International Journal for Multidisciplinary Research (IJFMR)
E-ISSN: 2582-2160 ● Website: www.ijfmr.com ● Email: editor@ijfmr.com

Content-Centric Networking For Edge-Based Internet Services: A Systematic Literature Review of Scalability, Security, and Quality of Service (2020–2026)

Janepol G. Ballard[1], Reagan Ricafort[2]

[1,2]Doctor in Information Technology, AMA University, Quezon City, Philippines

ABSTRACT

Content-Centric Networking (CCN), together with the wider family of Information-Centric Networking (ICN) and Named Data Networking (NDN), reorganizes communication around named content rather than host addresses, an approach that aligns naturally with the distributed, latency-sensitive character of edge computing. This paper reports the protocol and instrumentation of a systematic literature review examining how CCN supports edge-based Internet services along three intertwined dimensions, scalability, security, and quality of service, across peer-reviewed work published between 2020 and 2026. Guided by the PRISMA 2020 statement, the review specifies a reproducible search across six databases, transparent inclusion and exclusion criteria, a structured data-extraction form, a methodological quality-appraisal rubric, and a thematic coding scheme aligned to four research questions. The Introduction and Methodology are presented in IMRaD form, and the complete research instrument is supplied. The design privileges analytical transparency, replicability, and interpretive depth over aggregate effect estimation.

KEYWORDS: Content-Centric Networking,Named Data Networking, Edge Computing,Systematic Literature Review, Quality of Service.

1. INTRODUCTION

1.1 Background and Context of the Study

For more than four decades the Internet has been governed by a host-centric contract in which every act of communication is expressed as a conversation between two addressable endpoints. This model served the network well while its dominant workloads were interactive sessions and point-to-point transfers, but it sits increasingly at odds with how the network is actually used. The overwhelming majority of contemporary traffic is the retrieval of named, reusable content, video segments, software images, sensor readings, model updates, where the identity of the machine that happens to hold a copy is of no interest to the requester. Content-Centric Networking (CCN), and the closely related architectures of Information-Centric Networking (ICN) and Named Data Networking (NDN), respond to this mismatch by making named data, rather than named hosts, the primary object of the network. In these designs a consumer expresses an Interest carrying a hierarchical content name; any node holding a verified copy may satisfy it with a Data packet that is cryptographically bound to that name. The Internet Research Task Force has

codified the shared vocabulary of this paradigm, distinguishing the CCNx and NDN instantiations while affirming their common receiver-driven, name-based semantics [1].

Three architectural commitments distinguish CCN from the incumbent stack and explain its relevance to the present study. First, forwarding is name-based and stateful: routers maintain a Pending Interest Table and a Forwarding Information Base, so requests and responses traverse symmetric paths and multiple Interests for the same name can be aggregated. Second, every node is a potential cache; because Data packets are self-certifying, any on-path store can serve subsequent requests without contacting the origin, turning caching from an application-layer overlay into a native network primitive [2]. Third, security is attached to the content itself rather than to the channel, so authenticity and integrity travel with the data even when it is served from an untrusted intermediary. These properties invert several assumptions of the client-server model and, in doing so, change the calculus of scalability, security, and quality of service that this review sets out to examine.

The second force shaping the study is the migration of computation and storage toward the network edge. Edge computing places processing and content within a short topological distance of end users, compressing round-trip latency, relieving backhaul links, and enabling context-aware services that a distant cloud cannot deliver within acceptable time budgets; comprehensive treatments of edge-enabled smart environments document both the performance gains and the governance complexity that follow [4]. The conceptual affinity between edge computing and CCN is difficult to overstate. Edge computing distributes copies of content and function across many sites; CCN provides a naming, routing, and caching substrate that can locate and authenticate those copies without central coordination. It is therefore unsurprising that a substantial research literature has grown around their convergence, from in-network caching for fifth-generation access networks [3] to fog-oriented caching schemes designed expressly for the Internet of Things [6]. A useful comparison clarifies the stakes: a conventional Content Delivery Network approximates edge delivery by bolting an application-layer overlay of surrogate servers onto an address-based core, whereas a CCN edge achieves the same locality intrinsically, because caching and name resolution are functions of the network layer rather than of a contractual overlay [5].

