# Network Technologies for Heterogeneous and Converged Connectivity: A Comparative Survey

Niha Naineni
Department of Computer Science
Virginia Tech
nihanaineni@outlook.com

***Abstract*—Modern connectivity is supplied by a portfolio of technologies rather than by a single dominant network. Ethernet and passive optical networks provide stable high-capacity infrastructure; Wi-Fi supplies flexible local access; cellular systems add managed mobility and wide-area coverage; low-power wide-area networks support constrained sensing; and non-terrestrial networks extend service beyond economical terrestrial reach. Meanwhile, time-sensitive networking, deterministic networking, software-defined networking, network functions virtualization, slicing, and edge computing reshape how access technologies are assured and combined. This survey develops a three-layer taxonomy that separates transport, assurance, and control or service functions, then compares major technology families across capacity, latency, reliability, range, mobility, energy, economics, security, interoperability, and maturity. A structured narrative review of standards and peer-reviewed literature is used to identify recurring trade-offs and deployment boundary conditions. The synthesis shows that performance claims are meaningful only when topology, workload, channel, ownership, and failure assumptions are explicit. No family dominates the complete design space: emerging systems are heterogeneous by construction and increasingly depend on cross-domain orchestration. The paper concludes with application mappings and a research agenda centered on verifiable end-to-end guarantees, trustworthy automation, zero-trust identity, sustainable operation, and reproducible evaluation of 5G-Advanced and prospective 6G systems.**

***Index Terms*—5G-Advanced, 6G, deterministic networking, Ethernet, heterogeneous networks, low-power wide-area networks, network functions virtualization, non-terrestrial networks, software-defined networking, Wi-Fi.**

## I. Introduction

NETWORK engineering has shifted from selecting one access method to composing multiple technologies around an application. A factory can combine industrial Ethernet for programmable logic controllers, Wi-Fi for operator terminals, private cellular service for mobile robots, low-power wireless sensors for condition monitoring, and an optical uplink to enterprise or cloud systems. A smart-city platform may combine unlicensed low-power wide-area networking (LPWAN), public cellular service, fiber aggregation, and satellite failover. In each case, the service perceived by an application emerges from several physical media, protocol families, schedulers, trust domains, and operators.

This heterogeneity creates a basic problem for comparative surveys. Technologies at different layers are often placed in the same table as if they were substitutes. Ethernet, Wi-Fi, fifth-generation (5G) cellular radio, LPWAN, and satellite systems move frames or packets over a medium. Time-sensitive networking (TSN) and deterministic networking (DetNet) provide assurance mechanisms over selected media and paths. Software-defined networking (SDN), network functions virtualization (NFV), slicing, and edge computing instead change how networks are controlled, isolated, and used to host services. Comparing all of them by peak throughput alone produces a category error. It also obscures dependencies: a network slice still requires a transport, an edge platform still requires a path, and a deterministic application still requires every traversed segment to implement compatible timing and resource controls.

The second problem is that headline specifications are conditional. A wireless physical-layer rate depends on channel width, spatial streams, coding, signal quality, contention, and client capability. Latency depends on offered load, scheduling, queueing, retransmissions, core placement, and the percentile reported. Range depends on frequency, antenna, terrain, regulation, and link budget. Device lifetime depends on sleep policy and traffic more than on a protocol label. A useful survey must therefore distinguish capability envelopes from measured application performance and must record the assumptions under which comparisons are made.

The third problem is convergence. Fiber, Ethernet, Wi-Fi, cellular, LPWAN, and non-terrestrial networks (NTNs) are not progressing in isolation. Cellular releases incorporate satellite access and positioning; Wi-Fi and cellular systems share edge platforms; Ethernet extends into operational technology through TSN; DetNet seeks deterministic paths across routed domains; and SDN or NFV platforms orchestrate services over heterogeneous infrastructure. The relevant unit of design is increasingly an end-to-end service chain rather than a single link.

This paper addresses these problems through four contributions.

1) It presents a three-layer taxonomy separating access and transport, assurance, and control or service technologies.
2) It provides a multidimensional comparison of eight major network families using explicit criteria for performance, energy, economics, security, interoperability, and maturity.
3) It maps technology combinations to representative industrial, urban, enterprise, transportation, and emergency scenarios rather than declaring a universal winner.
4) It identifies research questions that arise specifically at technology boundaries: cross-domain service-level

objectives, trustworthy automation, identity continuity, deterministic wireless, terrestrial–NTN integration, sustainability, and reproducibility.

The remainder of the paper is organized as follows. Section II explains the scope, research questions, and evidence method. Section III introduces the taxonomy and comparison framework. Section IV surveys the technology families. Section V synthesizes cross-technology trade-offs. Section VI analyzes converged architectures and use cases. Section VII develops open research problems, and Sections VIII and IX discuss limitations and conclude.

## II. Scope, Research Questions, and Review Method

### A. Research Scope

The review covers technologies that materially shape contemporary local, access, wide-area, industrial, and emerging converged networks. Seven families are included: wired Ethernet and optical access; wireless local- and personal-area systems; cellular networks; LPWANs; NTNs; deterministic networking; and programmable or virtualized architectures. The IEEE 802 architecture provides a common reference for local and metropolitan area networks, including Ethernet, wireless LANs, and low-rate wireless systems [1]. Cellular capabilities are considered through Long Term Evolution (LTE), 5G, and 5G-Advanced, while sixth-generation (6G) claims are treated as research and standardization targets. The review does not compare individual commercial products, regional tariffs, proprietary radio planning tools, or vendor-specific orchestration stacks.

The analysis is guided by five research questions.

- *RQ1:* Which properties distinguish major technology families when deployment assumptions are controlled?
- *RQ2:* Which technologies are substitutes, and which operate at different layers or provide complementary functions?
- *RQ3:* How do reach, mobility, latency, reliability, energy, cost, security, and maturity trade against one another?
- *RQ4:* Which combinations best match representative application envelopes?
- *RQ5:* Which unresolved problems occur at the boundaries between technology and administrative domains?

### B. Evidence Collection and Selection

The study follows a structured narrative method rather than a statistical meta-analysis. Normative standards and architecture documents were prioritized when the question concerned protocol scope, service definitions, or standardization status. Sources include IEEE, the Internet Engineering Task Force (IETF), the International Telecommunication Union (ITU), the Third Generation Partnership Project (3GPP), the European Telecommunications Standards Institute (ETSI), the Open Networking Foundation (ONF), and the U.S. National Institute of Standards and Technology (NIST). Peer-reviewed surveys and foundational articles were selected when they established a technology direction, synthesized a large literature, or compared alternatives under explicit assumptions.

Three inclusion rules were applied. First, a source had to define an architecture, explain mechanisms, or present a comparative synthesis relevant to one or more research questions. Second, sources had to be sufficiently stable or influential to support family-level conclusions; transient product announcements were excluded. Third, emerging 6G claims were included only when they helped identify candidate capabilities or evaluation problems, not as evidence of deployed performance. The resulting bibliography combines normative references with peer-reviewed work and retains foundational papers where current systems inherit the same concepts.

The evidence base has deliberate limits. Standards describe supported behavior but do not guarantee implementation quality, coverage, or commercial availability. Survey papers aggregate studies that may use different workloads and metrics. Furthermore, mature families have decades of operational evidence, whereas 6G and some satellite-integrated architectures remain dominated by analytical studies, prototypes, and requirements work. Conclusions are therefore expressed as conditional comparative judgments rather than universal numerical rankings.

