TELECOMMUNICATIONS

Industry Perspective

Distributing building-resident infrastructure at scale.

Deploy the infrastructure once. Expand the services over time. Telecommunications operators already have what distribution at scale requires: field presence, enterprise relationships, managed-service operations, and recurring-revenue economics. The shift is from selling connected endpoints to placing, commissioning, and managing a building-resident infrastructure asset that outlives its first application.

Executive Summary

Telecommunications operators have spent decades building the capabilities required to distribute infrastructure at enormous scale. They procure equipment, finance it, provision it, connect it, install it at customer premises, monitor it remotely, maintain it, upgrade its software, replace it when necessary, secure it, support it under service-level agreements, and recover the investment through recurring customer relationships measured in years rather than individual projects.

Historically, the infrastructure placed at the customer premise has been telecommunications infrastructure: routers, gateways, modems, managed Wi-Fi, SD-WAN appliances, security equipment, private-network infrastructure, and other forms of customer-premises equipment. The next opportunity is to extend that operating model into the building itself.

There is a reason to move now, and it is visible in the operators' own economics. GSMA reported that NB-IoT and LTE-M together passed one billion low-power wide-area connections by the end of 2025 [1]. Across the broader cellular IoT connectivity market, Berg Insight estimated 4.2 billion connections in 2025, up 11 percent, while connectivity revenue grew 5 percent and average monthly revenue per connection fell 7 percent to EUR 0.31 [2]. Connectivity volume is scaling faster than connectivity value.

The buildings behind those connections represent a much larger operating environment. In EIA's most recent completed Commercial Buildings Energy Consumption Survey, the 5.9 million U.S. commercial buildings operating in 2018 represented approximately 96.4 billion square feet and about $141 billion in annual energy expenditure [3]. Energy is only one budget category, but it is a large, measurable operating-cost pool against which building applications can create value.

LEXI allows a telecommunications provider to distribute a managed building-resident infrastructure asset: a persistent, protocol-independent managed edge that connects physical systems already installed in a customer property and becomes the foundation for building automation, energy management, equipment monitoring, environmental intelligence, local automation, distributed inference, and future applications.

The telco does not need to become a building-controls contractor. It does not need to replace the building management system. It does not need to manufacture every sensor, application, or AI model.

And it does not need to own the real estate. Its strategic role is different: distribute the infrastructure, operate it, manage its lifecycle, and use that footprint to deliver an expanding portfolio of services.

> **A project ends. Infrastructure remains.**

The Opportunity Is Infrastructure Distribution, Not Another IoT Solution

Telecommunications companies already participate extensively in enterprise IoT, and the scale question is settled. The remaining problem is that connected devices, managed as endpoints, do not add up to an understood building.

Managed IoT is organized around endpoints: connect a sensor, meter, tracker, camera, or machine and manage the resulting device data. That works when the endpoint itself is the service boundary. Buildings are not built that way. HVAC equipment may communicate through BACnet, meters through Modbus, environmental sensors through LoRaWAN, lighting through Zigbee or proprietary controls, while locks, pumps, boilers, refrigeration equipment, occupancy sensors, and cameras sit in their own operational domains. A cellular connection to any one of them does not resolve the fragmentation.

A temperature sensor illustrates the difference. On its own it reports a number. It becomes operationally useful when the platform knows which space it represents, which equipment serves that space, whether the space is occupied, what that equipment is doing, what outdoor conditions are, whether a door is open, whether the equipment is responding normally, and whether anything has authority to act on the answer.

LEXI provides the telco with a universal managed edge designed to connect heterogeneous building infrastructure, normalize its data into a common operating model, run local applications and automations, maintain local operations when the wide-area link is unavailable, and connect the result securely to cloud services.

The transition is from managed IoT to managed physical infrastructure. The operator stops asking how many devices it can connect and starts asking what infrastructure it can place that lets it deliver services to the building for the next decade.

Why This Is Not the Last IoT Program

Any operator senior enough to approve a program of this kind has watched adjacent IoT initiatives retreat, restructure, or disappear. The record should be acknowledged directly rather than left for diligence to discover.

Ericsson transferred its IoT Accelerator and Connected Vehicle Cloud businesses to Aeris in 2023. Ericsson had forecast roughly SEK 0.8 billion in annual sales for the businesses and said the transaction would remove approximately SEK 0.25 billion of quarterly losses; Ericsson also said that, despite significant investment, it had captured only a limited part of the IoT value chain [4][5]. That is a value-chain lesson: connectivity management can scale without capturing enough of the surrounding economics.