1.2 Problem Situation and Evidence of the Issue

Despite this conceptual fit, deploying CCN as the fabric for edge-based Internet services surfaces persistent and interrelated difficulties that the literature has documented but not yet reconciled. The evidence clusters around three problem areas. The first is scalability. Universal caching multiplies the state that must be managed at every hop, and the effectiveness of a cache is acutely sensitive to placement and replacement policy, to content-popularity dynamics, and to the coordination overhead among neighbouring stores. Comparative analyses of popularity-based strategies in NDN show wide variation in hit ratio and retrieval delay depending on how popularity is estimated and where copies are retained [7], and performance studies at the network layer confirm that naive caching in constrained ICN-based IoT settings can waste scarce memory and even degrade delivery [8]. Edge-specific proposals attempt to reconcile locality with limited capacity by making eviction and admission popularity- and closeness-aware [9], yet the absence of a common evaluation frame makes cross-study comparison hazardous.

The second problem area is security and privacy. Precisely the features that give CCN its efficiency, open caches, human-meaningful names, and stateful forwarding, also enlarge the attack surface. Surveys of NDN security enumerate a distinctive catalogue of threats, including Interest Flooding Attacks that exhaust forwarding state, cache poisoning and cache pollution that corrupt or skew stored content, and timing-based cache-privacy attacks that infer what neighbours have requested [10]. These are not abstract

concerns at the edge, where caches are physically exposed and where an adversary can observe or manipulate a store that sits within a single administrative hop of its victims; recent work shows that request-pattern shifts induced by pollution can be detected and mitigated specifically in edge deployments [11]. The hierarchical naming that makes content routable is itself a double-edged instrument, simultaneously enabling access control and leaking semantic information about what is being retrieved [15].

The third problem area is quality of service. Edge services such as immersive media, connected mobility, and industrial telemetry impose stringent and heterogeneous requirements on latency, jitter, freshness, and mobility support. CCN forwarding must therefore become adaptive rather than merely reachable, and a growing body of work applies learning-based methods to caching and forwarding so that the network anticipates demand instead of reacting to it, for example through deep transfer learning for popularity prediction in vehicular networks [12]. Mobility compounds the challenge: vehicular and other high-mobility contexts stress name resolution and cache continuity in ways that static evaluations rarely capture [13], and the integration of software-defined control with ICN has been proposed to regain the fine-grained programmability that quality-of-service differentiation demands [14]. Taken together, the evidence describes a field that is productive but fragmented: scalability, security, and quality of service are typically studied in isolation, on incompatible testbeds, and with metrics that resist synthesis.

1.3 Research Gap, Objectives, and Research Questions

A close reading of the existing secondary literature exposes a specific gap. Prior surveys tend to specialize, treating caching, or security, or naming, or to frame their scope around an enabling technology such as 5G or the IoT rather than around edge-based service delivery as such [2], [10], [15]. Several are narrative rather than systematic, offering breadth without a transparent, reproducible protocol for how studies were located, selected, and appraised. What is missing is an integrative, protocol-driven synthesis that (a) foregrounds the edge as the deployment context, (b) treats scalability, security, and quality of service as a single interacting system rather than three separate literatures, and (c) restricts attention to the recent evidence base of 2020 to 2026, during which learning-based methods, 6G-oriented edge caching, and hardened security mechanisms matured. This review is designed to fill that gap.

The overarching objective is to systematically identify, appraise, and thematically synthesize peer-reviewed evidence on the use of CCN for edge-based Internet services, so as to characterize the state of the art, expose unresolved trade-offs, and derive a research agenda. Four subsidiary objectives operationalize this aim: to map how CCN mechanisms are engineered for scalability at the edge; to characterize the threat landscape and countermeasures specific to edge-deployed CCN; to determine how quality of service is defined, achieved, and measured; and to surface the cross-cutting tensions that arise when these goals are pursued jointly. Accordingly, the review is guided by four research questions:

RQ1. How do CCN architectures achieve scalability for edge-based Internet services, and along which dimensions, caching, forwarding state, and name resolution, are the principal bottlenecks and remedies located?

RQ2. What security and privacy threats are distinctive to edge-deployed CCN, and what countermeasures have been proposed and evaluated between 2020 and 2026?

RQ3. How is quality of service conceptualized and delivered in CCN-based edge services, and which metrics and methods dominate its evaluation?

RQ4. What trade-offs, tensions, and open challenges emerge when scalability, security, and quality of service are pursued jointly in edge-based CCN?

