kW (average IT load)	Model assumption (blended fleet-average IT draw; used only for normalised scaling)

Normalisation: operating hours	8,000	–	h/year	Model assumption
Electricity cost	€ 0.16–0.19/kWh	–	EU-average industrial electricity cost (used for monetisation)	[A14]
Annual IT energy (normalised)	≈ 5.6 TWh	–	Model output: N × kW × h	Derived
Annual facility energy (normalised)	≈ 7.6–7.7 TWh	≈ 6.4 TWh	Model output: IT energy × PUE	Derived; PUE inputs from Table 6
Relative energy per unit useful compute vs hyperscale	≈ 1.8×–3.0× (midpoint ≈ 2.1×)	1.0×	Defined as (PUE/U) ÷ (AWS PUE/AWS U)	Derived metric definition [A13] using PUE inputs [A1], [A11] and utilisation-differential framing [A6]
EU-wide electricity baseline (facility)	45–65 TWh/year (midpoint ≈ 50 TWh/year)	–	Used for scaling the opportunity	[A12]
Indicative avoidable electricity (EU-wide)	≈ 22–33 TWh/year (midpoint ≈ 26 TWh/year)	–	50 TWh × (1 – 1/relative factor)	Scaling baseline [A12] combined with the derived relative-factor definition [A13].
Indicative avoidable electricity cost (EU-wide)	≈ € 3.5–6.3 bn/year (midpoint ≈ € 4.4–5.2 bn/year; central estimate ≈ € 4.8 bn/year)	–	Avoidable TWh × € /kWh	[A12], [A14]

11.6.3 Efficiency Metric Definitions and Formulas

To ensure transparent and repeatable efficiency measurement, EUCLORA defines three primary quantitative indicators: Energy Efficiency (EE_IT), Effective Compute Utilisation (U), and Automation Density (AD).

Each can be derived directly from observable telemetry or audited ESG datasets.

Table 8: Efficiency Metric Definitions and Formulas

Metric	Definition	Formula	Units	Description
Energy Efficiency (EE_IT)	Compute output per unit of IT energy consumed	EE_IT = W_out / E_IT	Workload-hours / kWh	Measures how effectively IT energy is converted into useful compute output (excludes facility overhead).
Effective Compute Utilisation (U)	Initiated workload activity as a share of total available compute capacity	U = (C_active / C_total) × 100	%	Proxy for "useful work vs idle" across servers (CPU proxy; extendable to GPU/accelerators).
Automation Density (AD)	Number of servers managed per SRE/ops full-time equivalent	AD = N_servers / N_SRE-ops_FTE	servers / FTE	Captures operational automation and software leverage (higher AD implies fewer staff per managed server).
Facility Energy Efficiency (PUE)	Ratio of total facility energy to IT equipment energy	PUE = E_facility / E_IT	–	Standard industry metric for facility-level overhead (ISO/IEC 30134-2) [A15].
Composite Efficiency Index (CEI)	Normalised composite index combining IT efficiency, utilisation, and operational leverage, adjusted for facility overhead	CEI = (EE_IT × U × AD) / PUE	Relative index	Integrates the three dimensions into a single comparable indicator; normalise CEI to 1.0 for the hyperscaler benchmark case.

Notes

- W_out – total compute output in workload-hours (normalised across CPU, GPU, and accelerator time).

- E_IT – IT equipment energy consumption in kWh for the same period (servers, storage, networking).
- E_facility – total facility energy consumption in kWh for the same period (includes cooling, power distribution losses, lighting, etc.).
- C_active / C_total – measured as active versus available compute capacity (cycles, cores, or normalised resource units).
- N_SREops_FTE – operations personnel (SRE/ops) responsible for fleet operation (exclude unrelated corporate functions).
- Normalisation – CEI is normalised to a baseline of 1.0 for hyperscaler benchmarks.

These metrics form the quantitative backbone of EUCLORA's efficiency benchmarking framework and can be derived directly from the InnoFabric telemetry schema.

They provide a consistent, auditable way to measure software-driven infrastructure efficiency across heterogeneous providers.