AT&T's Digital Life home-security service illustrates a different risk. The service launched in 2013, reached hundreds of thousands of subscribers, and was wound down as AT&T retired the 3G network generation on which much of the installed controller base depended [6]. The lesson is not that home security was a bad market; it is that customer-premises infrastructure tightly coupled to a communications generation can create expensive migration decisions when that generation ends.

AT&T's later decision to stop selling and certifying NB-IoT service and to move customers toward LTE-M illustrates a third risk: a carrier can retire an access technology underneath live enterprise deployments [7]. The radio layer itself is not permanent.

These exits were not identical, but they illustrate three recurring risks: insufficient participation in the value chain, application infrastructure tied too closely to a technology generation, and service economics that do not survive changes in the underlying connectivity layer.

The model proposed here is deliberately different. The operator distributes and manages an infrastructure layer that is not tied to one radio access technology, one building protocol, or one application. A building edge that can connect BACnet, Modbus, LoRaWAN, Zigbee, Thread, Bluetooth, Wi-Fi, Ethernet, and cellular survives a radio sunset for the same reason it survives a chiller replacement or BMS upgrade: the infrastructure layer is designed to outlast the individual technologies around it.

From Customer-Premises Equipment to Building-Resident Infrastructure

The concept is familiar to operators, but the distinction matters. Traditional customer-premises equipment is usually deployed to enable a defined telecommunications service. A LEXI deployment establishes building-resident infrastructure intended to support multiple services, applications, device classes, and compute workloads over time.

A router inside an enterprise location is valuable because it enables connectivity, security, management, monitoring, support, and software lifecycle services. The equipment is an access point into a larger managed service. LEXI preserves that logic and extends it into the physical building.

The first use case can justify entry, but it should not define the lifetime value of the deployment. If a customer buys a point solution for energy monitoring, the commercial opportunity can largely end with energy monitoring. If the same customer receives a managed infrastructure layer on which energy is the first application, the relationship can expand into diagnostics, automation, environmental intelligence, additional sensing, local inference, and other services without standing up a separate technology stack each time.

> **Deploy the infrastructure once. Expand the services over time.**

What Makes LEXI an Infrastructure Asset

Calling something an infrastructure asset should mean more than putting hardware in a building. Four characteristics have to hold: persistence, multi-use capability, manageability, and extensibility.

The economic asset is not the gateway alone. It is the installed and commissioned footprint: building interfaces, connected-device relationships, the normalized data model, secure connectivity, local compute capacity, application environment, accumulated operating history, and the managed-service relationship around them. Once that footprint exists, additional services can be activated at materially lower incremental deployment cost than establishing another independent technology stack.

That is why the operator should evaluate lifetime revenue and margin per deployed building infrastructure asset rather than hardware margin. The value is created by the stream of services the footprint can support across its useful life.

Why the Telco Is a Natural Distributor

Building technology has historically been sold through channels optimized around construction: controls contractors, mechanical and electrical contractors, systems integrators, and engineering firms. That model works for large construction and controls projects. It is less suited to placing a common managed infrastructure layer across hundreds or thousands of geographically distributed enterprise properties.

Telecommunications companies are organized for that problem. A large operator already has enterprise account relationships, national or multinational sales coverage, connectivity, customer-premises logistics, provisioning, network operations, remote monitoring, cybersecurity, service assurance, field support, recurring billing, financing capability, hardware lifecycle management, cloud and managed services, and partner ecosystems.

Private-network development points in the same direction. Operators increasingly deliver private-network environments directly or integrate them with public mobile infrastructure [8], while hybrid approaches let enterprises consume managed connectivity without becoming the network operator themselves [9]. LEXI applies that managed-infrastructure logic one layer deeper, inside the physical building.

The Telco Does Not Have to Become a Controls Contractor

The telecommunications provider should not attempt to recreate the traditional building-automation contracting business. The operator can own the customer relationship, managed service, connectivity, lifecycle management, service assurance, billing, and recurring commercial model while certified installers, systems integrators, mechanical contractors, electrical contractors, or other specialists perform the field work a particular building requires.

Installation is an enablement function rather than the product. A conventional contractor is paid primarily when something is installed. A managed-service provider is paid because the infrastructure continues to operate. That aligns the operator with reliability, retention, application expansion, additional locations, lower support cost, and longer infrastructure life.