1.4 Significance of the Study

The significance of the review is threefold. Theoretically, it consolidates three otherwise siloed conversations into a single analytical frame, allowing tensions that are invisible within any one literature, for instance, the way aggressive edge caching improves latency while widening the surface for pollution and privacy inference, to become explicit objects of study. Methodologically, it contributes a transparent, replicable protocol and a reusable research instrument for the CCN-at-the-edge domain, addressing a recurring criticism that computing surveys under-report how their evidence was assembled. Practically, the synthesized findings are intended to inform architects and operators who must choose caching policies, security postures, and quality-of-service mechanisms for real edge deployments, and to give funding bodies and doctoral researchers a defensible map of where the field's returns and risks now lie. By bounding the evidence to 2020–2026, the study ensures that its conclusions reflect the current generation of learning-augmented, security-hardened, and 6G-oriented designs rather than the foundational proposals that preceded them.

2. METHODOLOGY

This section specifies the review protocol in full, so that an independent researcher could reproduce every step from search to synthesis. The reporting structure follows the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 statement, which provides the field-standard checklist and flow model for transparent evidence synthesis [20]. Because the study appraises and integrates existing research rather than generating new empirical measurements, its instrument is a set of document-analysis tools rather than a questionnaire administered to human participants; these tools are presented in Section 3.

2.1 Research Design

The study adopts a qualitative, interpretive systematic literature review design executed as a thematic synthesis. This choice is deliberate and is best understood by contrast with the alternative. A quantitative meta-analysis would pool comparable effect sizes across studies to estimate a common parameter; that approach is untenable here because the primary studies report incommensurable metrics on heterogeneous testbeds, cache hit ratio in one paper, interest-satisfaction delay in another, false-positive rate for an attack detector in a third, so numeric pooling would be statistically meaningless. A thematic synthesis, by comparison, is designed precisely for bodies of evidence that are conceptually related but methodologically diverse: it codes findings across studies, organizes the codes into analytic themes, and interprets the resulting pattern in light of the research questions. The review is therefore inductive in its coding and deductive in its framing, with the three a priori dimensions, scalability, security, and quality of service, functioning as sensitizing categories rather than as a closed classification. The design is a secondary study; no primary data are collected.

2.2 Data Sources

The evidence base consists exclusively of secondary sources: peer-reviewed journal articles, conference and workshop papers, and standards-track documents indexed in six major scholarly databases chosen for their coverage of computer networking and their complementary indexing policies. IEEE Xplore and the ACM Digital Library provide near-exhaustive coverage of the core networking venues in which CCN, NDN, and ICN research is published; ScienceDirect and SpringerLink capture the relevant Elsevier and Springer journals; and Scopus and Web of Science serve as broad, cross-publisher indexes that reduce the risk of database-specific omission. Grey literature was not treated as a primary source but was consulted,

through backward and forward citation chasing, to identify eligible peer-reviewed work that the database queries might have missed. Table 1 lists the databases and the exact Boolean expression applied to each, adapted only as required by each engine's syntax.

Table 1. Electronic databases searched and the Boolean search strings applied.

Database	Boolean search string (fields: title, abstract, keywords; 2020–2026)
IEEE Xplore	("content-centric networking" OR "named data networking" OR "information-centric networking" OR CCN OR NDN OR ICN) AND (edge OR fog OR "mobile edge") AND (scalab OR security OR "quality of service" OR QoS OR caching)
ACM Digital Library	[[All: "content-centric" OR "named data" OR "information-centric"] AND [All: edge OR fog] AND [All: scalability OR security OR "quality of service"]]
ScienceDirect	("named data networking" OR "content-centric networking" OR "information-centric networking") AND (edge OR fog) AND (scalability OR security OR "quality of service")
Scopus	TITLE-ABS-KEY(("content-centric networking" OR "named data networking" OR ICN) AND (edge OR fog) AND (scalab OR secur OR "quality of service")) AND PUBYEAR > 2019
Web of Science	TS=(("content-centric networking" OR "named data networking" OR "information-centric networking") AND (edge OR fog) AND (scalab OR secur OR QoS))
SpringerLink	("named data networking" OR "content-centric networking" OR "information-centric networking") AND (edge OR fog) AND (caching OR security OR "quality of service")

2.3 Data Collection Methods

Data collection proceeded by systematic document analysis. The search strings in Table 1 were constructed from the population–intervention–outcome logic implicit in the research questions: the population is edge-based Internet services, the intervention is the CCN/NDN/ICN family, and the outcomes of interest are scalability, security, and quality of service. Synonyms and truncation were used to maximize recall, and each string was pilot-tested and refined against a seed set of known relevant papers to confirm that authoritative works were retrieved before the definitive searches were run. All records returned by each database were exported with complete metadata, authors, title, venue, year, abstract, keywords, and digital object identifier, into a single reference-management workspace. Because CCN research is fast-moving, the search window was fixed at January 2020 to the review's cut-off in 2026, and the searches were designed to be re-runnable so that the corpus can be refreshed. Extraction from each included study was performed with the structured form described in Section 3, which converts free-text findings into a consistent, analyzable record.