11.6.4 Interpretation:

Using the baseline assumptions in this Appendix, the model indicates that a substantial share of EU data-centre electricity demand is attributable to (i) higher facility overheads (PUE) relative to leading hyperscale benchmarks and (ii) lower effective compute utilisation (idle capacity). If EU-owned operators were able to converge towards hyperscaler-level PUE and utilisation outcomes, the implied reduction in electricity required per unit of useful compute is material, yielding order-of-magnitude savings in the tens of TWh per year when scaled to EU-wide data-centre electricity consumption [A1], [A6], [A11]–[A13]. Monetised at EU non-household electricity prices, this corresponds to several billion euros per year in direct electricity costs, with the exact value depending on the price basis (tax treatment, contracting, and consumer band) [A14].

11.6.5 Economic Efficiency Gap

The quantitative assumptions used in this analysis draw on the data sources [A11], [A12], [A13], [A14] (and the utilisation differential framing in [A6]), which together provide the empirical basis for monetising the efficiency gap.

[A11] establishes an EU baseline for facility overhead via the energy-weighted average PUE from the first EU reporting period.

[A12] provides an aggregate baseline for EU data-centre electricity consumption used for scaling.

[A13] defines the model's derived efficiency indicator (energy per unit useful compute relative to AWS), including the utilisation midpoints applied transparently.

[A14] provides the electricity-price basis (€/kWh) used to monetise electricity impacts.

On this basis, the economic gap quantified in this Appendix is primarily the direct electricity-cost component implied by higher facility overheads and lower effective utilisation. The resulting euro value is therefore best interpreted as a range (driven by the electricity-price basis and utilisation assumptions), rather than a single precise figure.

11.6.6 Appendix Reference Mapping (Efficiency)

This Appendix distinguishes between two levels of referencing to ensure both readability and traceability: (1) The main References section contains the complete bibliographic sources with titles and URLs. (2) This Appendix provides a Reference Mapping linking each quantitative statement or data point to its specific source(s). Local identifiers [A#] are used to cross-reference individual entries within the Appendix.

Table 9: Reference Mapping

Ref.	Data or Statement Supported	Source(s)
A1	Hyperscaler facility-efficiency benchmark: AWS reports fleet-average PUE = 1.15.	Amazon Web Services (AWS). Sustainability Report 2024 – AWS Summary.
A2	Hyperscaler operations leverage (modelling band): Hyperscale operations literature reports system-to-operator ratios in the thousands; Hamilton reports ratios up to ~2,500 systems per administrator. This paper uses 3,000–5,000 servers per SRE/ops as an author-defined modelling extension beyond the published 2,500:1 observation, and treats the increase as an explicit assumption reflecting greater automation and platform tooling in hyperscale operations since 2007.	Hamilton, J. On Designing and Deploying Internet-Scale Services. Verma, A. et al. Large-scale cluster management at Google with Borg.
A3	Establishes that hyperscale operators typically sustain materially higher average utilisation than conventional enterprise estates due to automated scheduling, workload mixing, and large-scale operational tooling. Because auditable fleet-average utilisation is not consistently disclosed, this Appendix treats utilisation as a transparent model input and applies conservative midpoint assumptions in Table 6; the derived indicator is defined in [A13].	National Renewable Energy Laboratory (NREL). *Data Center IT Efficiency Measures Evaluation Protocol*. Kaffes, K. et al. *Leveraging Application Classes to Save Power in Highly-Utilized Data Centers*. Author analysis (derived metric definition in Section 11.6) [A13].

A4	EU operator PUE example (Germany): provides published PUE = 1.53 for T-Systems data centres in Germany.	Deutsche Telekom. Corporate Responsibility Report 2024.
A5	EU / non-hyperscale operations leverage benchmark: typical organisations operate at ~50–100 servers per administrator, while "world-class" may reach ~300–400; used to parameterise an EU-operator modelling band of 150–400 servers per SRE/ops.	Knorr, E. "Microsoft exec: We 'get' the cloud" (interview with Bob Muglia). IDC. Data Center Operations Efficiency Study 2023
A6	Compute efficacy (effective utilisation) differential underpinning the narrative: traditional estates exhibit low average server utilisation (large idle capacity), while hyperscale fleets can sustain high average machine utilisation enabled by automated scheduling and workload mixing.	National Renewable Energy Laboratory (NREL). Data Center IT Efficiency Measures Evaluation Protocol Kaffes, K. et al. Leveraging Application Classes to Save Power in Highly-Utilized Data Centers
A7	EU operator PUE example (Hetzner): provides published PUE for Hetzner data centres (average ≈ 1.13; reported range ≈ 1.10–1.16).	Hetzner Online GmbH. Environmental and Energy Statement 2023 (PUE disclosure).
A8	Provides published PUE (1.26) and fleet scale (≈ 450,000 servers in operation) for OVHcloud, used as an EU-operator reference point in Table 6.	OVHcloud. *ESG Report 2023* (PUE). OVHcloud. *OVHcloud presents its strategic plan, Shaping the Future, and its new financial targets for FY2026* (server-count disclosure).
A9	EU operator PUE example (Scaleway): provides published average PUE for Scaleway data centres (average PUE = 1.37, 2023 indicator set).	Scaleway. Impact Report 2024 (data-centre PUE indicator).
A10	Energy-equivalence inputs (for "country / households" contextualisations): provides a basis for national electricity totals and household electricity consumption assumptions used in conversions.	International Energy Agency (IEA). Electricity Information 2023. Eurostat. Energy Statistics datasets (incl. household electricity consumption).