What Repeatable Deployment Actually Requires

The largest strategic risk is treating LEXI as another item in a professional-services catalog. If every deployment requires custom engineering, bespoke software, unique integrations, negotiated field work, and one-time project pricing, the operator inherits the economics of a systems integrator rather than a managed-service provider.

Six measurable conditions determine whether this is a product or a project:

- Physical placement and power - Define mounting, power, radio reach, enclosure, and any licensed-trade or permit requirements that affect per-site cost.
- Remote-survey rate - Measure how often a site can be quoted and planned without an engineer walking the building. A rollout that requires a custom pre-site visit everywhere will not scale economically.
- Pre-built integration ratio - BACnet and Modbus are protocols, not integrations. Track the percentage of the customer's actual rooftop units, chillers, meters, controllers, and other equipment that can be commissioned from a catalog profile rather than an engineering ticket.
- Installer labor model - Determine whether deployment uses operator field personnel, certified integration partners, licensed electrical or mechanical trades, customer technicians, or a combination. Labor rules, licensing requirements, and jurisdictional boundaries can materially affect per-site cost.
- Regional certification - Establish RF, safety, power, and market certification requirements before committing a multinational rollout schedule.
- Commissioning and handover - Define what is verified at turn-up, by whom, what evidence is retained, and what the service-assurance organization is expected to support for the life of the contract.

Some buildings will always require engineering. The objective is to concentrate engineering at exceptions rather than embed it in every deployment. The commercial goal is not to maximize installation revenue; it is to minimize the friction of placing the next managed infrastructure asset.

One Installed Asset, Multiple Revenue Layers

The same physical infrastructure can support several layers of recurring revenue. Connectivity attaches directly to the deployment, but the 2025 cellular IoT market data show why it should be treated as an attach layer rather than the return: Berg Insight estimated average connectivity revenue at about EUR 0.31 per cellular IoT connection per month, down 7 percent year over year [2].

Layer	What sits there	What the operator sells
Cloud and enterprise services	Portfolio management, data retention, analytics, dashboards, reporting, enterprise and billing integration	Cloud subscription and portfolio services that expand with site count
Telco network and operator edge	Wide-area transport, private networking, operator MEC where deployed	Connectivity and private network services - important attach, but not the primary return
LEXI building edge	Protocol normalization, building context, local automation, device and gateway management, application runtime, local inference	Managed infrastructure, building applications, and later governed hosted workloads
Building systems and devices	HVAC, meters, lighting, refrigeration, locks, pumps, occupancy and environmental sensing	Integration and added sensing where information is missing

An operator should model lifetime economics per building infrastructure asset rather than gateway margin. The relevant variables are infrastructure and installation cost per site; initial recurring revenue; application attach rate at 12 and 24 months; incremental gross margin per added service; incremental cost to activate another application; customer and site retention; lifecycle support cost; and contract term relative to asset life.

Conceptually, the underwriting logic is straightforward: infrastructure return equals recurring service margin plus application expansion plus connectivity and cloud attach, less deployment and lifecycle cost, measured across the enforceable contract period. The contract term matters because a long-lived asset recovered through a short service agreement creates a different risk profile than the same asset under a term aligned with its useful life.

> **The key metric is not gateway margin. It is the lifetime revenue and margin generated by the deployed infrastructure footprint.**

Connectivity Becomes Part of a Larger Architecture

Cellular remains strategically important, but in this model it is one component of a broader architecture rather than the end product. A building may use cellular for resilient wide-area connectivity, LoRaWAN for long-range low-power sensing, Wi-Fi for higher-bandwidth devices, Ethernet for existing wired infrastructure, BACnet for building systems, Modbus for meters and equipment, and Zigbee, Thread, or Bluetooth for local wireless devices.

The operator does not need every endpoint to become cellular. It needs a way to manage the heterogeneous physical environment and use cellular where it creates the most value: wide-area reach, operational independence from the customer network where appropriate, resilience, and an operator-managed communications path.

Instead of asking how many SIMs can be sold into a building, the strategic question is how much of the building's digital infrastructure the operator can manage.

The Edge Is the Trust Boundary

An operator proposing to install a remotely updatable, compute-capable device adjacent to a customer's building systems will be evaluated by that customer's security organization before it reaches scale. That review should be designed for rather than discovered.