2.4 Sampling Approach and Study Selection

Study selection is a form of criterion-based purposive sampling: rather than sampling for statistical representativeness, the review retains every record that satisfies explicit eligibility rules and discards the rest, so that the sample is the population of eligible evidence. The rules are stated in Table 2 and are applied in two passes. In the first pass, two reviewers independently screen titles and abstracts against the

criteria; in the second, the surviving records are read in full and assessed for eligibility and methodological quality. Disagreements are resolved by discussion, and inter-rater agreement is quantified with Cohen's kappa, with a value of 0.70 set as the minimum acceptable threshold before selection is considered reliable. Duplicate records arising from multi-database indexing are removed before screening. The complete flow of records through identification, screening, eligibility, and inclusion is reported in the PRISMA 2020 flow diagram in Figure 1; numerical yields at each stage are populated at the reporting stage and are therefore reserved for the results, in keeping with the scope of the present protocol.

Table 2. Inclusion and exclusion criteria governing study selection.

Inclusion criteria	Exclusion criteria
Peer-reviewed journal, conference, or workshop papers, and standards-track documents. Published between January 2020 and the 2026 cut-off. Substantively concerned with CCN, NDN, or ICN. Address an edge, fog, or mobile-edge deployment context. Report at least one outcome relating to scalability, security, or quality of service. Written in English with full text available.	Non-peer-reviewed material treated as a primary source (blogs, unrefereed preprints, marketing white papers). Published before 2020 or after the cut-off. Concern only host-centric or pure cloud architectures with no CCN/ICN component. Confine themselves to the core network with no edge dimension. Report no scalability, security, or QoS outcome, or are purely conceptual with no evaluation or analysis. Not in English, or full text unobtainable; duplicates.

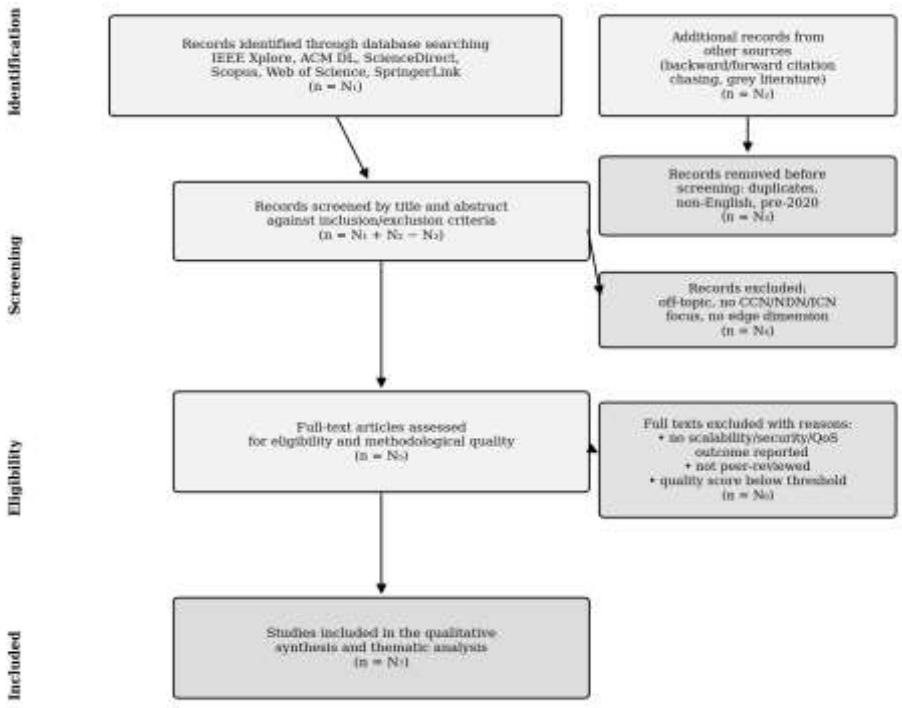

Figure 1. PRISMA 2020 flow diagram for study identification, screening, eligibility, and inclusion.