A11	Establishes the EU baseline for facility energy overheads (PUE) using the first EU data-centre reporting cycle: the report defines the KPI treatment and presents an EU-average PUE computed using an energy-consumption-weighted aggregation (EU average ≈ 1.36), including breakdowns by Member State and data-centre size category. This EU-average PUE is used as the baseline input for "EU-owned / non-hyperscale" scenarios in Section 11.6.	European Commission, DG ENER. *Assessment of the Energy Performance and Sustainability of Data Centres in EU: First Technical Report*
A12	Aggregate EU data-centre electricity consumption baseline used for scale framing (order-of-magnitude).	European Commission / JRC. Energy Consumption in Data Centres and Broadband Communication Networks in the European Union (EU-27) in 2022 (and related Commission reporting).
A13	Derived metric definition (calculation method): "Indicative energy per unit useful compute vs AWS". Computed as (PUE / U) ÷ (AWS PUE / AWS U), using midpoint utilisation assumptions (AWS $U = 55\%$; operator midpoints within model bands: T-Systems 25%, Hetzner 35%, OVHcloud 40%, Scaleway 35%, EU-average 35%). Utilisation values are model inputs anchored by the utilisation-differential framing in [A6], because operator-specific fleet-average compute utilisation is generally not publicly disclosed. The resulting ratios shown in Table 6 are calculated directly from the Table 6 midpoint inputs and rounded to two decimals (T-Systems 2.93×, Hetzner 1.54×, OVHcloud 1.51×, Scaleway 1.87×, EU-average 1.86×).	Author calculation defined in Section 11.6, using parameters from A1(AWS PUE), [A4], [A7]–[A9], [A11] (operator/EU PUE inputs), and A6(utilisation framing).
A14	Provides the EU-average electricity price for non-household consumers used to monetise energy impacts (€/kWh). For medium-sized non-	Electricity price statistics and underlying dataset Electricity prices for non-household consumers (nrg_pc_205).

	household consumers (annual consumption 500–2,000 MWh), Eurostat reports EU weighted-average prices in 2024 of € 0.1885/kWh (H1 2024) and € 0.1941/kWh (H2 2024) including non-recoverable taxes; and € 0.1575/kWh (H1 2024) and € 0.1629/kWh (H2 2024) excluding taxes (energy + supply + network). These figures support an industrial electricity-price modelling band of approximately € 0.16–0.19/kWh depending on tax treatment and contracting assumptions.	
A15	Defines Power Usage Effectiveness (PUE) and its calculation boundary for data centres, providing the normative definition used in Table 8 and Appendix 11.6.	ISO/IEC. ISO/IEC 30134-2: Information technology – Data centres – Key performance indicators – Part 2: Power Usage Effectiveness (PUE).

12 Imprint

EUCLORA

European Cloud Computing Research Alliance AISBL

(Association internationale sans but lucratif)

Registered office

Boulevard Brand Whitlock 132

1200 Brussels

Belgium

Contact

hello@euclora.org

This publication is issued by EUCLORA as part of the EUCLORA Research Series on Cloud Efficiency and Sovereignty.

The views expressed in this document reflect the analysis and policy position of EUCLORA at the time of publication and do not necessarily represent the views of individual members, funding bodies, or associated institutions.

This document may be cited for policy, research, and institutional purposes, provided that proper attribution is given.