Deloitte.

Together makes progress

Future of infrastructure

Deloitte Luxembourg's point of view

November 2025

The world is changing, and infrastructure—like all technology—is evolving with it.

TODAY'S CHALLENGES

- **Geopolitical constraints:** Fragmented data sovereignty and restrictive regulations.
- **Resource limitations:** Shortages in scalable compute, storage, and energy.
- **Operational inefficiency:** Challenges in managing and integrating multi-cloud and edge systems.
- **Cybersecurity risks:** Rising complexity in defending against advanced threats.
- **Talent gap:** Shortage of skilled professionals for evolving IT demands.
- **Data sovereignty risks:** Limited control and visibility over cross-border or third-party data.

“

There is a need to **build IT infrastructure as a sovereign, AI-ready, sustainable, and intelligent ecosystem** to drive automation, trust, and resilience.

TRENDS SHAPING THE FUTURE

- **Autonomous infrastructure/AIOps:** AI-driven automation enabling self-healing systems and faster deployments.
- **Sustainable operations:** Innovations promoting energy-efficient and eco-friendly infrastructures.
- **Edge computing:** Seamless integration across hybrid, edge, and multi-cloud environments.
- **Quantum technologies:** Advancing quantum-safe security and exploring cloud-based quantum computing for optimization and simulation.
- **AI-Ready infrastructure:** Architectures optimized for compute- and energy-intensive AI workloads.

© 2025 Deloitte Tax & Consulting, SARL

As IT infrastructure evolves, teams must focus on several key aspects.

Role of Infrastructure

INFRASTRUCTURE 1.0 ▶ **INFRASTRUCTURE 2.0** ▶ **INFRASTRUCTURE 3.0**

Keep-the-lights-on foundation

- **Stability over speed**
- On-premises, manual runbooks, ticketed changes, perimeter security, CapEx-heavy.
- Success = uptime, control, and risk mitigation.

Cloud-first growth engine

- **Agility and scale as strategic levers**
- Hybrid/multi-cloud, DevOps/SRE, identity/Zero Trust, OpEx/FinOps.
- Success = velocity + reliability + cost-to-serve

Intelligent platform

- **Acceleration on business outcomes**
- Platform engineering, AIOps and ZeroOps, IaC and autonomous ops, AI/edge-ready
- Success = time-to-value + sovereign resilience

Key considerations to define the road ahead

Eight topics to build the future of Infrastructure

Put digital sovereignty at the core

Geopolitical concerns and rising regulatory demands raise the criticality of sovereignty.

Build AI-optimized infrastructure

AI is reshaping compute and energy demand while enabling dynamic infrastructure management.

Embrace hybrid and edge architectures

Design an architecture leveraging cloud, on-premises, and edge to meet latency and data-locality requirements.

Make cybersecurity and resilience non-negotiable

With threats escalating, security cannot be bolted on.

Automate and observe relentlessly with ZeroOps

Scale and complexity require end-to-end intelligent observability.

Elevate financial governance with FinOps

The economics of cloud and AI make cost optimization a strategic concern.

Reinvent the operating model

The future of IT infrastructure cannot be managed with yesterday's ways of working.

Drive sustainability by design

The carbon footprint of digital infrastructure is under scrutiny.

© 2025 Deloitte Tax & Consulting, SARL

1. Put digital sovereignty at the core

Digital infrastructure no longer operates in a borderless environment. Sovereignty has become a guiding principle as governments tighten rules on data residency, critical infrastructure, and technology supply chains. Initiatives like the EU's GAIA-X, GDPR, and Digital Markets Act reflect Europe's push for autonomy and trusted frameworks. Meanwhile, laws such as the U.S. CLOUD Act and China's Cybersecurity Law show how data and digital asset control is increasingly politicized. For multinational enterprises, this creates a fragmented landscape where compliance, trust, and operations must be carefully balanced.

Why it matters?

Regulatory compliance and market access
Failure to meet sovereignty rules risks fines, reputation loss, and exclusion from key markets.

Resilience and security
Over-reliance on a single hyperscaler, vendor, or region increases exposure to geopolitical, trade, supply, and cyber risks.

Trust and competitiveness
Customers, citizens, and partners expect assurance that their data and services comply with national regulations and remain under local governance.