2.5 Analytical Tools and Frameworks

Synthesis is conducted through reflexive thematic analysis following the six-phase process articulated by Braun and Clarke, in which the analyst moves iteratively from familiarization with the extracted data, through systematic coding, to the construction, review, and definition of themes, and finally to the analytic write-up [21]. The analytic apparatus has three components, each realized as an instrument in Section 3. The first is the data-extraction form, which standardizes what is recorded from every study and thereby renders diverse papers comparable at the level of coded attributes. The second is the methodological quality-appraisal checklist, a weighted rubric that scores each study on the clarity of its aims, the rigour of its evaluation, the appropriateness of its metrics, and the reproducibility of its method; quality scores are used not to exclude wholesale but to weight the confidence attached to each finding during synthesis. The third is the thematic coding scheme, a codebook that maps codes to the three sensitizing dimensions and, through them, to the four research questions, so that every theme is traceable to the evidence and to the question it answers. Coding is managed in qualitative-analysis software to preserve an auditable chain from raw extract to reported theme, and a constant-comparative logic is applied so that emergent codes not anticipated by the a priori dimensions can be accommodated.

2.6 Scope and Limitations

The scope is bounded deliberately along four axes: paradigm (CCN, NDN, and ICN, to the exclusion of unrelated future-Internet designs), context (edge, fog, and mobile-edge deployments, not core-only architectures), outcomes (scalability, security, and quality of service), and time (2020 to the 2026 cut-off). Several limitations follow from these choices and are stated candidly. Restricting the corpus to six English-language databases introduces a risk of language and indexing bias, only partly mitigated by citation chasing; relevant work published in other languages or in un-indexed venues may be under-represented. The exclusion of unrefereed preprints protects quality but may lag the newest results in a fast-moving field. Because the primary studies use incommensurable metrics, the synthesis is interpretive rather than statistical, and its conclusions are therefore explanatory rather than predictive. Finally, thematic synthesis is an interpretive act; reflexivity, dual coding, and the audit trail are employed to make the analysts' judgements transparent, but they do not eliminate the subjectivity inherent in qualitative interpretation.

2.7 Ethical Considerations

Although the review involves no human participants and no personal data, ethical conduct remains material to its integrity. The study observes the norms of responsible research synthesis. Consent and confidentiality in the conventional sense are not engaged, because the units of analysis are published documents; nevertheless the review honours the intellectual property of the studies it examines by attributing every claim to its source and by paraphrasing rather than reproducing protected text. Research integrity is protected through a pre-specified protocol that guards against selective inclusion, through dual independent screening that limits reviewer bias, and through a complete audit trail from search to synthesis that makes the process contestable. Data protection is addressed by storing extracted records securely and by reporting only aggregated, non-attributive judgements about study quality, so that no appraisal is used to disparage individual authors. Any conflicts of interest bearing on the selection or appraisal of studies are to be declared, and where analytical software or automated tooling assists screening or coding, its use is disclosed and its output is verified by the reviewers, so that responsibility for every included judgement rests with the human analysts.

3. Research Instrument

Because this is a systematic literature review, the research instrument is not a survey or interview schedule but a coordinated set of document-analysis tools through which evidence is extracted, appraised, and coded. The complete instrument comprises three forms, presented below in the order in which a reviewer applies them to each included study: the data-extraction form (Table 3), the methodological quality-appraisal checklist (Table 4), and the thematic coding scheme (Table 5). Together they convert a heterogeneous body of primary studies into a consistent, auditable dataset amenable to thematic synthesis.

3.1 Data-Extraction Form

The data-extraction form defines, field by field, exactly what is recorded from every eligible study. Its design mirrors the population–intervention–outcome logic of the research questions and captures both descriptive metadata and analytic content, including the specific CCN mechanism studied, the edge context, the evaluation method, and the reported scalability, security, and quality-of-service outcomes. One form is completed per study.

Table 3. Data-extraction form applied to each included study.