### C. Comparison Discipline

Each family is assessed against twelve criteria listed in Table I. The criteria separate properties that are frequently collapsed into a single notion of performance. Capacity concerns sustained useful throughput under a stated load, while delay includes both central tendency and tail behavior. Reliability refers to delivery probability and recovery behavior; availability additionally depends on redundancy, maintenance, and correlated failure. Reach is not equivalent to coverage probability, and mobility is not equivalent to portability. Similar distinctions apply to endpoint energy, infrastructure energy, capital expenditure, operating expenditure, security, interoperability, and maturity.

Quantitative maxima are not used as universal scores. Instead, qualitative envelopes are reported and the dominant trade-offs are identified. For a concrete deployment, the comparison can be formalized as a constrained multi-criteria decision problem. Let $s_{i,t} \in [0,1]$ denote the normalized suitability of technology $t$ on criterion $i$, $w_i$ an application-specific importance weight, and $p_{j,t}$ a penalty for violating constraint $j$. A transparent ranking can be written as

$$U(t) = \sum_{i=1}^{n} w_i s_{i,t} - \sum_{j=1}^{m} \lambda_j p_{j,t}, \quad (1)$$

where $\lambda_j$ is large for a hard constraint such as a safety deadline, coverage minimum, or energy budget. Equation (1) is not presented as an empirical universal model; it illustrates why rankings change when application weights or constraints change. A low-power sensor and a mobile robot should not produce the same ordering.

## III. A Three-Layer Taxonomy

### A. Access and Transport

The first layer moves frames or packets across physical media and topologies. It includes copper and optical Ethernet,

TABLE I
COMPARISON CRITERIA AND REQUIRED INTERPRETATION

| Criterion | Question Asked | Boundary Conditions That Must Be Reported |
|---|---|---|
| Capacity | What useful traffic volume can the system sustain? | Channel or link rate, topology, overhead, concurrency, traffic mix, and backhaul |
| Latency and jitter | How quickly and predictably is a packet or transaction completed? | Load, percentile, queue discipline, retransmissions, path length, and processing placement |
| Reliability | How likely is correct delivery within a requirement? | Packet size, deadline, channel state, diversity, redundancy, and failure model |
| Reach | What geographic or topological span is economical? | Frequency, power, antenna, propagation, regulation, density, and civil infrastructure |
| Mobility | How well are moving endpoints and sessions supported? | Speed, handover domain, interruption tolerance, authentication, and route continuity |
| Energy | What energy is consumed by endpoints and infrastructure? | Sleep duty cycle, reporting interval, computation, cooling, and embodied equipment |
| Economics | What resources are required to deploy and operate? | Spectrum, installation, subscription, staff, maintenance, renewal, and service lifetime |
| Security and privacy | What trust, identity, and isolation mechanisms exist? | Device capability, administrative domains, key lifecycle, update path, and threat model |
| Interoperability | Can components and management systems be substituted or combined? | Standards profile, certification, APIs, data models, spectrum rules, and vendor extensions |
| Maturity | How much deployment, tooling, and operational evidence exists? | Release status, ecosystem depth, test coverage, skill availability, and backward compatibility |

passive optical networking (PON), Wi-Fi, IEEE 802.15.4-based systems, Bluetooth Low Energy, ultra-wideband, cellular radio and core transport, LPWAN, and satellite or airborne links. Technologies within this layer can sometimes substitute for one another at an interface, but they occupy very different operating points. Ethernet and fiber concentrate capacity and predictability in installed infrastructure. Wi-Fi trades deterministic behavior for low-cost local flexibility. Cellular combines mobility, licensed-spectrum coordination, authentication, and managed service. LPWAN sacrifices throughput to extend reach and endpoint lifetime. NTN sacrifices some delay, terminal simplicity, or capacity efficiency to reach locations where terrestrial deployment is uneconomical.

### B. Assurance

The second layer adds timing, loss, availability, or isolation properties to a selected transport. TSN is a family of IEEE 802.1 mechanisms for time synchronization, traffic shaping, scheduled transmission, frame preemption, resource reservation, and redundancy. DetNet extends deterministic service concepts across Layer-3 routed segments and explicitly targets paths that can traverse both routed and bridged subnetworks [2]. Assurance is not intrinsic merely because a technology has a low unloaded latency. Determinism requires admission control, queue isolation, path computation, synchronization where needed, and compatible behavior at every hop. Research on ultra-low-latency networks shows that TSN, DetNet, and cellular low-latency mechanisms must be coordinated rather than evaluated independently [3].

### C. Control and Service

The third layer programs, virtualizes, isolates, and places functions over one or more transports. SDN separates control logic from packet-forwarding devices and exposes programmable abstractions [4]. NFV moves functions from dedicated appliances to virtual machines, containers, or cloud-native execution environments [5]. Network slicing creates logical service contexts over shared resources, while edge or fog computing places application and network processing closer to users or sensors. These mechanisms can accelerate service creation and improve resource use, but they also introduce dependencies on controllers, orchestrators, software supply chains, telemetry, and cross-domain interfaces.

The layered model connects applications and service-level objectives to control, assurance, and transport mechanisms. It also exposes a common failure mode in convergence: an orchestrator may reserve resources inside one domain but lack a compatible abstraction for the next. A service can consequently appear isolated in a dashboard while traversing a shared link whose queue, spectrum, or failure state is not visible.

## IV. TECHNOLOGY FAMILIES

### A. Wired Ethernet and Optical Access

IEEE 802.3 Ethernet remains the default wired local-area technology because it combines mature interoperability, predictable point-to-point links, strong backward compatibility, and a broad rate ecosystem [6]. Modern switched full-duplex Ethernet removes the collision behavior of early shared media. Link aggregation, redundant topologies, virtual LANs, quality-of-service mechanisms, and mature operational tooling make it suitable for enterprise, data-center, campus, and industrial infrastructure. TSN profiles extend Ethernet toward applications that require bounded delay, accurate time distribution, or seamless redundancy.

At data-center and HPC scale, topology becomes a first-class network design variable. Low-diameter designs such as Slim Fly and Dragonfly demonstrate trade-offs among path length, bandwidth, resilience, switch cost, cabling, and power [7], [8]. Special-purpose cluster interconnects provide another operating point. For example, the Tofu Interconnect D uses a multidimensional mesh/torus architecture and hardware-supported communication mechanisms designed for large-scale systems [9].

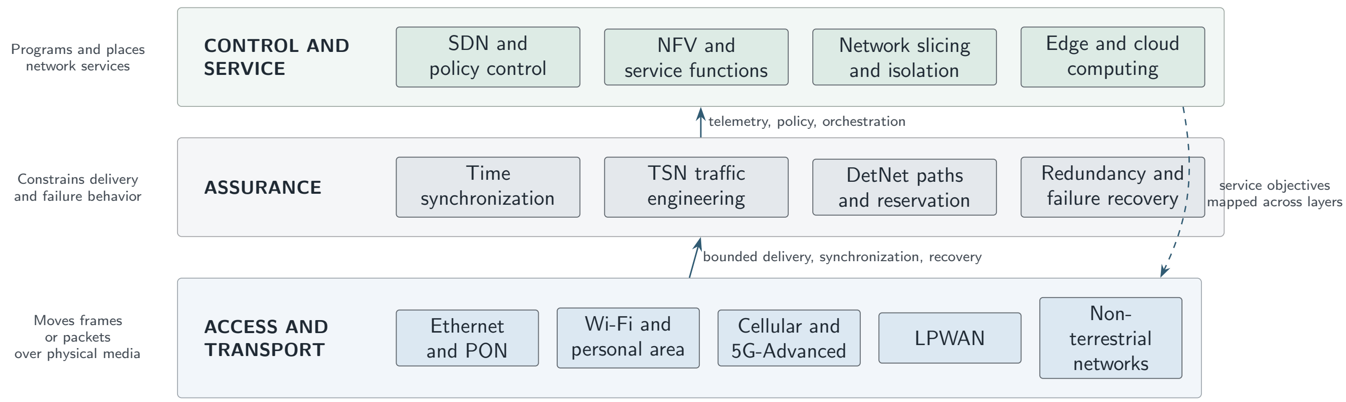


Fig. 1. Three-layer taxonomy separating access and transport technologies from assurance mechanisms and programmable control and service functions.

