Amsterdam Economic Board

# LEAP TECHNOLOGY TRENDS & SCENARIOS

## *Building future-proof and sustainable digital infrastructures*

**JUNE 2021**

**The Lower Energy Acceleration Program (LEAP) aims to accelerate the transition to a sustainable digital infrastructure. This document provides an overview of existing and new solutions and scenarios to speed up the transition to energy efficient and effective digital infrastructures. The solutions range from the environmental to the technical to the social. Four scenarios, which have been carved out by combining solutions, demonstrate a shift towards a sustainable future. The aim is to inspire and to set direction.**

Digital infrastructures have to deal with massive amounts of data. High bandwidth data transfers, affordable data plans, cloud migrations and the increasing popularity of streaming services mean data consumption is growing at unprecedented rates. While over the past few decades the energy efficiency of computing hardware has drastically improved and software performance and usability have become far more effective, they still cannot keep up with demand. The growing energy needs of ICT is especially important in the Netherlands—a prominent European "data hub" distributed over a relatively small geographic area.

The transition to a sustainable digital infrastructure need to speed up. Using renewable energy resources is only part of the solution, since its production still has limitations. The innovative solutions available need to leveraged, new ones stimulated and barriers hindering adoption removed. The digital system will need to be more integrated within our energy infrastructure and environment.

VU Vrije Universiteit Amsterdam

PhotonDelta

### Table of contents

## Three horizons

This article provides an overview of technological, social and environmental solutions on three horizons: solutions that are available today (H1), solutions for the near future (H2) and solutions further away (H3). Then key barriers and different scenarios that could lead to a future generation of digital infrastructures are explored. Finally, the future: what is needed to accelerate the adoption of existing and new technological solutions in our digital infrastructure?

In this article, a number of key solutions and challenges are presented. For a full overview, please see the report: The LEAP Technology Landscape for 2030 and Beyond – Solutions, Adoption Factors, Impediments, Open Problems, and Scenarios by Roberto Verdecchia (VU Amsterdam), Patricia Lago (VU Amsterdam) and Carol de Vries (PhotonDelta). Their research is based on interviews and focus groups with more than 40 stakeholders from the data center industry.