Strategic autonomy
In critical sectors like healthcare, finance, and public services, sovereignty safeguards long-term independence and control over the digital future.

What should be done?

Data architecture and residency
Design, manage, and govern data flows, ensuring compliance with residency rules.

- Data flow mapping and residency management
- Sovereign cloud adoption and operation
- Multi-region storage design and deployment
- Critical data governance

Vendor and cloud diversification
Reduce dependency on single providers.

- Multi-cloud and vendor diversification
- Contractual transparency management
- Vendor exit and lock-in mitigation
- Interoperability and open standards

Governance and compliance by design
Embed sovereignty requirements into architecture, procurement, and operations.

- Sovereignty-by-design governance
- Regulatory compliance automation
- Security, observability and audit alignment

Supply chain resilience
Secure supply chains and prepare for geopolitical disruptions.

- Supply chain risk assessment
- Trusted vendor sourcing and certification
- Geopolitical contingency planning

Strategic alignment with business and regulators
Position sovereignty as a business advantage and build trust.

- Sovereignty as a business enabler
- Regulator and consortium engagement
- Customer trust and transparent commitment

© 2025 Deloitte Tax & Consulting, SARL

2. Embrace hybrid and edge architectures

The IT landscape has become decentralized. Enterprises now operate across on-premises, multi-cloud, and edge environments, making hybrid models the standard. This approach blends cloud scalability and innovation with on-premises control and governance for sensitive workloads. Meanwhile, edge computing is accelerating as IoT devices, connected vehicles, and real-time analytics drive the need to process data closer to its source. The result is a highly distributed infrastructure requiring seamless orchestration to ensure performance, resilience, and trust at scale.

Why it matters?

Latency and performance
Real-time applications (e.g., IoT, autonomous vehicles, healthcare monitoring, manufacturing) require near-instant edge processing.

Flexibility and agility
Hybrid architectures allow organizations to deploy workloads in the most suitable environment: cloud for scale, on-premises for control, edge for responsiveness.

Resilience and continuity
Distributed architectures reduce single points of failure and strengthen business continuity.

Regulatory and sovereignty compliance
Hybrid models support compliance by ensuring sensitive data and workloads remain within required jurisdictions.

Cost optimization
Balancing workloads across cloud, edge, and on-premises environments enables organizations to optimize cost-performance trade-offs.

What should be done?

Workload placement strategy
Classify workloads by latency, sensitivity, cost, and compliance, and deploy them in the most suitable environment.

- Workload classification and placement framework
- Hybrid deployment decision model

Orchestration and interoperability
Adopt container orchestration platforms, standardize APIs, and ensure E2E observability across environments.

- Multi-cloud
- Container orchestration
- API and interface standardization
- Cross-environment observability

Edge infrastructure
Deploy edge nodes for local AI inference, strengthened hardware, and automated remote management for distributed assets.

- Edge node deployment
- Local AI inference enablement
- Remote edge asset management
- Remote edge automatic patching

Zero-trust across distributed system
Apply identity-based access control, encryption, and continuous monitoring to all endpoints and environments.

- Identity-based access management
- End-to-end encryption
- End-to-end monitoring
- Distributed Compliance Enforcement

Cost and performance
Use FinOps practices and policy-driven workload placement to optimize cost-to-value across environments.

- Hybrid/Edge FinOps cost tracking
- Policy-driven workload optimization
- Cost-to-value performance alignment

© 2025 Deloitte Tax & Consulting, SARL

3. Automate and observe relentlessly with ZeroOps

Modern IT infrastructures—hybrid, multi-cloud, and edge-distributed—have grown too complex for manual management. Traditional monitoring tools cannot keep pace with workload scale, rapid change, or sophisticated threats. Automation and observability have become essential: automation streamlines operations and reduces errors, while observability provides real-time visibility and insights across distributed systems. Together, they enable proactive, efficient, and resilient infrastructure management.

Why it matters?

Complexity management
Hybrid and edge infrastructures generate volumes of operational data that cannot be managed manually.

Reliability and performance
Observability delivers end-to-end visibility, ensuring systems perform as expected and issues are quickly resolved.

Speed and agility
Automation speeds up provisioning, scaling, and remediation, enabling faster innovation and time-to-market.