Optical fiber complements Ethernet in access, aggregation, data centers, and cores. It provides high capacity, long reach, low attenuation, and immunity to electromagnetic interference. PON shares an optical distribution network among endpoints, reducing active equipment in the field. Current ITU work specifies higher-speed PON systems above 10 Gbit/s per wavelength channel [10]. The shared architecture improves deployment economics but introduces scheduling, split-ratio, and failure-domain considerations. A nominal access rate is shared among subscribers, and the optical distribution plant can become a common point of physical damage.

For industrial communication, the deeper trend is convergence between information technology and operational technology. Traditional fieldbuses and specialized industrial Ethernet profiles coexist with commodity Ethernet, Internet Protocol (IP), and TSN. Wollschlaeger *et al.* describe this transition as part of the broader integration associated with the Internet of Things (IoT) and Industry 4.0 [11]. The engineering challenge is not simply increasing line rate; it is preserving safety, timing, maintenance, and long equipment lifetimes while introducing shared infrastructure and software-defined control.

### *B. Wi-Fi and Personal-Area Networks*

Wi-Fi is the dominant low-cost local wireless access family. It operates primarily in unlicensed spectrum, benefits from a large device ecosystem, and can be deployed without mobile-operator infrastructure. IEEE 802.11ax introduced mechanisms aimed at dense, high-efficiency wireless LANs, including orthogonal frequency-division multiple access, improved multi-user operation, and scheduling features [12], [13]. IEEE 802.11be-2024, associated with Wi-Fi 7, adds extremely-high-throughput mechanisms at the physical and medium-access layers [14]. Multi-link operation, wider channels, and additional spatial capabilities can increase flexibility and peak performance.

Application performance nevertheless depends on contention, interference, access-point placement, client implementation, power-save behavior, and backhaul. Unlicensed operation lowers entry barriers but does not assign an operator exclusive control over neighboring transmitters. Retransmissions and contention can produce long-tail delay even when the average is low. Managed enterprise deployments mitigate these effects through channel planning, access-point density, admission policy, roaming coordination, and spectrum analysis. These measures improve service but also reduce the apparent simplicity and low cost that make Wi-Fi attractive.

Personal-area technologies occupy lower-power and shorter-range operating points. IEEE 802.15.4 supports low-rate wireless networks and underpins ecosystems such as Zigbee and Thread [15]. Bluetooth Low Energy is optimized for peripherals, wearables, beacons, and intermittent exchange. Ultra-wideband is used for accurate ranging and localized communication. Mesh routing can extend coverage but consumes airtime and energy, increases path variability, and complicates troubleshooting. These technologies are effective when short packets, localized communication, low device cost, or battery operation matter more than sustained capacity.

### *C. Cellular Networks and 5G-Advanced*

Cellular systems combine wide-area radio planning, licensed-spectrum coordination, managed mobility, authentication, charging, and an operator-grade core. LTE remains important because of its coverage, device ecosystem, and embedded deployment base. Early 5G work framed densification, millimeter-wave operation, massive multiple-input multiple-output, device-centric design, and virtualization as mutually reinforcing directions [16], [17]. The deployed 5G system consequently represents both radio evolution and a shift toward a service-based, software-driven core.

Three broad service categories shaped 5G discussion: enhanced mobile broadband, massive machine-type communication, and ultra-reliable low-latency communication. These labels represent design directions rather than guarantees available uniformly across bands, releases, operators, and devices. Low-band deployments improve coverage but cannot offer the same capacity as dense mid-band or millimeter-wave systems. Millimeter-wave links can provide large bandwidth but require dense infrastructure, beam management, and blockage mitigation. Low latency depends on radio scheduling, transport, core placement, application processing, and load.

Network slicing seeks to create logical networks with distinct policies and resource treatment over shared infrastructure. Foukas *et al.* identify resource abstraction, isolation, admission, orchestration, and multi-domain management as core challenges [18]. A slice identifier does not by itself guarantee isolation: compute, storage, transport queues, spectrum, and control-plane dependencies must all be considered. The problem becomes more difficult when a slice traverses infrastructure owned by several organizations or uses non-cellular access.

3GPP Release 18 marks the first 5G-Advanced release and extends radio performance, positioning, sidelink, automation, energy-related functions, and NTN integration [19]. These extensions strengthen cellular support for industrial, transportation, and integrated access scenarios, but feature presence is not synonymous with end-to-end availability. Device categories, spectrum authorizations, operator implementation, core capabilities, and local design all influence what an application can use.

### D. Prospective 6G and IMT-2030

6G should be treated as a research and standardization trajectory, not as a mature deployment class. Research visions emphasize integrated sensing and communication, sub-terahertz operation, native or pervasive artificial intelligence, new topology models, extreme localization, and tighter terrestrial–non-terrestrial integration [20], [21]. These themes indicate where research investment is concentrated, but they do not establish commercial performance.

ITU-R Recommendation M.2160 defines the IMT-2030 framework, usage scenarios, and capability directions [22]. In 2026, ITU experts agreed draft technical performance requirements spanning 20 dimensions, with formal approval expected later in the process [23]. The requirements phase is important because it determines how candidate radio interface technologies will later be evaluated. It also exposes a recurring measurement problem: aspirations such as sustainability, intelligence, sensing, resilience, and ubiquitous coverage require definitions that extend beyond a peak data-rate test.

### E. Low-Power Wide-Area Networks

LPWANs serve devices that transmit small amounts of data over long distances while operating for extended periods with limited energy. IETF RFC 8376 characterizes the family by broad coverage, low bandwidth, small packets, and long battery life [24]. Common use cases include metering, environmental monitoring, asset status, agriculture, infrastructure alarms, and low-frequency telemetry. These applications tolerate delays or low duty cycles that would be unacceptable for control loops or interactive media.

The family includes both unlicensed and licensed-spectrum systems. LoRaWAN and Sigfox emphasize low-cost endpoints and infrastructure in unlicensed bands. NB-IoT and LTE-M reuse cellular spectrum, sites, identity systems, and operator cores. Comparative work shows that the choice depends on interference, link budget, capacity, deployment ownership, mobility, and quality-of-service needs rather than on range alone [25]–[28]. LoRa or Sigfox can provide attractive cost and battery properties, whereas NB-IoT can provide operator integration and managed service under suitable coverage. LTE-M supports higher data rates and mobility than the most constrained LPWAN profiles at a corresponding energy and complexity cost.