# HORIZON 1: SOLUTIONS FOR TODAY

## Solutions across time horizons

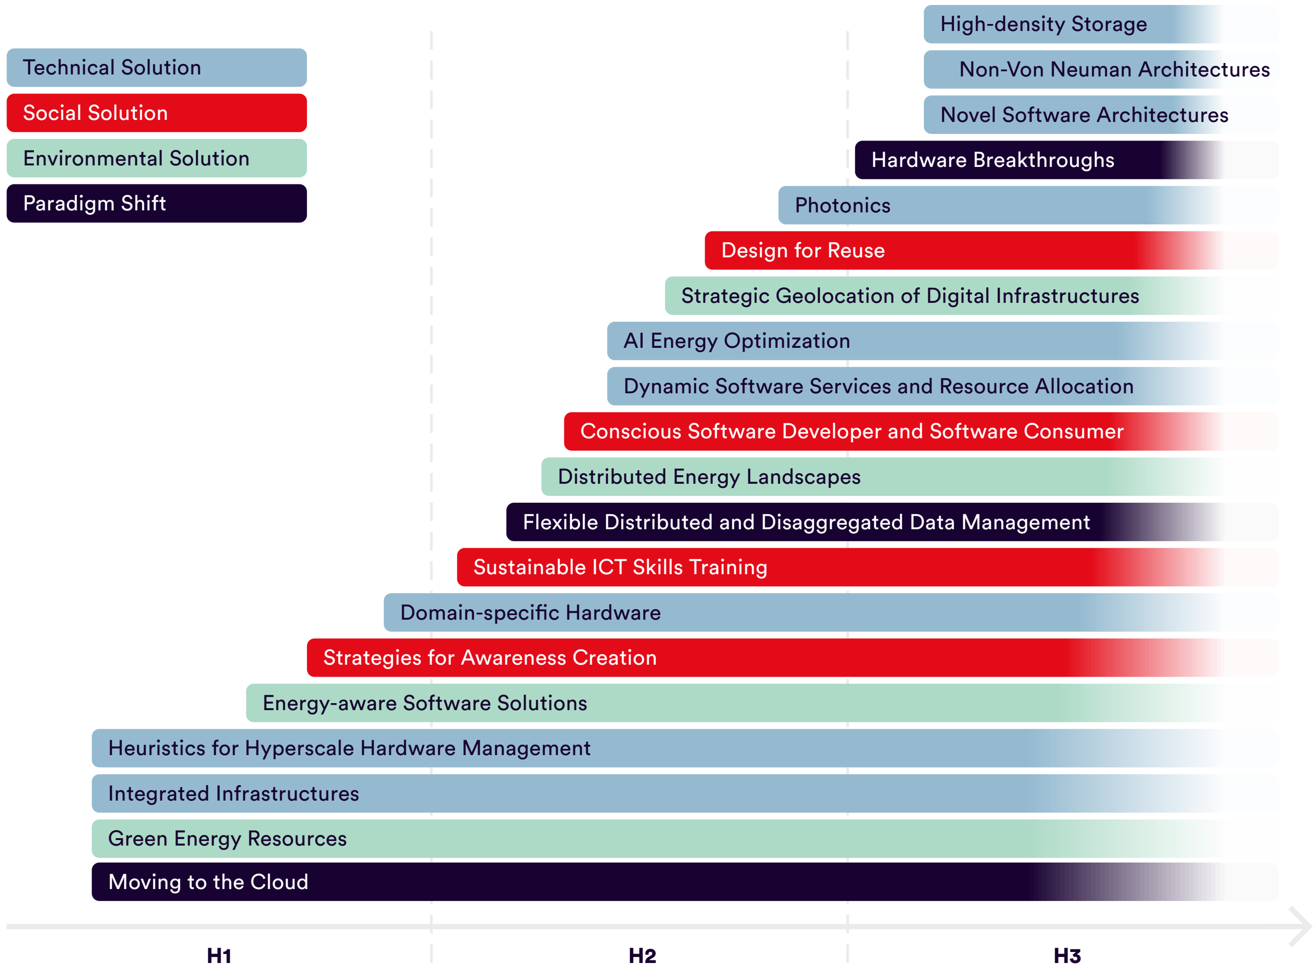




**Horizon 1 solutions are already available for adoption. Although these proven technologies are well known, there is still significant room for implementation of their energy-saving potential. These solutions are expected to reach further maturity or full potential in H2.**

### Moving to the cloud

Cloud services are the on-demand availability of computer system resources, especially data storage and computing power, without direct active management by the user. During H1, a key trend is the movement to cloud services, contributing to a more energy-efficient digital infrastructure. It entails moving data, computational and software capabilities from on-premise to the cloud. It increases the popularity of energy-efficient software applications specifically designed to be deployed in the cloud, such as cloud-native and serverless applications.

### Heuristics for hyperscale hardware management

Migrating to the cloud entails a growing centralization of both software and hardware resources in hyperscale digital infrastructures. Heuristics for hyperscale hardware management are solutions for the energy optimization of this type of digital infrastructure. Examples are heat management—such as efficient cooling strategies and the reuse of dissipated thermal heat—and energy-aware storage optimization.

### Energy-aware software optimizations

There are a number of energy-aware software optimizations in Horizon 1 that spill over to Horizon 2. On one hand, these solutions are software-centric, such as event-based software that only consume resources when certain triggers appear in an event stream. On the other hand, these solutions are also cloud-centric, for example kill zombie systems, that detect and shut down idle servers to ensure no energy waste.

In H2, these optimizations can be used for flexible distribution and disaggregation and smart virtualization. Current trends predict the widespread popularization of innovative software optimizations, such as fine-grained dynamic load balancing and AI-enabled optimization of software energy consumption.

### Strategies for awareness creation

On the social level, it is imperative to communicate the environmental impact of data production, data manipulation and data usage. Building awareness about the use of energy in data consumption could lead to behavioral changes in data consumption patterns.