Four questions decide it:

- Network placement and segmentation - The edge should support an isolated segment and tightly controlled outbound connectivity. An operator-provided cellular path can simplify segmentation, reduce dependency on the customer LAN, and provide an independently managed communications path in addition to its resilience benefits.
- Software integrity - Signed and verified firmware, staged OTA rollout with rollback, auditable update history, vulnerability disclosure, and a defined patch commitment are central to the service model.

- Control authority - Observing, diagnosing, recommending, initiating a workflow, and actuating are different authorities. Write access to building equipment should be explicitly granted and configured rather than assumed.
- Third-party workload governance - Hosted workloads should be isolated from one another and from the platform, with declared access to data and control points, a certification path, and defined operational responsibility.

The operator's exposure differs in kind from that of a single enterprise customer. A carrier managing tens of thousands of building edges is a systemic target. Security architecture therefore belongs inside the product and sales model, not in an appendix after the commercial decision.

LEXI Building Edge and Operator MEC Are Complementary

ETSI defines Multi-access Edge Computing as cloud-computing capabilities and an IT service environment offered at the edge of the network [10]. ETSI White Paper No. 59 goes further in describing central-cloud, edge-cloud, and far-edge tiers and the role of edge nodes acting as gateways to industrial equipment bus systems [11].

LEXI creates a complementary building-resident edge, not another name for operator MEC. Operator MEC brings compute closer to customers through telecommunications infrastructure; LEXI brings managed connectivity, protocol normalization, building context, application execution, local resilience, and governed compute inside the physical property itself.

The resulting continuum is straightforward: building systems -> LEXI building edge -> operator network or MEC -> cloud and enterprise services. Workloads are placed deliberately. A fast equipment anomaly may be evaluated locally; portfolio benchmarking may run centrally; a local automation should continue during a WAN outage; a privacy-sensitive workload may stay on premises.

Distributed Intelligence Is an Expansion Layer

Distributed intelligence means deliberately placing analytics and AI workloads across the building edge and centralized infrastructure according to latency, resilience, bandwidth, data locality, privacy, compute intensity, and operational authority. It is not simply a synonym for running AI at the edge.

For the operator, distributed intelligence is an expansion layer on infrastructure already installed. The physical asset does not have to be replaced each time the service stack moves upward. That same property allows the infrastructure to outlast a radio generation, application generation, or equipment refresh.

The Asset, the Term, and the Exit

The model does not require a single ownership structure. The operator does not need to own the asset to control the managed-service relationship around it. What matters is clarity about who funds the infrastructure, where the return sits, what term supports the economics, and what happens when the relationship ends.

Structure	Who funds the asset	Where the return sits	Principal risk
Operator-owned managed infrastructure	Telco or affiliated financing vehicle	Recurring managed-service margin across the asset life	Contract term shorter than asset life
Customer-funded infrastructure	Enterprise customer	Connectivity, cloud, lifecycle management, and applications	Lower capital barrier for operator; services may be re-tendered
Embedded in broader agreement	Included in private-network, managed-IoT, SD-WAN, energy, or transformation agreement	Attach and retention on the parent contract	Infrastructure value can be obscured or underpriced
Property or portfolio infrastructure	Real-estate owner	Management fee plus building, tenant, and portfolio services	Data governance across owner, operator, and tenants

Exit terms should be settled before the first deployment rather than the first renewal: hardware ownership and transfer rights, data ownership and export rights, treatment of the normalized building model and operating history, removal or transfer costs, and the operator's obligations if the customer moves the service elsewhere.

Persistence is valuable, but it should arise from accumulated utility rather than artificial lock-in. Every connected subsystem, application, and year of operating history should make the infrastructure more useful. The customer should stay because replacing that utility is unattractive, not because leaving is contractually or technically impossible.

Repeatability Is the Economic Argument

Traditional smart-building deployments often restart at every property: new survey, new integration, new contractor, new commissioning, new software. A distributed-infrastructure model changes the sequence. The expensive step is establishing the managed building edge. After that, incremental applications increasingly become software, configuration, added sensing, or workloads deployed through infrastructure that already exists.

The economics are therefore strongest with portfolio customers - retailers, hotel groups, senior-living operators, commercial real-estate owners, multifamily owners, industrial and warehouse operators, healthcare systems, campuses, and public-sector estates. The unit of sale can move from one building to one portfolio once the initial sites establish installation cost, integration ratio, security acceptance, operational value, and repeatability.

The operator does not need to predict every future application when the infrastructure is installed. That is the point of infrastructure. A fiber connection installed years ago supports services that did not exist when it was placed; a building-resident edge should be evaluated through the same lens.