Cost control
Automated optimization reduces overprovisioning and idle resource waste, improving efficiency and savings.

Resilience
Proactive monitoring with automated responses reduces downtime and lowers mean time to recovery.

What should be done?

Infrastructure as Code (IaC)
Adopt a tool such as Terraform to codify provisioning and configuration.

- Infrastructure as code
- Standardized environment configuration
- Compliance enforcement

End-to-end observability
Instrument logs, metrics, and traces across the stack. Deploy unified observability platforms and AI/ML to detect anomalies.

- Full-Stack telemetry (logs, metrics, traces)
- Unified observability platform
- AI/ML-driven anomaly detection
- AI/ML-driven predictive analytics

Automated remediation
Use AIOps and event-driven automation to resolve incidents without human intervention.

- Event-driven automation and AIOps
- Automated scaling, failover and patch management
- Self-healing orchestration integration (e.g.,Kubernetes)
- Observability integrated with orchestration tools

Security and compliance
Use observability data for audit-ready records and integrate threat detection with automated incident response.

- Automated security policy enforcement
- Audit-ready observability records
- Automated threat detection and incidence response

Data-driven operations culture
Enable teams to use telemetry and automation instead of manual intervention.

- Operations dashboards for IT and business
- Shared visibility between Teams
- Telemetry-driven decision-making culture

© 2025 Deloitte Tax & Consulting, SARL

4. Reinvent the operating model

The traditional IT model—siloed teams, ticket queues, and manual governance—is no longer suited for the digital era. Hybrid cloud, edge, and AI workloads demand speed, resilience, and continuous delivery that legacy structures cannot support. The challenge is both organizational and technical. Infrastructure must evolve from a back-office function to a product and platform. Reinventing the operating model means rethinking teams, processes, and culture to embed agility, automation, and business alignment, transforming infrastructure into a strategic enabler.

Why it matters?

Agility
Legacy operating models slow innovation and create bottlenecks for development and business teams.

Talent utilization
Outdated roles and manual practices waste scarce skills that could be applied to higher-value, automated work.

Governance
Without modernized processes, organizations struggle to balance compliance, cost, and speed.

Cross-functional collaboration
Infrastructure, development, security, and finance must adopt integrated ways of working to deliver at scale.

Strategic relevance
The operating model determines whether infrastructure is a cost center or a business differentiator.

What should be done?

Platform teams
Adopt platform engineering principles to deliver infrastructure as a product.

- Platform team design and management
- Infrastructure-as-product roadmapping and SLAs
- Developer self-service enablement

Continuous governance
Embed security, FinOps, and compliance into workflows, shifting from reactive audits to continuous compliance.

- Embedded security, FinOps and compliance workflows
- Policy enforcement via automation and observability
- Continuous compliance operations

Talent and skills
Embed security, FinOps, and compliance into workflows, shifting from reactive audits to continuous compliance.

- Cloud-native, automation, AI and cybersecurity upskilling
- Automation-driven work
- Talent attraction and retention through innovation
- Spend-to-Value Benchmarking

Data-driven decision making
Leverage observability and financial data to guide resource allocation and prioritization.

- Observability and FinOps-driven decision-making
- Executive dashboards and KPI alignment
- Data transparency for resource prioritization

Collaboration and innovation culture
Form cross-functional teams and foster experimentation with modern architectures.

- Cross-functional teams
- Experimentation with new architectures (serverless)
- Infrastructure team as partners in innovation

© 2025 Deloitte Tax & Consulting, SARL

5. Build AI-optimized infrastructure

AI has moved from experimentation to enterprise-scale deployment. Generative AI, large language models, and advanced machine learning are now central to business strategy across industries. These workloads place unprecedented demands on compute, storage, data throughput, and energy efficiency. Unlike traditional applications, AI training and inference require specialized infrastructure—GPUs, TPUs, accelerators, low-latency interconnects—and massive datasets spread across cloud, edge, and on-premises environments.

Why it matters?

Business growth
AI is a key differentiator for customer engagement, automation, and innovation. Without optimized infrastructure, enterprises risk falling behind.

Performance and scalability
Training and deploying large models requires massive compute and low-latency data pipelines; without scale, AI initiatives stall or underperform.

Financial impact
AI workloads are expensive, and inefficient architectures drive uncontrolled costs and undermine ROI.