Security must account for devices that may be physically exposed, computationally constrained, and expected to operate for years. Key provisioning, secure boot, replay protection, update mechanisms, ownership transfer, and decommissioning are as important as link encryption. Comparative studies that consider only initial radio energy overlook the lifecycle cost of failed updates, truck rolls, battery replacement, or compromise.

### F. Non-Terrestrial Networks

NTNs use satellites or airborne platforms to extend coverage over oceans, remote regions, aviation routes, disaster zones, and areas where terrestrial infrastructure is uneconomical. Geostationary Earth orbit systems provide large, relatively stationary footprints but incur substantial propagation delay. Low Earth orbit constellations reduce propagation distance and can lower delay, but require moving-beam management, handover, gateway diversity, tracking, and large-scale constellation operation. Medium Earth orbit systems occupy an intermediate trade space.

3GPP Release 17 established initial support for New Radio and IoT operation over NTN, and Release 18 introduced further enhancements [29]. Standardized integration can reduce the need for entirely separate satellite terminals and service stacks, particularly for low-rate IoT and messaging. However, on-board processing, inter-satellite routing, Doppler, timing, spectrum coordination, gateway placement, and validation remain active challenges [30].

NTN is best understood as complementary infrastructure. It adds reach, geographic diversity, and rapid service restoration, but generally cannot match the dense capacity economics of fiber or terrestrial radio in populated areas. Terminals can face higher energy and antenna constraints, and shared beams can become capacity bottlenecks. For resilience, independence must be examined carefully: a satellite link may bypass a damaged local cable while still sharing power, a gateway, cloud control, spectrum coordination, or an identity provider with the failed service.

### G. Deterministic Networking

Deterministic networking addresses traffic whose correctness depends on delivery within a bounded time and loss envelope. Examples include industrial motion control, vehicle systems, utility protection, professional media, and selected avionics or transportation functions. A low average delay is insufficient because rare queueing or recovery events can violate a deadline. The design objective is consequently expressed in worst-case or high-confidence terms.

TSN introduces mechanisms such as precise time synchronization, scheduled traffic, credit-based shaping, frame preemption, per-stream filtering, and redundancy. Different profiles select subsets suitable for industrial automation, automotive, or audio-video applications. DetNet extends the concept across routed networks using explicit paths, resource allocation, service protection, and bounded congestion loss [2]. The two are complementary when a deterministic IP flow traverses TSN-enabled local segments.

Wireless determinism remains especially difficult. Fading, interference, mobility, retransmission, and shared spectrum produce uncertainty not present on a controlled full-duplex wire. Scheduled cellular and industrial wireless systems can improve predictability, but end-to-end guarantees still include transport and processing. The survey of ultra-low-latency networking by Nasrallah *et al.* demonstrates why TSN, DetNet, and 5G mechanisms must be analyzed as a coordinated system [3].

### H. SDN, NFV, Slicing, and Edge Computing

SDN separates network control from forwarding and exposes abstractions through which software can program policy and paths [4]. The global view can improve traffic engineering, access control, and service automation. Akyildiz *et al.* show, however, that monitoring scale, rule placement, controller responsiveness, and recovery remain coupled traffic-engineering problems [31]. The broad SDN literature identifies scalability, resilience, verification, security, and migration as continuing concerns [32].

NFV implements network functions on virtualized or cloud-native infrastructure instead of requiring a dedicated appliance for each function. ETSI work covers virtualized functions, container infrastructure, management, and orchestration [5]. Surveys identify performance isolation, state management, function placement, service chaining, acceleration, and carrier-grade reliability as core challenges [33], [34]. Cloud-native implementation improves deployment flexibility but increases the importance of software supply-chain security, image provenance, dependency management, and observable state.

Slicing combines programmable control, virtualization, policy, and resource allocation to present logical networks with different service characteristics. Slices can help separate enterprise tenants, public-safety services, or traffic classes, but isolation must be demonstrated across radio, transport, compute, storage, and management. A shared controller, hypervisor, acceleration device, timing source, or physical path can create correlated failures even when logical addresses are distinct.

Edge and fog computing move processing toward users, machines, or data sources. Mobile edge computing research emphasizes the joint optimization of radio, computation, and task placement [35]; fog computing extends the concept across distributed resources between devices and clouds [36]. Local processing can reduce latency, backhaul use, and data exposure. It can also reduce economies of scale, create many small failure domains, complicate updates, and move data across jurisdictional boundaries. Function placement must therefore consider compute capacity, network state, energy, trust, resilience, and service migration together.

RDMA-accelerated serverless systems such as rFaaS also demonstrate that function-placement performance depends jointly on compute resources, invocation mechanisms, and the underlying network interface [37].

## V. Cross-Technology Comparative Synthesis

### A. Qualitative Operating Envelopes

Table II compares representative operating envelopes. The symbols H, M, and L denote relative high, medium, and low positions within the surveyed families, not guaranteed measurements. A range such as L–H indicates that deployment choices dominate the family label. For example, cellular capacity varies markedly across spectrum and density, while NTN delay varies by orbit and routing.

Three broad patterns follow. First, reach and concentrated capacity are usually optimized by different infrastructure choices. Fiber and dense radio systems provide capacity where investment can be concentrated; LPWAN and NTN extend service over sparse areas or to constrained endpoints. Second, mobility and predictability require active coordination. Cellular networks manage handover and scheduling, while deterministic systems engineer paths and queues. Contention-based local wireless offers flexibility but not automatic bounds. Third, the programmability families cannot be assigned link metrics because they operate across links. Their value must be evaluated through service deployment time, policy correctness, utilization, recovery, and lifecycle cost.

### B. Latency, Reliability, and Failure Semantics

Latency comparisons should report a distribution and a path definition. A radio-interface measurement excludes core, transport, security gateway, and application processing. A round-trip median hides the tail relevant to control. A deterministic claim may apply only after admission and under a specified fault model. At minimum, studies should report median, a high percentile, maximum observation, loss, load, and recovery behavior.

Experiments with RDMA data-center networks further show that congestion-control decisions can improve short-flow behavior while penalizing long flows, making fairness and high-percentile flow completion time essential evaluation metrics [38], [39].

Reliability also has several meanings. Packet delivery probability measures an individual transfer. Availability measures the fraction of time a service is usable. Resilience measures the ability to continue or recover under faults. Safety may require a system to fail predictably rather than merely remain available. Ethernet redundancy, cellular dual connectivity, Wi-Fi multi-link operation, multipath transport, and satellite backup all provide diversity, but the benefit depends on failure independence. Two paths routed through the same duct, gateway, power supply, controller, or cloud region are not fully redundant.

Cross-domain assurance is limited by the least observable segment. A deterministic TSN flow can lose its bound when mapped into a best-effort IP tunnel. A cellular slice can traverse a shared PON or cloud network without corresponding isolation. A satellite backup can preserve connectivity while increasing delay beyond an application's timeout. End-to-end objectives therefore require explicit translation between local assurance models.