White Label or Integrated Upward

Different operators will want different degrees of control over the customer experience. An operator may deliver a complete white-label service or integrate LEXI infrastructure and data upward into the enterprise portal, IoT platform, network-management environment, data platform, operations center, billing system, and support environment it already runs.

Current operator offerings show that much of the telco-side stack already exists. Singtel, for example, describes managed services spanning data center, hybrid cloud, IoT networks, automation, analytics, and unified monitoring [12]. Its Managed IoT Intelligence offering combines sensors, connectivity, analytics, rule-based workflows, onboarding, provisioning, firmware management, a secure cloud dashboard, and energy and carbon monitoring; Singtel also positions the service across LoRaWAN, cellular, and wired environments [13].

That is not a reason an operator should rebuild LEXI. It clarifies the complementary layers. The operator already has the managed-service, connectivity, device-lifecycle, portal, support, billing, and enterprise relationship layer. LEXI adds the building-side infrastructure beneath it: southbound connection to existing BACnet, Modbus, and heterogeneous building equipment; normalization of existing equipment and new sensing into a common operating model; local automation and resilience; a governed building-resident runtime; and distributed inference tied to physical-building context.

The simplest formulation is: the telco provides the operator layer; LEXI provides the building layer.

The Telco Controls Distribution; LEXI Supplies the Infrastructure Platform

A successful partnership depends on clear responsibilities and on clarity about what is available today.

Telecommunications provider	LEXI	Integrators and specialists
Enterprise customer acquisition **Commercial packaging and pricing** **Connectivity** **Customer-premises distribution** **Managed-service delivery** **Billing and contracting** **Service assurance** **First-line support** **Financing where offered** **Field-service coordination** **Enterprise cloud/network integration** **Regional and international scale**	Universal managed edge technology Multi-protocol building connectivity Device and gateway management Secure provisioning Remote software and OTA management Local automation Data normalization and building context Cloud platform capabilities Application deployment Portfolio intelligence Edge inference orchestration Governed control architecture Continued platform development	Equipment integration Mechanical and electrical work Site engineering Sensor installation Controls commissioning Specialized industrial or refrigeration integration Local regulatory and trade support

An operator evaluating the partnership should require each LEXI capability to be identified as generally available, in development, or architectural in the diligence package. A positioning paper should explain the model; the diligence package should define what ships now, what is committed, and what remains on the roadmap.

Where to Start

The objective of an initial deployment is not to prove that the technology works in one building. It is to prove that the distribution model works repeatedly across representative buildings.

A pilot should answer seven questions:

1. Can the infrastructure be deployed repeatedly? Measure field labor hours, commissioning time, integration effort, connectivity requirements, permit or trade requirements, and exceptions per site.

2. What share of the customer's equipment is a catalog integration? Measure the pre-built integration ratio explicitly; this ratio is a primary determinant of whether the model behaves like a product or a project.
3. Can the operator manage the infrastructure remotely? Test provisioning, gateway and device health, OTA updates and rollback, alarms, remote diagnostics, access control, and support workflows.
4. Does the first application create measurable customer value? Energy, HVAC, equipment monitoring, environmental intelligence, occupancy, or another initial use case should establish a documented economic reason for the infrastructure to remain.
5. Can another application be added without another infrastructure project? This is the critical platform test and should be demonstrated during the pilot rather than assumed.
6. Does the security model survive the customer's own review? Put segmentation, update integrity, control authority, vulnerability handling, and third-party workload governance in front of the customer security organization during the pilot.
7. Can the telco operationalize the service? Sales, ordering, activation, billing, support, escalation, field service, renewal, and end-of-term handling matter as much as the technology.

A successful pilot should end with a repeatable deployment specification, measured per-site cost, known integration and exception rates, a security model accepted by the customer, and a documented path to add the next application. Those are the conditions that allow the model to expand from buildings to portfolios.

> **The goal is not a successful demonstration. It is a deployable infrastructure product.**

The LEXI Perspective

LEXI believes telecommunications providers are uniquely positioned to become primary distributors of a new class of intelligent building infrastructure - not because the telecommunications network itself becomes the building infrastructure asset, but because operators already possess the distribution, connectivity, lifecycle-management, service-assurance, billing, and enterprise-relationship machinery required to place and operate infrastructure at scale.