# HORIZON 2: SOLUTIONS FOR THE NEAR FUTURE

**Technological innovations and solutions expected to be adopted in Horizon 2 (H2) - within the next**

**four to six years - are currently being developed. Attention needs to be paid to these solutions over the next few years to enable them to grow.**

### Software enabling distributed and edge computing

In the coming years, paradigm shift from hyperscale data centers to distributed and edge computing will take place. With steady advancements in communication technologies, 5G, the Internet of Things and the growing affordability of computational power, edge computing is expected to gain widespread popularity in the near future. Edge computing will vary from static on-premise mini-clouds to flexible 'follow-the-need services'. A shift of tasks from the cloud back to smart devices including smartphones is also taking place.

The shift towards a distributed paradigm enables the use of distributed energy landscapes. These are supported by smart energy grids that locally produce and consume energy. This also allows for the strategic geolocation of digital infrastructures. Digital infrastructures can be strategically positioned close to their end users in order to ensure high bandwidth and low energy consumption of data flow and communication.

In addition, distribution and disaggregation supports more precise profiling of energy consumption patterns, allowing for dynamic software services and resource allocation. This solution enables the profiling of data usage patterns and the dynamic allocation of services and resources. This shift will require the creation of novel software architectures realizing the "software intelligence" necessary to enable adaptation and flexibility.

### AI energy optimization

Artificial intelligence (AI) will also play a prominent role in Horizon 2, with a focus on distribution and disaggregation. Examples of AI for energy efficient digital infrastructures include:

- novel federated learning algorithms, supported by the appearance of edge AI.
- data optimization/compression strategies allowing the transfer of high volumes of curated information rather than raw data.
- approximate computing, which is the provisioning of results of acceptable quality, rather than optimal, to reduce energy consumption.

Energy efficiency software engineering for AI-based systems is also rapidly gaining traction. Promising prototypes showcase how utilizing AI-dedicated hardware components can reduce energy consumption of AI-based systems.



### Design for reuse

The lifecycle expectancy of digital infrastructure hardware components in H1 is approximately two to three years. In H2, a growing design trend for reuse and hardware lifecycle management practices is expected, which will reduce the carbon footprint of hardware production and the need for critical materials. Design for reuse, however, implies a delay in using newer and more efficient technologies that drive hardware waste, meaning trade-offs must be carefully considered.

### Sustainable ICT skills training and conscious software developers and consumers

H2 is characterized by the popularization of sustainable ICT skills training—in both academic and industrial settings. This is necessary to address the energy-effectiveness and sustainability of digital infrastructures across all types of industrial sectors. An example is using the eco-settings of data servers as a tool to reduce energy consumption by about 10%.

The rise of both responsible software developers and software consumers leads to more sustainable software solutions by design, as consumers demand more sustainable digital infrastructures from developers.