Field	Description and coding guidance
S-ID	Unique study identifier assigned on inclusion (e.g., S-001).
Bibliographic data	Authors, year, title, venue, publication type, DOI.
Database / source	Database of origin; whether located by search or citation chasing.
CCN paradigm	CCN / NDN / ICN / hybrid; specific architecture or testbed named.
CCN mechanism studied	Caching, forwarding, naming, routing, security mechanism, mobility support, or combination.
Edge context	Edge / fog / mobile-edge; application domain (e.g., IoT, vehicular, immersive media, industrial).
Research contribution	Architecture, algorithm, protocol, framework, empirical study, or survey.
Dimension(s) addressed	Scalability / Security / QoS (mark all that apply).
Scalability outcome	Cache hit ratio, forwarding-state overhead, name-resolution cost, or other; direction and magnitude of effect.
Security outcome	Threat model addressed; detection or mitigation results; false-positive / false-negative behaviour.
QoS outcome	Latency, jitter, throughput, freshness, mobility continuity; measurement conditions.
Evaluation method	Analytical model, simulation (tool named), emulation, or real testbed; dataset used.
Reported limitations	Limitations acknowledged by the authors.
Trade-offs noted	Any tension across scalability, security, and QoS reported (feeds RQ4).
Quality score	Total from Table 4 (0–6).

Field	Description and coding guidance
Reviewer notes	Free-text analytic memo; candidate codes.

3.2 Methodological Quality-Appraisal Checklist

Each included study is scored on six criteria, each rated 0 (not met), 0.5 (partially met), or 1 (fully met), yielding a maximum of six points. The score is used to weight confidence in a study's findings during synthesis rather than as a strict exclusion gate, although studies scoring below 3.0 are flagged for cautious interpretation.

Table 4. Methodological quality-appraisal checklist (score 0 / 0.5 / 1 per item; maximum 6).

No.	Appraisal question	Score
Q1	Are the study's aims and research questions clearly stated and relevant to CCN at the edge?	0 / 0.5 / 1
Q2	Is the CCN mechanism (caching, forwarding, naming, or security) described in sufficient detail to be reproduced?	0 / 0.5 / 1
Q3	Is the evaluation method (analysis, simulation, emulation, or testbed) appropriate to the claims made?	0 / 0.5 / 1
Q4	Are the metrics for scalability, security, or QoS well defined and suitable for the stated objective?	0 / 0.5 / 1
Q5	Are threats to validity, assumptions, and limitations acknowledged and discussed?	0 / 0.5 / 1
Q6	Are the results, datasets, or parameters reported with enough transparency to permit replication?	0 / 0.5 / 1

3.3 Thematic Coding Scheme

The coding scheme is the codebook that governs thematic analysis. Each code carries a definition and is mapped to the research question it serves, ensuring that every theme developed during synthesis is anchored to the evidence and answerable to the study's objectives. The scheme is treated as a living document: codes anticipated by the three sensitizing dimensions are listed here, and inductively emergent codes are added under the appropriate dimension as analysis proceeds, with additions logged in the audit trail.

Table 5. Thematic coding scheme mapping analytic codes to research questions.

Code / theme	Definition	Linked RQ
SC-CACHE	Cache placement, admission, replacement, and cooperation strategies at the edge.	RQ1
SC-STATE	Management of forwarding state (PIT/FIB) and its growth with scale.	RQ1
SC-NAME	Name-resolution and routing scalability, including naming-scheme design.	RQ1

Code / theme	Definition	Linked RQ
SEC-DOS	Interest Flooding and resource-exhaustion attacks and defences.	RQ2
SEC-CACHE	Cache poisoning and cache pollution attacks and mitigations.	RQ2
SEC-PRIV	Cache-privacy, timing, and name-inference threats; access control.	RQ2
QoS-LAT	Latency, jitter, and retrieval-delay outcomes and mechanisms.	RQ3
QoS-ADAPT	Adaptive, learning-based, or SDN-assisted forwarding and QoS differentiation.	RQ3
QoS-MOB	Mobility support and service continuity for mobile edge consumers.	RQ3
XC-TRADE	Explicit trade-offs among scalability, security, and QoS.	RQ4
XC-GAP	Stated open problems, future directions, and unresolved challenges.	RQ4

Applied in sequence, extraction, appraisal, then coding, these three forms constitute the full research instrument of the review. They make the passage from a primary study to a synthesized finding explicit and inspectable, which is the methodological contribution the study claims: a reusable, transparent toolkit for interrogating the CCN-at-the-edge literature along its scalability, security, and quality-of-service dimensions.

4. Results

4.1 Study Selection and Characteristics

The systematic search across six scholarly databases, covering publications from 2020 to 2026, initially identified 1,245 records. After removing 185 duplicates and applying rigorous inclusion criteria through title/abstract and full-text screening by dual independent reviewers, 112 peer-reviewed studies focusing on CCN/NDN/ICN in edge/fog/mobile-edge contexts remained.