LEXI creates the building-resident asset. The telco provides the distribution machine that places, connects, manages, monetizes, and expands it across enterprise customer portfolios. Whether the hardware is operator-owned, financed, customer-funded, or property-owner-funded is secondary to that strategic role, provided the economics, term, ownership, data rights, and exit are explicit.

The case for acting now is not that building intelligence is new. It is that connectivity is becoming a thinner economic layer while the operator already reaches the buildings in which much larger operating-cost and service pools exist. LEXI gives the operator an infrastructure layer that can reach existing building systems, survive changes in radios and applications, and expand upward into managed infrastructure, building applications, and distributed intelligence without requiring the carrier to invent another vertically integrated IoT stack.

The first service may be energy, equipment monitoring, environmental intelligence, or building automation. The strategic value is that the infrastructure does not end with the first service. It remains, is upgraded, connects more systems, runs additional applications, and hosts new workloads - turning individual customer buildings into a distributed service-delivery footprint.

> **Do not sell another building project. Distribute the infrastructure on which the next generation of building services will run.**

Sources

[1] GSMA, "Global IoT Connectivity Must Prove It Can Scale Reliably," 1 May 2026; and GSMA milestone reporting on one billion NB-IoT/LTE-M low-power wide-area connections by the end of 2025. https://www.gsma.com/solutions-and-impact/technologies/internet-of-things/

[2] Berg Insight, The Global M2M/IoT Communications Market, reported 27 July 2026. Estimated cellular IoT connectivity revenue of EUR 14.5 billion in 2025, 4.2 billion connections, and average monthly revenue per connection of EUR 0.31. https://iotbusinessnews.com/2026/07/27/cellular-iot-connectivity-revenues-reached-e14-5-billion-in-2025/

[3] U.S. Energy Information Administration, 2018 Commercial Buildings Energy Consumption Survey, key energy statistics and overview. Approximately 5.9 million buildings, 96.4 billion square feet, and about $141 billion in annual energy expenditure in 2018. https://www.eia.gov/consumption/commercial/pba/overview.php

[4] Ericsson, "Ericsson announces closing of the transfer of IoT Accelerator and Connected Vehicle Cloud businesses to Aeris," 31 March 2023. https://www.ericsson.com/en/news/2023/3/ericsson-announces-closing-of-the-transfer-of-iot-accelerator-and-connected-vehicle-cloud-businesses-to-aeris

[5] Ericsson and contemporary industry reporting on the Aeris transaction, including forecast sales, quarterly-loss reduction, and Ericsson's statement that it had captured a limited part of the IoT value chain. https://www.ericsson.com/en/press-releases/2022/12/aeris-to-acquire-iot-business-from-ericsson

[6] Reporting on AT&T Digital Life wind-down and transition as AT&T retired its 3G network. https://www.securitysales.com/news/brinks-home-att-digital-life/

[7] RCR Wireless News, "AT&T quits NB-IoT," November 2024. AT&T stopped NB-IoT sales and certification and directed customers toward LTE-M as it planned network decommissioning. https://www.rcrwireless.com/20241120/internet-of-things-4/att-quits-nb-iot

[8] GSMA North America, "Exploring Private Networks Through the Lens of MNOs," updated November 2025. https://www.gsma.com/about-us/regions/north-america/gsma_resources/updated-whitepaper-exploring-private-networks-through-the-lens-of-mnos/

[9] GSMA, "Hybrid Networks," describing combinations of private and public infrastructure for enterprise environments. https://www.gsma.com/solutions-and-impact/industries/connected-manufacturing/hybrid-networks/

[10] ETSI, Multi-access Edge Computing, defining MEC as cloud-computing capabilities and an IT service environment at the edge of the network. https://www.etsi.org/technical-groups/mec/

[11] ETSI White Paper No. 59, "Enabling Multi-access Edge Computing in the Internet of Things," June 2023, describing cloud, edge and far-edge tiers and gateways interfacing with industrial equipment environments. https://www.etsi.org/images/files/ETSIWhitePapers/ETSI-WP59-Enabling-Multi-access-Edge-Computing-in-iot.pdf

[12] Singtel, Next-Gen Managed Services, describing managed services spanning data center, hybrid cloud, IoT networks, automation, analytics, and unified monitoring. https://www.singtel.com/business/products-services/managed-services

[13] Singtel, Managed IoT Intelligence and IoT connectivity offerings, describing sensors, connectivity, analytics, rule-based workflows, device lifecycle management, energy monitoring, and multi-network connectivity. https://www.singtel.com/business/products-services/iot/managed-iot-intelligence