TABLE II
QUALITATIVE COMPARISON OF MAJOR TECHNOLOGY FAMILIES

| Family | Reach / Mobility | Capacity / Delay | Endpoint Energy | Dominant Strength | Principal Limitation |
|---|---|---|---|---|---|
| Ethernet / PON | Site to core; fixed | High; low and stable | Moderate | Predictable capacity and mature tooling | Installation cost and no physical mobility |
| Wi-Fi | Local; portable | Medium–high; load-sensitive | Moderate | Economical local flexibility | Contention, interference, and tail latency |
| Personal-area | Short; limited mobility | Low; application-dependent | Very low | Low-cost localized devices | Limited rate, reach, and integration |
| Cellular | Wide area; strong mobility | Medium–high; managed | Moderate–high | Coverage, identity, and handover | Spectrum, infrastructure, and recurring cost |
| LPWAN | Wide area; limited mobility | Low; delay-tolerant | Very low | Long-lived constrained sensing | Small payloads and limited real-time support |
| NTN | Regional–global; mobile | Variable; orbit-dependent | Moderate–high | Geographic reach and path diversity | Delay, terminal constraints, shared capacity |
| TSN / DetNet | Engineered domain; fixed | High; bounded when admitted | Moderate | Bounded delay and low loss | Configuration and cross-domain complexity |
| SDN / NFV / edge | Spans underlying transports | Determined by hosted path and compute | N/A | Programmability and service placement | Software concentration and orchestration risk |

### C. Energy and Sustainability

Endpoint energy depends strongly on duty cycle, listening, retransmission, and application computation. LPWAN excels when devices sleep for long intervals and exchange small messages. Personal-area systems are effective for short-range intermittent communication. Wi-Fi and cellular devices can use power-saving features but consume more energy when maintaining high throughput, continuous reachability, or poor-coverage links. Satellite terminals face demanding link budgets, although specialized direct-to-device and IoT modes reduce the requirement for selected services.

Infrastructure comparisons are more complex. Fiber is energy efficient per transported bit at high utilization, but installation has embodied impacts. Dense cellular and Wi-Fi deployments increase site, backhaul, and cooling requirements. Edge computing may reduce transport while duplicating underutilized compute across sites. Satellite systems require launch, replenishment, gateway, and space-segment resources. Reporting only joules per bit can favor a high-capacity system even when the workload does not need that capacity. Sustainability should be evaluated per delivered service over the infrastructure lifecycle.

### D. Economics, Ownership, and Maturity

Unlicensed systems reduce spectrum and operator barriers but shift planning, interference diagnosis, security, and maintenance to the site owner. Public cellular and satellite services transfer much operation to providers, but create recurring charges, coverage dependencies, and external policy constraints. Private cellular offers control at the cost of spectrum arrangements, radio planning, core operation, and specialist skills. Ethernet and PON require physical installation but provide long-lived capacity and clear ownership.

Capital expenditure and operating expenditure should be separated. A low-cost sensor can create high operating cost if batteries, gateways, or proprietary subscriptions require frequent intervention. A more expensive wired installation can be economical over a long industrial lifecycle. Migration cost, vendor portability, certification, failure recovery, workforce skills, and expected service life are therefore first-class comparison criteria.

Maturity affects risk independently of nominal performance. Ethernet, Wi-Fi, LTE, and established PON profiles have broad ecosystems and operational knowledge. New 5G-Advanced features, integrated NTN functions, and 6G concepts have less field evidence. A mature solution may be preferable when spare parts, diagnostic tools, certification, and long-term support outweigh a newer feature.

### E. Security and Trust Boundaries

Heterogeneous networks combine constrained sensors, roaming users, virtual functions, third-party clouds, remote gateways, and several operators. Security cannot be reduced to link encryption. Device identity, credential provisioning, secure boot, update mechanisms, management access, logging, segmentation, privacy, and decommissioning must persist across the service lifecycle.

IRTF RFC 8576 notes that constrained computation, small packets, sleep behavior, and long device lifetimes complicate established Internet security mechanisms [40]. NIST zero-trust architecture removes implicit trust based on network location and instead emphasizes explicit users, assets, resources, and continuous policy decisions [41]. Applying zero trust to a converged network requires identity continuity across Wi-Fi, cellular, LPWAN, satellite, edge, and cloud domains. It also requires policies that degrade safely when connectivity to an identity or policy service is lost.

Programmability changes the threat model. SDN controllers and NFV orchestrators can improve consistency but become high-value targets. A software supply-chain compromise can affect many virtual functions. Telemetry can reveal sensitive operational behavior, while inadequate telemetry hides policy violations. Isolation should therefore be tested against shared compute, accelerators, storage, timing, management, and human administration rather than inferred from logical configuration.

## VI. Converged Architectures and Application Mapping

### A. Industrial Automation

Industrial systems contain traffic classes with very different requirements. Motion control may need bounded delay, synchronization, and high availability. Machine vision requires high throughput and local compute. Mobile robots need handover and controlled radio coverage. Handheld terminals favor Wi-Fi economics, while environmental sensors favor low-power operation. A converged architecture can use TSN-enabled Ethernet for fixed control, private 5G for mobile machines, Wi-Fi for people and general devices, IEEE 802.15.4 or LPWAN for sparse sensing, and edge compute for inspection or control analytics.

The primary risk is not technology diversity itself but uncontrolled interaction. Shared backhaul can create correlated congestion. Inconsistent identity models can leave operational equipment exposed. An orchestration change intended for enterprise traffic can invalidate a deterministic schedule. Industrial deployments therefore require explicit zones, safety boundaries, change control, time distribution, and independent fallback behavior. Wireless links should be assigned only to traffic whose deadline and failure behavior can be demonstrated under interference and mobility.

### B. Smart Cities, Utilities, and Agriculture

City and utility systems span buildings, streets, substations, waterways, and rural areas. LPWAN is well matched to small, infrequent readings from meters or environmental sensors. Cellular service supports richer telemetry, mobile workforces, and managed wide-area connectivity. Fiber or microwave systems aggregate high-capacity sites. Satellite links provide remote reach or backup where terrestrial service is absent. Edge processing can filter camera or sensor data before it reaches central platforms.

The main constraints are lifecycle, ownership, and public exposure. Devices may remain deployed for a decade, change operators, or be physically accessible. A city can inherit incompatible platforms from different departments and procurement cycles. Coverage measured during a pilot can change with construction, foliage, interference, or provider refarming. Architecture should therefore favor open interfaces, replaceable gateways, explicit data ownership, credential rotation, and a documented offline mode.

### C. Enterprise and Campus Networks

Enterprise campuses typically use Ethernet and optical aggregation as the capacity substrate, Wi-Fi for user access, and public cellular for wide-area continuity. Private cellular becomes attractive for outdoor areas, industrial mobility, or devices that benefit from subscriber identity and managed handover. SDN-based campus fabrics can automate segmentation and policy, while edge platforms host latency-sensitive applications.

The design question is often operational rather than physical. Wi-Fi has a large client ecosystem and site-owned economics; private cellular adds a second radio, core, spectrum, and device-certification domain. A dual-access strategy can improve resilience only if authentication, addressing, policy, and application sessions survive handover. Observability should correlate user identity, radio state, switch path, virtual function, and application experience without collecting unnecessary personal data.

### D. Transportation, Maritime, and Aviation

Vehicles, trains, ships, aircraft, and logistics assets move across administrative and coverage boundaries. Cellular systems provide terrestrial mobility and capacity near population centers. Wi-Fi serves local cabin, depot, or station access. Satellite systems provide continuity over oceans and sparse regions. TSN or specialized deterministic buses support in-vehicle control, while wide-area links carry non-safety-critical operations, entertainment, updates, and telemetry.

Applications must distinguish onboard control from remote connectivity. A vehicle should not make a safety function depend on a wide-area path with variable coverage. Remote services need disruption-tolerant behavior, cached policy, and secure update mechanisms. Multipath operation can combine cellular and satellite links, but scheduling must consider cost, delay, loss, and data sensitivity. Mobility management across terrestrial and NTN systems remains a significant integration challenge.