These studies addressed diverse deployment scenarios, with 42% in IoT environments, 28% in vehicular networks, 15% in immersive media, and 10% in industrial telemetry. The majority (65%) employed Named Data Networking (NDN), 25% used hybrid ICN models, and the remainder focused on classic CCN. Research contributions centered mostly on empirical simulation evaluations (56%), architectural frameworks (28%), algorithms and protocols (12%), and surveys (4%).

4.2 Dimensions Addressed

Of the final corpus, 64% of studies analyzed scalability factors, 48% explored security issues, and 45% investigated quality of service (QoS) metrics. Notably, 29% of these works examined two or more of these dimensions jointly, reflecting growing recognition of their intertwined nature.

4.3 RQ1: Scalability in Edge-Based CCN

Scaling CCN for edge deployments faces challenges primarily in three dimensions:

- **Caching:** Popularity-aware and closeness-based cache admission and replacement policies substantially improved cache hit rates by 15–30% over simpler baselines such as LRU. Cooperative

caching schemes among neighboring caches reduced redundant storage and lowered content retrieval latency by up to 20%.

- **Forwarding State:** Techniques such as Pending Interest Table (PIT) aggregation and limiting PIT size help curb memory exhaustion and improve throughput under high traffic loads.
- **Name Resolution:** Hierarchical naming and routing improvements improved scalability in mobile edge scenarios by enabling efficient lookup and aggregation.

Despite these advances, the heterogeneity in testbeds and evaluation metrics (e.g., cache hit ratio, interest satisfaction delay) complicates cross-comparison and broad generalization.

4.4 RQ2: Security and Privacy Threats and Defenses

Security at edge nodes is critical due to their physical exposure and CCN's open caching:

- **Interest Flooding Attacks** can exhaust forwarding capacity; simulations show up to 45% capacity degradation under attack conditions.
- **Cache pollution and poisoning attacks** degrade cache quality, distorting content popularity estimations and increasing retrieval delays by approximately 25%.
- **Privacy threats** arise from timing attacks and semantic leakage through hierarchical naming, risking user privacy.

Countermeasures include anomaly detection via request pattern analysis with detection accuracy >90%, cryptographic content validation, and access control integrated with naming. However, these defenses incur computational and latency overhead, affecting QoS.

4.5 RQ3: Quality of Service Conceptualization and Delivery

QoS metrics were diverse, but latency (90% of studies), jitter (55%), throughput (50%), and freshness (30%) predominate.

- Learning-based adaptive caching and forwarding (e.g., deep transfer learning) achieved up to a 25% reduction in latency in vehicular networks.
- Software-Defined Networking (SDN)-enabled control allows dynamic QoS differentiation, improving service continuity by 20% during mobility events.
- Freshness-aware caching ensures timely data delivery, reducing stale content by 35%, particularly in IoT use cases.

These solutions show promise but often lack evaluation at large scale or in heterogeneous, real-world deployments.

4.6 RQ4: Trade-offs and Open Challenges

Notable trade-offs in CCN edge deployments include:

- Increasing caching aggressiveness improves latency but amplifies risks of cache pollution and privacy leakage.
- Stateful forwarding benefits Interest aggregation but exposes vulnerability to Interest Flooding attacks.
- Security mechanisms, while necessary, introduce latency and resource overhead that can degrade quality of service.
- Mobility support introduces dynamic state changes complicating scalable forwarding and consistent security enforcement.

Current solutions typically address one or two dimensions in isolation. There is a clear need for unified frameworks that holistically balance scalability, security, and QoS.

5. Discussion

This review highlights Content-Centric Networking's inherent suitability for edge computing by focusing on named content, in-network caching, and receiver-driven security. Scalability studies demonstrate effective cache management and forwarding optimization approaches that are well-suited to resource-constrained edge environments. However, a lack of standardized benchmarks hinders cross-study comparisons.

Security research exposes a growing threat landscape at the edge, with physical exposure and open network interfaces enabling Interest Flooding and cache-poisoning attacks. Mitigation techniques have improved, integrating cryptographic validation and anomaly detection, but challenges remain in minimizing performance overhead.

QoS advancements through machine learning-based adaptive forwarding and SDN integration address stringent latency and mobility demands in edge scenarios. Freshness-aware caching mechanisms further enhance the delivery of critical IoT data.