Sustainability
AI's energy footprint faces increasing regulatory and public scrutiny, making performance and sustainability a critical balance.

Trust and compliance
Sensitive AI data must meet strict security, privacy, and sovereignty requirements to ensure adoption and stakeholder trust.

What should be done?

High-performance compute
Deploy GPU/TPU clusters or leverage on-demand cloud capacity.

- High-performance compute provisioning
- GPU/TPU cluster deployment and management
- Low-latency interconnect optimization
- On-demand cloud capacity

Robust data pipelines
Enable fast ingestion, processing, and distribution of large datasets.

- Data pipeline design and orchestration
- Scalable data lakehouse
- End-to-end pipeline observability

Balance cost with value with FinOps
Track spend per workload, use auto-scaling and link costs to business outcomes.

- AI FinOps cost tracking and optimization
- Auto-scaling
- Hybrid cloud management
- Spend to value alignment

Energy-efficient and sustainable AI
Align innovation with sustainability goals, such as using renewable-powered data centers.

- Energy-efficient data center operations
- Advanced cooling and workload Scheduling
- Carbon footprint monitoring and reporting

Security and compliance
Apply encryption and zero-trust principles while ensuring data meets sovereignty and regulatory requirements.

- Zero-trust pipeline enforcement
- Regulatory and sovereignty compliance for AI Data
- AI Disaster Recovery
- AI resilience planning

© 2025 Deloitte Tax & Consulting, SARL

6. Make cybersecurity and resilience non-negotiable

As IT infrastructure becomes hybrid, cloud-based, and edge-distributed, the attack surface grows significantly. Organizations face sophisticated ransomware, supply chain breaches, and state-sponsored attacks. Rising downtime costs threaten business continuity, customer trust, and national security in critical sectors. In this environment, cybersecurity and resilience are no longer optional—they are strategic foundations for infrastructure in a volatile digital era.

Why it matters?

Business continuity
Even brief disruptions can trigger major financial losses, reputational damage, and regulatory penalties.

Trust and brand protection
Customers and partners expect systems and data to always remain secure and available.

Regulatory pressure
Governments are imposing stricter requirements around critical infrastructure protection and incident disclosure.

Evolving threat landscape
Attackers now leverage automation, AI, and supply chain infiltration, rendering traditional defenses insufficient.

Geopolitical tensions
Infrastructure is increasingly targeted in hybrid warfare, making resilience a national and economic security priority.

What should be done?

Zero-trust architecture
Assume breaches, enforce identity-centric security, and continuously validate access for users, devices, and workloads.

- Zero-trust identity and access enforcement
- Continuous session monitoring
- Policy enforcement automation

Resilience and continuity
Design redundant, distributed systems. Implement automated multi-region disaster recovery with automated failover.

- Redundant and distributed architecture design
- Automated multi-region disaster recovery
- Regular business continuity and recovery testing
- Automated failover

Security across environments
Apply consistent policies across environments, including endpoints and IoT, with unified observability.

- Hybrid and edge security policy extension
- Endpoint and IoT device protection
- Cross-environment observability and monitoring
- Remote edge automatic patching

Automation and AI in defense
Leverage AIOps, SIEM/SOAR, and ML to detect, predict, and remediate threats in real time.

- AI/ML-drive anomaly detection
- Real-time incident detection and response (SIEM/SOAR)
- Automated patch and vulnerability management

Governance and compliance
Align continuously with evolving regulations and ensure audit-ready controls.

- Regulatory compliance alignment
- Audit-ready security control visibility
- Board-level resilience reporting

© 2025 Deloitte Tax & Consulting, SARL

7. Elevate financial governance with FinOps

Cloud and AI adoption have reshaped IT economics. Traditional CapEx models are being replaced by variable, consumption-based spending across multiple providers, services, and business units. This offers flexibility but introduces unpredictable costs, hidden charges, and potential overspend. AI and data-intensive workloads, with their massive demand for compute, storage, and energy, amplify the challenge. Financial discipline is now as critical as technical excellence. FinOps unites engineering, finance, and business leaders to deliver cost visibility, accountability, and optimization as a core capability.

Why it matters?

Cost transparency
Cloud and AI consumption models are complex, and without visibility, overspending is inevitable.