### E. Public Safety and Disaster Recovery

Disasters can damage power, fiber, cellular sites, and terrestrial backhaul simultaneously. Rapidly deployable Wi-Fi, private cellular cells, microwave links, airborne platforms, and satellite terminals can restore selected services. The objective is graceful service prioritization, not immediate recreation of normal capacity. Voice, location, command, and sensor traffic may take precedence over bulk data.

Resilience planning should model common dependencies. A satellite terminal still needs local power, clear placement, credentials, and a working gateway path. A portable cellular system requires spectrum authority and device compatibility. Wi-Fi is easy to deploy but may face interference and limited coverage. Pre-provisioned identities, offline authorization, interoperable equipment, exercises, and documented fallback procedures are often more important than peak link rate.

Table III emphasizes composition rather than replacement. In every scenario, the preferred design assigns technologies to traffic classes based on requirements and then engineers the boundaries. A procurement process that asks one vendor or technology to satisfy all rows is likely to overpay for simple traffic and under-protect critical traffic.

## VII. Open Research Problems

### A. Cross-Domain Intent and Verifiable Guarantees

Network control increasingly starts with an intent such as isolate this tenant, keep delay below a bound, or survive one link failure. Each domain translates that intent into local mechanisms: a radio scheduler, TSN gate schedule, DetNet path, optical allocation, virtual-network policy, or cloud

TABLE III
REPRESENTATIVE TECHNOLOGY COMPOSITIONS BY APPLICATION

| Scenario | Likely Composition | Critical Requirements | Primary Integration Risks |
|---|---|---|---|
| Industrial automation | TSN Ethernet, private 5G, Wi-Fi, low-power sensing, edge compute | Bounded delay, safety separation, mobility, local autonomy, long lifecycle | Timing translation, shared backhaul, unsafe change, identity mismatch |
| Smart city / utility | LPWAN, public cellular, fiber aggregation, edge/cloud, satellite backup | Broad coverage, low endpoint energy, maintainability, public exposure | Device lifecycle, fragmented ownership, coverage drift, proprietary platforms |
| Enterprise campus | Ethernet/PON, Wi-Fi, public or private cellular, SDN policy | Capacity, user mobility, segmentation, manageable cost | Policy inconsistency, duplicated operations, shared identity dependencies |
| Transportation | In-vehicle deterministic networks, cellular, Wi-Fi, NTN | Mobility, intermittent reach, local safety, secure updates | Handover, variable delay, cost-aware multipath, cross-border regulation |
| Public safety / recovery | Portable Wi-Fi/cellular, microwave, NTN, local edge | Rapid deployment, priority, resilience, offline operation | Power, credentials, spectrum, device compatibility, common failure |

placement. The translations are rarely equivalent. Research is needed on machine-readable service semantics, compositional verification, and evidence that a deployed path satisfies the original objective.

An end-to-end guarantee should include assumptions, scope, and failure model. It should identify which segments are admitted, which resources are reserved, how clock error is bounded, how failover changes delay, and what happens when telemetry is incomplete. Formal verification can check configurations, but models must be kept synchronized with dynamic infrastructure. Runtime assurance can detect violations, but detection after a safety deadline may be insufficient. Hybrid design-time and runtime methods are therefore required.

### B. Deterministic Wireless and Terrestrial–NTN Integration

Deterministic wireless must manage interference, fading, blockage, mobility, and retransmission without making wasteful worst-case reservations. Promising directions include coordinated scheduling, multi-connectivity, predictive resource allocation, frequency or path diversity, and application-aware degradation. Evaluation should report tail latency and loss under interference, handover, overload, and failure, not only under controlled average conditions.

Terrestrial–NTN integration adds dynamic topology, long or variable delay, Doppler, moving beams, and gateway dependencies. Routing and congestion control must distinguish propagation from queueing and must react without instability. Mobility systems must coordinate terrestrial cells, satellite beams, and possibly multiple operators. Security and lawful-access policies cross jurisdictions. Testbeds that combine real terrestrial networks, satellite emulation, and reproducible mobility traces are needed to validate analytical results.

### C. Trustworthy AI-Assisted Control

AI can assist channel prediction, anomaly detection, traffic forecasting, energy management, placement, and fault diagnosis. The risk is that a model optimized for average throughput violates a tail-latency, safety, or fairness requirement. Distribution shift is unavoidable as topology, spectrum, software, and user behavior change. Adversaries can manipulate training data, telemetry, or control inputs.

Trustworthy network automation needs constrained decision spaces, calibrated uncertainty, interpretable evidence, rollback, and safe fallback. Models should be evaluated on rare events and out-of-distribution conditions. Training and inference data require provenance and privacy protection. Human operators need explanations tied to network state rather than generic feature importance. Formal policy can bound what a learned component may change, creating a separation between optimization and safety.

### D. Identity, Attestation, and Policy Continuity

Zero-trust principles are difficult to apply when devices sleep, roam, operate offline, or cannot support conventional agents. Research is needed on lightweight attestation, post-quantum migration for constrained devices, privacy-preserving identity, and authorization that remains safe during disconnection. Ownership transfer and decommissioning are especially underdeveloped for long-lived IoT systems.

Policy continuity across domains also remains open. A device can move from private Wi-Fi to public cellular and then to satellite while the application session persists. Each network can use different identifiers, privacy rules, and telemetry. Mapping identities must avoid creating a global tracking mechanism. Cross-domain authorization should disclose the minimum information required for service and should produce auditable evidence of policy decisions.

### E. Resilient Programmable Infrastructure

SDN, NFV, slicing, and edge computing concentrate control in software platforms. Logical redundancy does not guarantee failure independence if controller replicas share a cloud, database, software image, or credential authority. Research should model correlated software and supply-chain failures alongside physical failures. Orchestrators need transactional changes, dependency-aware rollback, and verified recovery procedures.

Observability must follow a service across physical links, virtual functions, containers, radio bearers, and administrative domains. Current telemetry often uses incompatible naming, timing, and data models. Normalization can lose semantics, while centralized collection creates cost and privacy concerns. Selective, trustworthy telemetry and causal diagnosis are needed to explain whether an application failure originated in radio, transport, compute, policy, or external dependency.