## HORIZON 3: SOLUTIONS FURTHER AWAY

**The impact of technological innovations in Horizon 3, which are expected to drastically change**

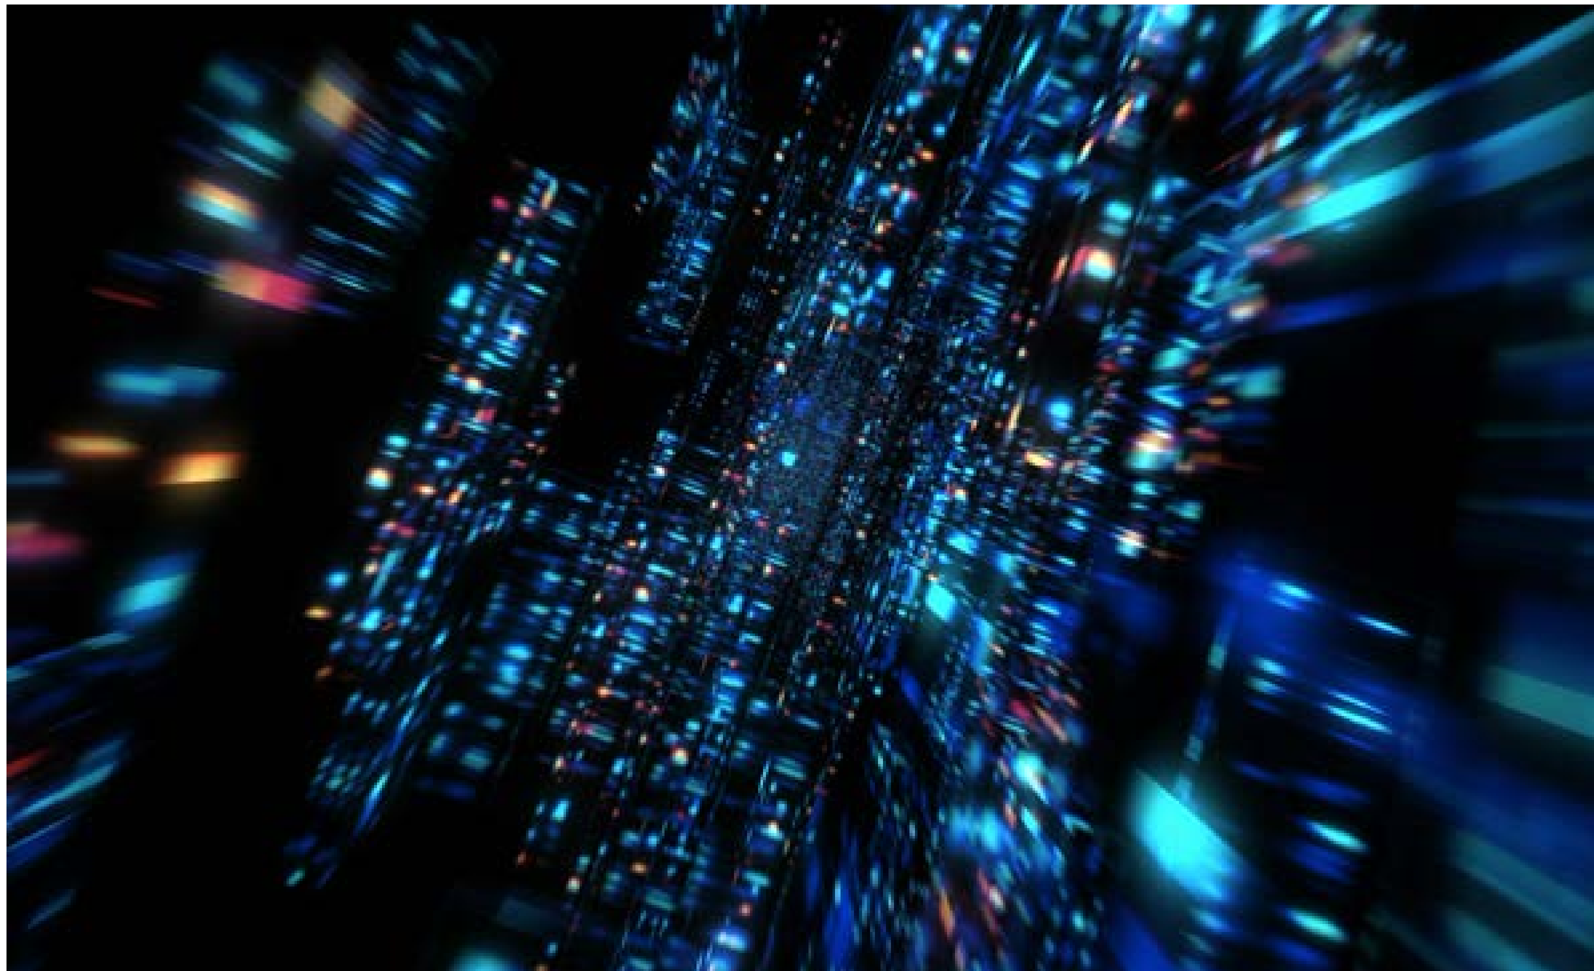

**how digital infrastructures are designed and operated, are currently in the research phase and expected to be operational in about seven years.**

### Photonics

The most important hardware breakthrough in H3 will be the widespread use of photonics. While photonics is already used in H1, it is expected to become much faster and ubiquitous during H2 and will be of widespread and consolidated use in about 10 years. In addition to inter-server/client communication (co-packaged optics), which is already showing promising results, the widespread transition towards optic communication is also expected to occur within data centers, with integrated photonics leading to the replacement of micro-electronic hardware with optics-based ones.