Business alignment
Infrastructure spending should drive business outcomes and value, not be treated as mere overhead.

Scalability and innovation
FinOps enables organizations to innovate at speed while keeping costs predictable and sustainable.

Investor and management expectations
Boards and investors require financial discipline and clear ROI to justify digital transformation investments.

AI economics
AI workloads demand massive compute and storage, and without governance, costs can quickly become unsustainable.

What should be done?

Cost visibility and accountability
Deploy platforms for real-time spend visibility and allocate costs by workload and team.

- Cloud and hybrid cost visibility platforms
- Workload and team-level cost allocation
- Dashboards with tailored insights

FinOps operating model
Form a cross-functional FinOps team with clear roles, responsibilities, and review cycles.

- Cross-functional FinOps team
- Defined roles and accountability
- Regular cost review
- Cost alignment cycles with business

Optimization practices
Apply auto-scaling and rightsizing to continuously identify and eliminate underutilized resources.

- Resource rightsizing and auto-scaling
- Reserved and spot instances optimization
- Automated decommission of idle resources
- Spend-to-value benchmarking

Observability and automation
Connect financial and operational telemetry for cost-per-service insight.

- Cost-integrated observability platforms
- Policy automation for cost efficiency
- ML-driven spend forecasting and anomaly detection

Educate and align culture
Train engineering teams on cost-aware design and operations.

- Cost-awareness trainings
- Incentives for efficiency improvements
- Innovation Enabled Through Financial Governance

© 2025 Deloitte Tax & Consulting, SARL

8. Drive sustainability by design

The rapid growth of cloud services, AI workloads, and data-intensive applications is driving a surge in IT energy use and carbon emissions. Data centers already consume a significant share of global electricity, with demand set to rise further due to AI and edge computing. Meanwhile, regulators, investors, and customers are pressing for stronger environmental responsibility. Sustainability has thus become a core design principle for modern infrastructure, not a secondary consideration.

Why it matters?

Positive impact and responsibility
Sustainable infrastructure cuts carbon emissions, conserves resources, and aligns technology growth with social benefit.

Regulatory compliance
Governments are imposing stricter energy-efficiency standards and sustainability disclosure requirements for digital infrastructure.

Cost efficiency
Energy is one of the largest operational expenses in IT and efficiency gains deliver immediate benefits.

Investor and market expectations
Strong sustainability performance increasingly affects valuation, brand reputation, access to capital, and appeal to environmentally conscious customers.

Future-proofing
Embedding sustainability ensures long-term, environmentally and financially viable digital growth, including AI and edge computing.

What should be done?

Energy-efficient architectures
Optimize resources with virtualization, containerization, and right-sizing; adopt efficient cooling methods.

- Workload consolidation
- Legacy system modernization
- Server virtualization and containerization
- Advanced cooling and energy-efficient data centers

Renewable energy
Source power from renewable-backed cloud providers and explore benefits from opting for green power.

- Renewable energy sourcing
- Green cloud and colocation centers
- Electricity purchased from green sources

AI and high-performance workloads
Schedule energy-aware workloads and run AI tasks in low-carbon regions; link carbon impact to business outcomes.

- Energy-aware scheduling
- AI workload efficiency monitoring
- FinOps for energy and cost alignment
- Low-carbon workload placement

Measurement and reporting
Monitor and report energy use and emissions at the workload level with dashboards for accountability.

- Carbon and energy observability tooling
- Standards-based sustainability reporting
- Real-time sustainability dashboarding

Sustainable culture
Set sustainability KPIs and use sustainability as a driver of innovation and competitive advantage.

- Sustainability KPIs in Leadership Governance
- Green Coding & Efficient Infrastructure Operations
- Sustainability as a Driver of Innovation

© 2025 Deloitte Tax & Consulting, SARL

To embrace this shift, it is essential to stabilize each layer of the infrastructure...