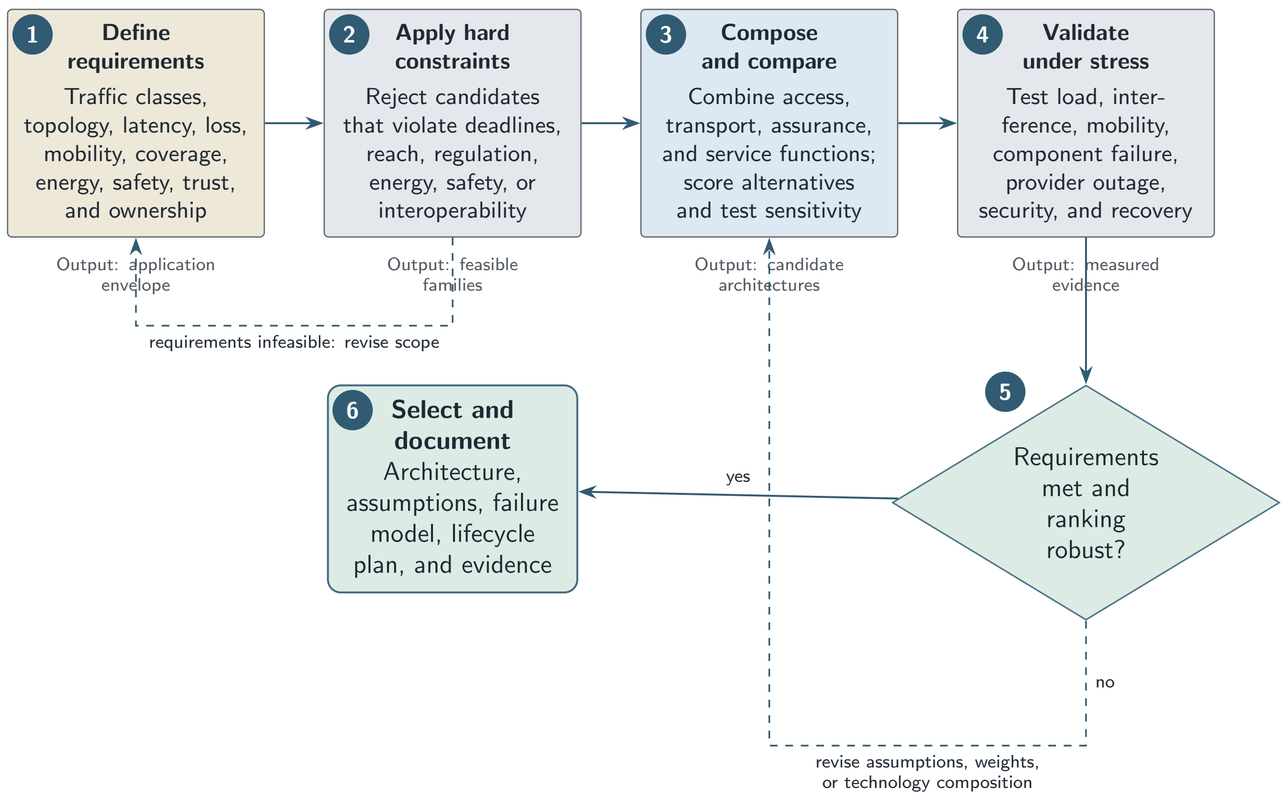


Fig. 2. Practical workflow for selecting and validating a network technology portfolio under application-specific constraints. Feedback paths represent revision after infeasibility or unsuccessful robustness testing.

### *F. Sustainable Networking*

Energy optimization frequently shifts consumption rather than reducing it. Moving computation to the edge can save transport energy but duplicate servers. Increasing radio density can reduce handset transmit power while increasing site power and embodied equipment. Larger satellite constellations can extend reach while increasing launch and replenishment impacts. A credible comparison must define the service delivered, geographic boundary, time horizon, load, and lifecycle stages.

Future work should report absolute energy, carbon intensity, utilization, and embodied impact in addition to energy per bit. Optimization should consider equipment lifetime, repairability, sleep states, cooling, workload placement, and renewable availability. Sustainability also has a resilience dimension: efficient systems that lack spare capacity can fail under emergencies, while excessive redundancy wastes resources. Multi-objective methods should expose this trade-off.

### *G. Reproducible and Comparable Evaluation*

Network results are difficult to compare when papers omit topology, channel conditions, load, traffic distributions, firmware, scheduler configuration, background interference, or confidence intervals. Emerging 6G work intensifies this problem because targets, simulations, prototypes, and standardized systems are frequently discussed together.

GPCNeT illustrates a reproducible methodology that deliberately induces network contention and measures its effects across production HPC interconnects [42].

At minimum, evaluations should publish topology, spectrum, power, antenna assumptions, workload traces, queueing and retry settings, software versions, statistical methods, and raw or derived data when possible. Tail metrics require enough observations to support the reported percentile. Comparisons should use a common application envelope rather than each technology's most favorable benchmark. Open testbeds, reference workloads, and artifact review would materially improve cumulative knowledge. Reconfigurable physical testbeds can complement simulation: the SDT platform demonstrates how commodity switches and software-defined topology control can support repeatable experiments across congestion-control, routing, and traffic-engineering configurations [43].

## VIII. Limitations and Practical Selection Procedure

This survey compares families rather than products, frequency plans, or regional deployments. Its qualitative envelopes cannot replace site measurements, and standards do not guarantee implementation quality or feature availability. Evidence is also asymmetric because mature systems have more operational history than emerging technologies.

For practical selection, an organization should define traffic, deadline, loss, mobility, coverage, energy, safety, trust, and ownership requirements. Hard constraints should eliminate infeasible candidates before weighted scoring. Survivors should be tested under representative interference, load, mobility, and failure. Lifecycle analysis should include spectrum, installation, subscriptions, monitoring, credentials, updates, batteries, spares, and staff. The design should also document how identity, assurance, telemetry, and recovery cross technology boundaries.

Equation (1) can structure the decision, but sensitivity analysis is essential. If a small change in weights reverses the ranking, the choice is not robust. Scenario analysis should include normal operation, peak load, degraded coverage, component failure, provider outage, credential-service loss, and migration. A mixed architecture often dominates a single-family design because it assigns critical, mobile, low-power, and bulk traffic to different resources. The benefit appears only when the integration cost and common dependencies are explicitly managed.

## IX. Conclusion

Contemporary network technologies occupy distinct points in a multidimensional design space. Ethernet and optical access lead in predictable fixed capacity; Wi-Fi leads in economical local flexibility; personal-area systems support localized low-power interaction; cellular systems combine mobility, identity, and managed wide-area service; LPWAN extends the life and reach of constrained devices; NTN supplies geographic continuity; and TSN or DetNet adds engineered timing and loss behavior. SDN, NFV, slicing, and edge computing operate across these families by making control, isolation, and service placement programmable.

Heterogeneity is a design property, not a temporary stage before one technology wins. End-to-end quality depends on the least controlled link, compute node, identity dependency, or administrative boundary. Progress therefore depends on trustworthy cross-domain orchestration, verifiable assurance, sustainable lifecycle design, and reproducible evidence as 5G-Advanced, NTN integration, and IMT-2030 expand the capability and complexity of converged networks.

## References


[1] IEEE, "IEEE std 802-2024, IEEE standard for local and metropolitan area networks: Overview and architecture," 2025. [Online]. Available: https://standards.ieee.org/ieee/802/10894/

[2] N. Finn *et al.*, "Deterministic networking architecture," Internet Engineering Task Force, RFC 8655, Oct. 2019. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc8655

[3] A. Nasrallah *et al.*, "Ultra-low latency networks: The IEEE TSN and IETF DetNet standards and related 5G research," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 88–145, 2019.

[4] Open Networking Foundation, "Software-defined networking definition," accessed: Aug. 2026. [Online]. Available: https://opennetworking.org/sdn-definition/

[5] European Telecommunications Standards Institute, "Network functions virtualisation (NFV)," technical-group overview; accessed: Aug. 2026. [Online]. Available: https://www.etsi.org/technical-groups/nfv/

[6] IEEE 802.3 Working Group, "Ethernet standards and working-group resources," accessed: Aug. 2026. [Online]. Available: https://grouper.ieee.org/groups/802/3/

[7] M. Besta and T. Hoefler, "Slim fly: A cost effective low-diameter network topology," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2014, pp. 348–359.