### Non-von Neumann Architectures

Any computer architecture in which the underlying model of computation is different from what has come to be called the standard von Neumann model are called Non-von Neumann architectures. Non-von Neumann architectures can execute complex computational tasks using a fraction of the energy consumed by today's hardware.

### Neuromorphic

Neuromorphic computing is the use of very-large-scale integration of systems containing electronic circuits to mimic neuro-biological architectures. The evolution of neuromorphic computing and similar solutions can lead to groundbreaking energy savings when it comes to computational-intensive tasks.

### Novel Software Architectures

With hardware expected to present considerable breakthroughs in H3, software will also need to evolve and adapt. Novel software architectures must be created to fit the drastic technological changes resulting from H3's hardware advances.

## CHALLENGES

**These technological developments are promising, but to take advantage of the opportunities they present, several challenges need to be met and conditions created for continued innovation.**

Problems and impediments fall into the following four categories:

### Lack of knowledge and awareness

The sustainability of digital infrastructures must become a higher priority. Often, due to a lack of knowledge in the business context, it is neglected in favour of other goals such as reaching business targets and customer satisfaction. Education and training programs dedicated to the impact ICT sustainability can have on a business will help alleviate this lack of knowledge and awareness.

Without a clear understanding of the (business or technological) impact of these new technologies, business leaders may be adverse to change. Reshaping consolidated technologies may lead to uncharted situations, which means most organizations will require clear analyzes, understanding and proof before investing in these technologies on a substantial scale.

Conducting systematic experiments to evaluate and measure the impact of the presented solutions will, from a business point of view, help mitigate the vague understanding of their impact on quality, energy savings and the evolution of technology ecosystems. In addition, a clear understanding of the potential tradeoffs between implementation and deployment costs is critical. There should be an understanding of—and support for—the tradeoff, as energy savings will not deteriorate the quality of services provided. Research into future innovations needs business cases and use cases before they can be successfully adopted in industrial contexts.

### Lack of policies, KPIs and activating strategies

Currently, there are barely any regulations, standards or policies regarding the sustainability of digital infrastructures. The complexity of this problem can be traced to the current absence of KPIs, such as sustainability labels for software and hardware components, that can track and guide the progress of stakeholders developing sustainable digital infrastructures.

The lack of activating strategies is another barrier to innovation. Activating strategies can include higher taxation of fossil energy sources, dynamic pricing based on real-time energy demands and more (tax) legislation on e-waste production. Activating taxation strategies could trigger stakeholders and target other areas beyond electricity and carbon emissions, such as the circular use of critical materials.

**Need for coordinated change and leading champions**

Moving to a sustainable digital infrastructure is a complex transition involving many stakeholders. Historical data shows that technological leaps often take between 10-20 years to gain full maturity; that's how much time it takes for business leaders to clearly understand the pros and cons of the technology and for it to be adopted by a wider audience.

To accelerate this process, a coordinated change enables stakeholders to progress jointly, sharing the investments and risks. This avoids a scattered landscape, characterized by compartmentalized technology silos adopted by only a few companies.

Many companies perceive a lack of guidance in their sustainability journeys and miss systematic support in the adoption of energy efficient solutions as well as leading champions. Guidance could come from governmental institutions, research consortiums or even private companies that champion and supervise the common endeavour towards more sustainable digital infrastructures. Leading champions should take initiative within their organizations and steer change towards the next generation of sustainable technologies and share their journey with other organizations to learn from.

**Technical barriers**

One of the biggest challenges underpinning the proposed solutions is the availability of electricity. For example, it is estimated that the electric grid in the Netherlands will need to at least double in capacity during H2—and even more as other industries transition to electric energy. Capacity planning of the energy infrastructure, and the impending scarce supply of renewable energy, remain big challenges.