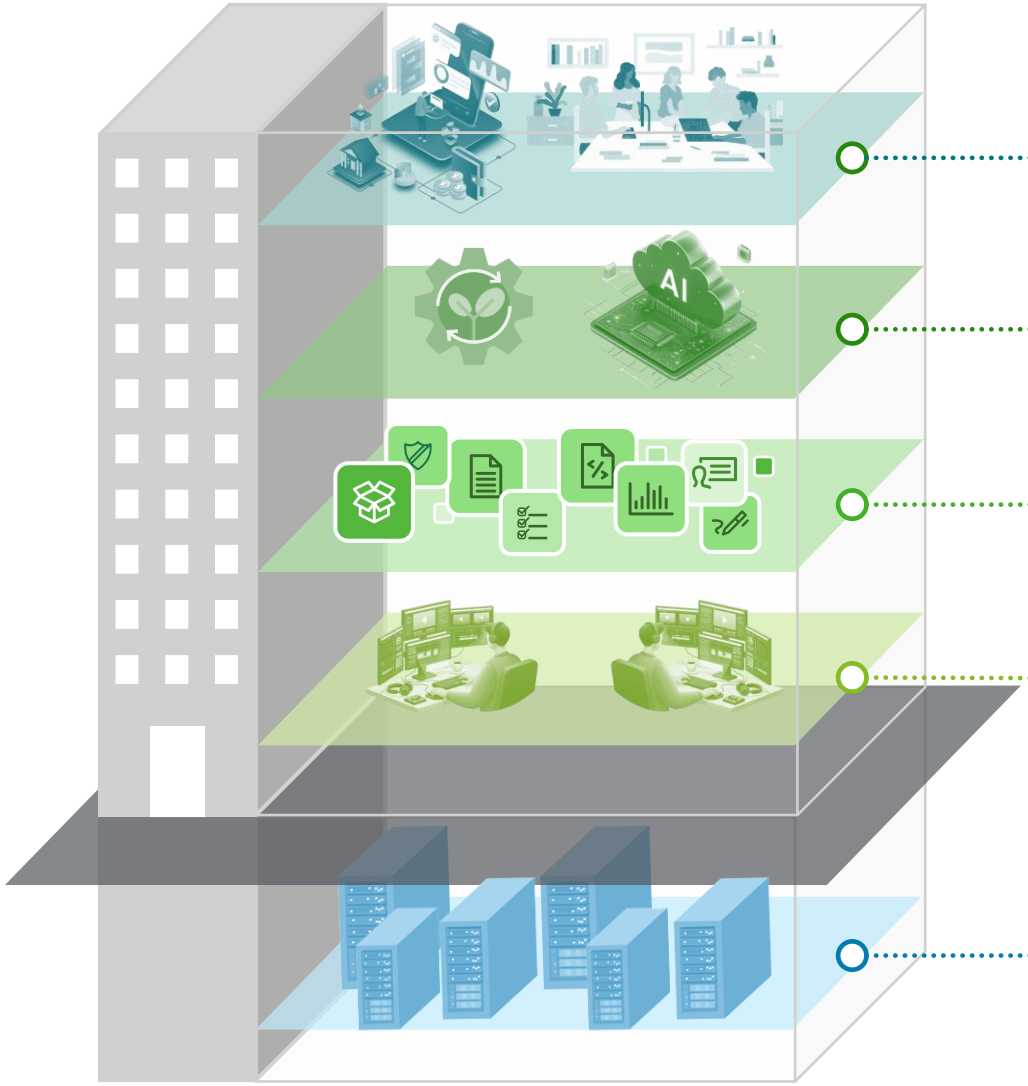

TALENT & OPERATION

Future operations will use adaptive **operating models** and **FinOps** strategies, aligning talent, resources, and cost efficiency for seamless collaboration and optimized infrastructure governance.

TECHNOLOGY ENABLED

AI-driven automation, **ZeroOps**, and **sustainability** will redefine efficiency—maximizing resource use, minimizing overhead, reducing environmental impact, and boosting system intelligence.

APPLICATIONS

Applications will be architected for **hybrid and edge** environments, blending on-site and cloud resources to ensure scalability, low latency, and adaptability.

MONITORING

Infrastructure must be continuously monitored to maintain **resilience**, enhance **cybersecurity**, and enable **observability** that detects and mitigates threats before they affect operations.

PHYSICAL INFRASTRUCTURE

As data centers become increasingly distributed, **sovereignty** becomes a core principle, ensuring secure data handling while meeting compliance and geopolitical requirements.