[8] Y. Kang, X. Wang, and Z. Lan, "Study of workload interference with intelligent routing on dragonfly," in *SC22: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2022, pp. 1–14.

[9] Y. Ajima, T. Kawashima, T. Okamoto, N. Shida, K. Hirai, T. Shimizu, S. Hiramoto, Y. Ikeda, T. Yoshikawa, K. Uchida, and T. Inoue, "The Tofu Interconnect D," in *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, 2018, pp. 646–654.

[10] ITU-T, "G.9804.1: Higher speed passive optical networks—general requirements," 2025, with amendments through 2025. [Online]. Available: https://www.itu.int/rec/T-REC-G.9804.1

[11] M. Wollschlaeger, T. Sauter, and J. Jasperneite, "The future of industrial communication: Automation networks in the era of the internet of things and industry 4.0," *IEEE Industrial Electronics Magazine*, vol. 11, no. 1, pp. 17–27, 2017.

[12] B. Bellalta, "IEEE 802.11ax: High-efficiency WLANs," *IEEE Wireless Communications*, vol. 23, no. 1, pp. 38–46, 2016.

[13] E. Khorov, A. Kiryanov, A. Lyakhov, and G. Bianchi, "A tutorial on IEEE 802.11ax high efficiency WLANs," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 1, pp. 197–216, 2019.

[14] IEEE, "IEEE std 802.11be-2024, wireless LAN MAC and PHY specifications—amendment: Enhancements for extremely high throughput," 2025. [Online]. Available: https://standards.ieee.org/ieee/802.11be/7516/

[15] ——, "IEEE std 802.15.4-2020, IEEE standard for low-rate wireless networks," 2020. [Online]. Available: https://standards.ieee.org/standard/802_15_4-2020.html

[16] J. G. Andrews *et al.*, "What will 5G be?" *IEEE Journal on Selected Areas in Communications*, vol. 32, no. 6, pp. 1065–1082, 2014.

[17] F. Boccardi, R. W. Heath, Jr., A. Lozano, T. L. Marzetta, and P. Popovski, "Five disruptive technology directions for 5G," *IEEE Communications Magazine*, vol. 52, no. 2, pp. 74–80, 2014.

[18] X. Foukas, G. Patounas, A. Elmokashfi, and M. K. Marina, "Network slicing in 5G: Survey and challenges," *IEEE Communications Magazine*, vol. 55, no. 5, pp. 94–100, 2017.

[19] 3GPP, "Release 18: 5G-advanced system enhancements," accessed: Aug. 2026. [Online]. Available: https://www.3gpp.org/specifications-technologies/releases/release-18

[20] M. Giordani, M. Polese, M. Mezzavilla, S. Rangan, and M. Zorzi, "Toward 6G networks: Use cases and technologies," *IEEE Communications Magazine*, vol. 58, no. 3, pp. 55–61, 2020.

[21] W. Saad, M. Bennis, and M. Chen, "A vision of 6G wireless systems: Applications, trends, technologies, and open research problems," *IEEE Network*, vol. 34, no. 3, pp. 134–142, 2020.

[22] ITU-R, "Recommendation M.2160-0: Framework and overall objectives of the future development of IMT for 2030 and beyond," Nov. 2023. [Online]. Available: https://www.itu.int/rec/R-REC-M.2160-0-202311-I/en

[23] International Telecommunication Union, "IMT-2030: Technical requirements for the 6G future," Mar. 2026. [Online]. Available: https://www.itu.int/hub/2026/03/imt-2030-technical-requirements-for-the-6g-future/

[24] S. Farrell, "Low-power wide area network (LPWAN) overview," Internet Engineering Task Force, RFC 8376, May 2018. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc8376

[25] K. Mekki, E. Bajic, F. Chaxel, and F. Meyer, "A comparative study of LPWAN technologies for large-scale IoT deployment," *ICT Express*, vol. 5, no. 1, pp. 1–7, 2019.

[26] N. Naik, "LPWAN technologies for IoT systems: Choice between ultra narrow band and spread spectrum," in *2018 IEEE International Systems Engineering Symposium*, 2018, pp. 1–8.

[27] M. Centenaro, L. Vangelista, A. Zanella, and M. Zorzi, "Long-range communications in unlicensed bands: The rising stars in the IoT and smart city scenarios," *IEEE Wireless Communications*, vol. 23, no. 5, pp. 60–67, 2016.

[28] U. Raza, P. Kulkarni, and M. Sooriyabandara, "Low power wide area networks: An overview," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 2, pp. 855–873, 2017.

[29] 3GPP, "Non-terrestrial networks (NTN)," May 2024, technology overview. [Online]. Available: https://www.3gpp.org/technologies/ntn-overview

[30] O. Kodheli *et al.*, "Satellite communications in the new space era: A survey and future challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 1, pp. 70–109, 2021.

[31] I. F. Akyildiz, A. Lee, P. Wang, M. Luo, and W. Chou, "A roadmap for traffic engineering in SDN–OpenFlow networks," *Computer Networks*, vol. 71, pp. 1–30, 2014.

[32] D. Kreutz *et al.*, "Software-defined networking: A comprehensive survey," *Proceedings of the IEEE*, vol. 103, no. 1, pp. 14–76, 2015.

[33] R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. De Turck, and R. Boutaba, "Network function virtualization: State-of-the-art and research challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 236–262, 2016.

[34] B. Han, V. Gopalakrishnan, L. Ji, and S. Lee, "Network function virtualization: Challenges and opportunities for innovations," *IEEE Communications Magazine*, vol. 53, no. 2, pp. 90–97, 2015.

[35] Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "A survey on mobile edge computing: The communication perspective," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2322–2358, 2017.

[36] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing*, 2012, pp. 13–16.

[37] M. Copik, K. Taranov, A. Calotoiu, and T. Hoefler, "rFaaS: Enabling high performance serverless with RDMA and leases," in *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2023, pp. 897–907.

[38] J. Snyder and A. R. Lebeck, "Fast convergence to fairness for reduced long flow tail latency in datacenter networks," in *2022 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2022, pp. 1007–1017.

[39] Y. Kang, X. Wang, and Z. Lan, "Q-adaptive: A multi-agent reinforcement learning based routing on dragonfly network," in *Proceedings of the 30th International Symposium on High-Performance Parallel and Distributed Computing*, 2021, pp. 189–200.

[40] O. Garcia-Morchon, S. Kumar, and M. Sethi, "Internet of things (IoT) security: State of the art and challenges," Internet Research Task Force, RFC 8576, Apr. 2019. [Online]. Available: https://datatracker.ietf.org/doc/html/rfc8576

[41] S. Rose, O. Borchert, S. Mitchell, and S. Connelly, "Zero trust architecture," National Institute of Standards and Technology, NIST Special Publication 800-207, Aug. 2020.

[42] S. Chunduri, T. Groves, B. Mendygral, B. Austin, J. Balma, K. Kandalla, K. Kumaran, G. Lockwood, S. Parker, S. Warren, N. Wichmann, and N. J. Wright, "GPCNeT: Designing a benchmark suite for inducing and measuring contention in HPC networks," in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (SC)*, 2019, pp. 1–33.

[43] Z. Chen, Z. Zhao, Z. Li, J. Shao, S. Liu, and Y. Xu, "SDT: A low-cost and topology-reconfigurable testbed for network research," in *2023 IEEE International Conference on Cluster Computing (CLUSTER)*, 2023, pp. 343–353.