There are as yet not many off-the-shelf solutions for optimizing energy consumption in data centers. This leads to creating or adapting ad-hoc solutions in lieu of efficiently and effectively applying permanent solutions.

As the reduction of the power consumption of micro-electronics becomes physically impossible, there is little freedom left to further optimize the energy consumption of currently available technology. Current data processing paradigms need to be transformed to enable a transition towards the next generation of hardware technologies. Many data centers still use so-called 'brown software' that is very energy inefficient. This problem is increasing with the ever-growing adoption of AI, which at the moment is characterized by severe software energy inefficiencies.

Lastly, implementing and integrating solutions without drastic change to the normal functioning of data centers prevents forward motion. The normal functioning of digital infrastructures cannot be interrupted while integrating a new solution, and the return of investment of adopting a certain solution cannot be hindered by the cost of integrating the solution.

## FOUR SCENARIOS FOR AN ENERGY-EFFICIENT INFRASTRUCTURE

**Four scenarios have been developed to move to a future-proof and sustainable digital infrastructure. They are based on combining solutions from Horizons 1, 2 and 3, and they are ordered chronologically from today to what is expected during Horizon 3.**

### Overview of Sustainable Digital Infrastructure Scenarios

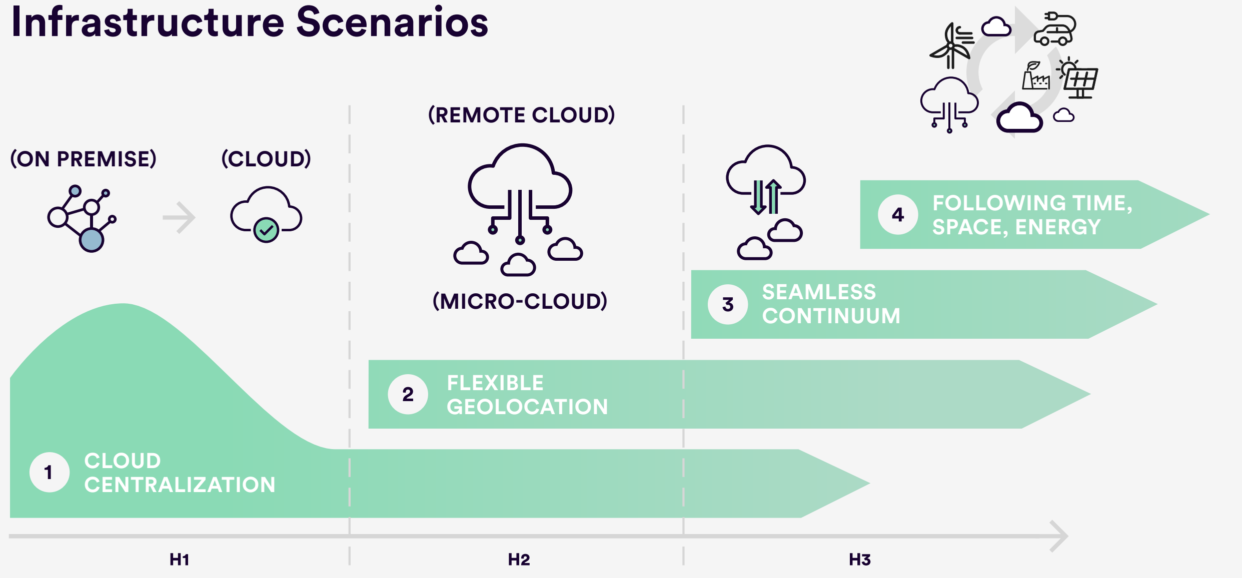


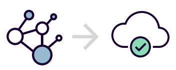

**Scenario 1: Cloud centralization**

Today, organizations migrate their software and hardware from on-premise to a centralized remote cloud. Optimization and on-demand use enable an energy-efficient digital infrastructure. The centralization requires vast volumes of energy and physical space to facilitate hyperscale premises already in competition with other societal needs on the energy grid. With an expected exponential data growth in the next couple of years, cloud centralization will very likely be only one solution going forward.