© 2025 Deloitte Tax & Consulting, SARL

… By implementing interconnected, complementary capabilities supported by the right technologies

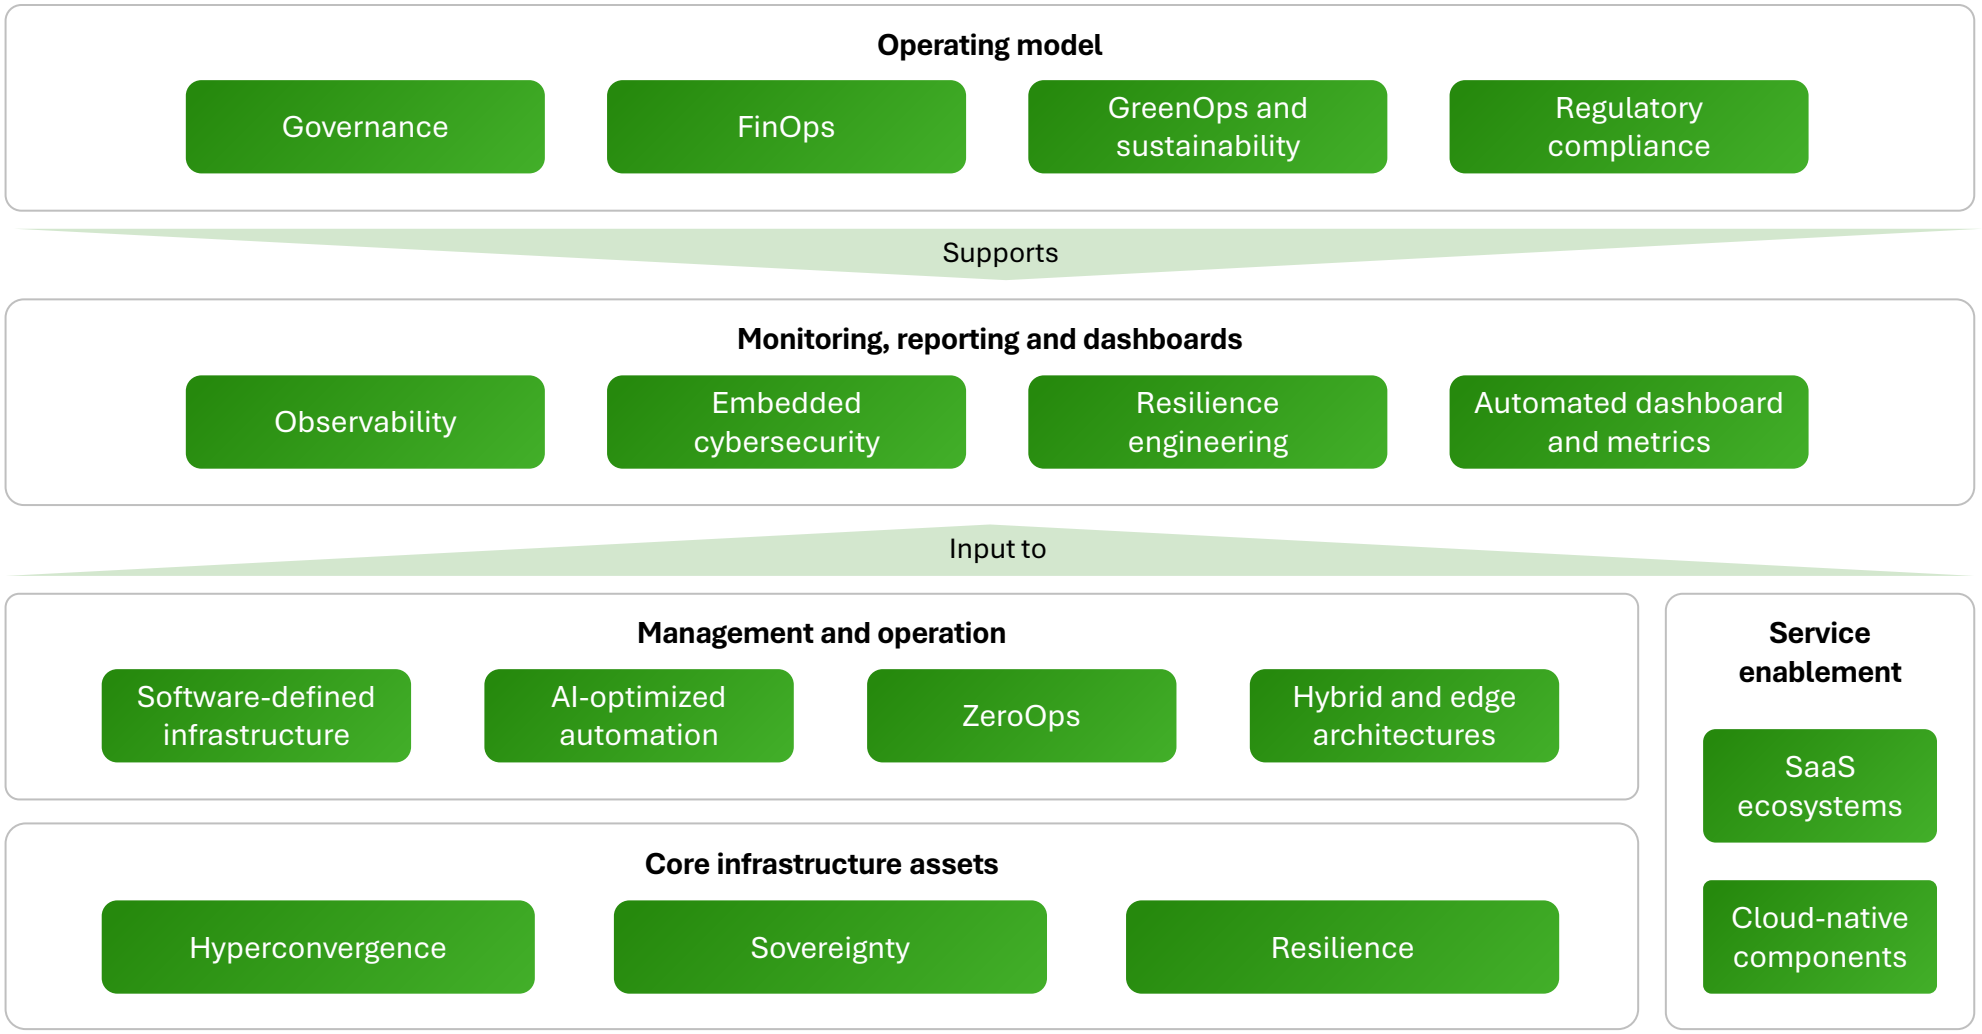

© 2025 Deloitte Tax & Consulting, SARL

How to move forward?

Define your vision...

- Strategy
- Business case definition
- Operating model design
- Discovery and design

Build and migrate...

- Architecture
- Program management
- Quality assurance
- Change management

Operate...

- Continuous improvement

© 2025 Deloitte Tax & Consulting, SARL

Deloitte.

Deloitte refers to one or more of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms, and their related entities (collectively, the “Deloitte organization”). DTTL (also referred to as “Deloitte Global”) and each of its member firms and related entities are legally separate and independent entities, which cannot obligate or bind each other in respect of third parties. DTTL and each DTTL member firm and related entity is liable only for its own acts and omissions, and not those of each other. DTTL does not provide services to clients. Please see www.deloitte.com/about to learn more.

Deloitte provides leading professional services to nearly 90% of the Fortune Global 500® and thousands of private companies. Our people deliver measurable and lasting results that help reinforce public trust in capital markets and enable clients to transform and thrive. Building on its 180-year history, Deloitte spans more than 150 countries and territories. Learn how Deloitte’s approximately 460,000 people worldwide make an impact that matters at www.deloitte.com.

This communication contains general information only, and none of Deloitte Touche Tohmatsu Limited (“DTTL”), its global network of member firms or their related entities (collectively, the “Deloitte organization”) is, by means of this communication, rendering professional advice or services. Before making any decision or taking any action that may affect your finances or your business, you should consult a qualified professional adviser.

No representations, warranties or undertakings (express or implied) are given as to the accuracy or completeness of the information in this communication, and none of DTTL, its member firms, related entities, employees or agents shall be liable or responsible for any loss or damage whatsoever arising directly or indirectly in connection with any person relying on this communication. DTTL and each of its member firms, and their related entities, are legally separate and independent entities.

MAKING AN IMPACT THAT MATTERS since 1845

© 2025 Deloitte Tax & Consulting, SARL