**Scenario 2: Flexible geolocation**

Moving to a distributed system, which takes place in Horizon 2, both remote clouds and micro-clouds will coexist in a hybrid model. It is based on the paradigm shift to flexible distributed and disaggregated data management, where edge computing, distributed energy landscapes, dynamic software services and resource allocation gain traction. These technologies enable the distribution of both the computational and energy consumption load over different geographical areas, mitigating energy consumption centralized in specific geographical areas as outlined in Scenario 1.

### Scenario 3: Seamless continuum

This scenario is characterized by a pool of shared hardware and software resources that are available in both micro-clouds and remote clouds. This shared pool of resources facilitates a so-called seamless continuum, where hardware and software resources are allocated at runtime. If a certain resource is needed, the ones that are the best fit—both in terms of performance and energy—are chosen.

### Scenario 4: Follow time, space and energy

Building on Scenario 3, in this scenario resources are allocated based on their software and hardware capabilities, the availability of energy needed, the proximity of resources and the timeliness of the task at hand. Any digital infrastructure resource can be used seamlessly, linking anything and everything.

## WHAT IS NEEDED TO ACCELERATE?

**The aim of this document is to set direction and to inspire. There are many different ways to become more energy-efficient and -effective, but investment and implementation depend on business conditions being met. One of the key conditions to accelerate is to stimulate innovations, best practices, pilots and demonstrations and share learnings amongst all stakeholders. It's critical for front running champions to take the lead with a joint mission.**

The centralisation of data centers is increasingly difficult to integrated in the energy system and spatial planning. Until novel hardware breakthroughs occur in horizon 3, a distributed architecture allows for sustainable progress and better integration. But once here, the hardware will set new standards of low-energy data storage, communication and data processing with the least impact on the environment. The most likely scenario is that these new technologies will coexist with older technologies for a while, with more and more focus on sustainable business.

Besides technical challenges being met, an increased awareness of the sustainable digital services is crucial for progress. If stakeholders throughout the value chain are made aware of their roles, impact and opportunities, they will also be more willing and able to put sustainability higher on their agendas and being held accountable for their actions.

Lastly, public-private partnership are needed and a proactive government plays a pivotal role in stimulating innovations towards sustainable digital infrastructures.



## LEAP AND THE TECHNOLOGY TRENDS AND SCENARIOS

**This overview of solutions and scenarios for the near or longer-term future is developed as part of LEAP. As a next step, the LEAP coalition will focus on a number of solutions brought forward in this document.**

LEAP believes that further developing distributed concepts has the potential to create new connections and collaborations. In addition, LEAP is identifying which specific (technological) solutions need extra support.

The energy used for production and a material's scarcity, toxicity or waste has not been factored into these solutions and scenario's. Circular ICT hardware is an important topic for LEAP, where trade-off between energy efficiency and hardware waste are taken into account as well as focus on awareness-building and monitoring strategies.

In order to progress, evidence, awareness, transparency and accountability throughout the whole value chain are needed and collaboration to build a sustainable future of digital infrastructures. LEAP's role is to mobilize partners to collaborate and stimulate others to move forward. LEAP is a coalition of the willing, and every LEAP partner—business, government and knowledge institution—participates from a shared and well-understood self-interest, with a desire to collaborate and the ability to connect.

# About LEAP

**LEAP is a platform to support the future generation of data centres. Its aim is to speed up the transition to a sustainable digital infrastructure by deploying and accelerating existing and new technologies that contribute to energy efficiency and integration into the power system and landscape, and crucially with the circular use of materials.**

LEAP is a coalition of data centres, commercial end users, telecom providers, hardware suppliers, network organisations, universities and research centres, startups and governments in the Amsterdam Metropolitan Area (AMA); a crucial internet hub for the international data centre sector as it strives for national, European and global impact.

Because LEAP believes that new public-private partnerships are imperative, we are building an international platform for collaboration, knowledge exchange and communication. LEAP aims to make a positive contribution to a green, smart and energy efficient society in which growth, the environment, people and community go hand in hand, vital materials are preserved, and the Netherlands maintains its lead in innovation for a future-proof digital infrastructure.

**Visit the LEAP website.**