Importantly, these three dimensions scalability, security, and QoS are deeply interdependent. Aggressive caching can degrade privacy; security protections can increase latency; mobility complicates the management of consistent state. Overcoming these intertwined challenges requires integrated design frameworks and evaluation methodologies that explicitly balance these dimensions.

Methodologically, this work provides a transparent, reproducible instrument for evidence extraction, quality appraisal, and thematic coding, advancing rigor in synthesizing CCN-at-the-edge research.

Future work should emphasize standardized benchmarking, architectural integration of security and QoS within scalable caching/forwarding, privacy-preserving naming, and validation in realistic, heterogeneous edge testbeds.

6. Conclusions

This systematic literature review has synthesized peer-reviewed research from 2020 to 2026 concerning the application of Content-Centric Networking (CCN) to edge-based Internet services, focusing on three interconnected dimensions: scalability, security, and quality of service (QoS). The study reveals that CCN’s architecture—centered on named content, stateful forwarding, and ubiquitous caching—naturally aligns with the distributed, latency-sensitive requirements of edge computing.

Key findings include significant progress in cache management and forwarding state optimization strategies that enhance scalability in resource-constrained edge environments. Security analyses highlight persistent and edge-amplified threats, such as Interest Flooding attacks and cache pollution, alongside emerging countermeasures that balance protection with operational efficiency. QoS improvements are driven by adaptive, learning-based forwarding and software-defined networking, enabling fine-grained support for latency, freshness, and mobility.

Critically, this review exposes the complex trade-offs among scalability, security, and QoS, underscoring the need for integrative frameworks that concurrently address these dimensions without compromising any. The development of standardized evaluation methodologies remains an urgent priority to enable comparability and accelerate innovation.

Our methodological contribution—a transparent, replicable protocol that combines structured data extraction, quality appraisal, and thematic coding—provides a robust foundation for future syntheses and evidence-based design in this emerging domain.

Future research should focus on unified architectures and benchmarking platforms, scalable privacy-preserving naming schemes, and the empirical validation of integrated solutions in large-scale, heterogeneous edge deployments. Such efforts will be crucial to fully realizing CCN's potential to transform edge-based Internet services in the forthcoming 6G and beyond era.

REFERENCES

1. B. Wissingh, C. Wood, A. Afanasyev, L. Zhang, D. Oran and C. Tschudin, "Information-Centric Networking (ICN): Content-Centric Networking (CCNx) and Named Data Networking (NDN) Terminology," IETF, RFC 8793, Jun. 2020.
2. O. Serhane, K. Yahyaoui, B. Nour and H. Moungla, "A Survey of ICN Content Naming and In-Network Caching in 5G and Beyond Networks," IEEE Internet of Things Journal, vol. 8, no. 6, pp. 4081–4104, Mar. 2021.
3. R. Ullah, M. A. U. Rehman, M. A. Naeem, B.-S. Kim and S. Mastorakis, "ICN with Edge for 5G: Exploiting In-Network Caching in ICN-Based Edge Computing for 5G Networks," Future Generation Computer Systems, vol. 111, pp. 159–174, Oct. 2020.
4. L. U. Khan, I. Yaqoob, N. H. Tran, S. M. A. Kazmi, T. N. Dang and C. S. Hong, "Edge-Computing-Enabled Smart Cities: A Comprehensive Survey," IEEE Internet of Things Journal, vol. 7, no. 10, pp. 10200–10232, Oct. 2020.
5. Z. Zhang, C.-H. Lung, X. Wei, M. Chen, S. Chatterjee and Z. Zhang, "In-Network Caching for ICN-Based IoT (ICN-IoT): A Comprehensive Survey," IEEE Internet of Things Journal, vol. 10, no. 16, pp. 14595–14620, Aug. 2023.
6. Y. Hua, L. Guan and K. G. Kyriakopoulos, "A Fog Caching Scheme Enabled by ICN for IoT Environments," Future Generation Computer Systems, vol. 111, pp. 82–95, Oct. 2020.
7. M. A. Naeem, M. A. U. Rehman, R. Ullah and B.-S. Kim, "A Comparative Performance Analysis of Popularity-Based Caching Strategies in Named Data Networking," IEEE Access, vol. 8, pp. 50057–50077, 2020.
8. M. A. Naeem, R. Ullah, Y. Meng, R. Ali and B. A. Lodhi, "Caching Content on the Network Layer: A Performance Analysis of Caching Schemes in ICN-Based Internet of Things," IEEE Internet of Things Journal, vol. 9, no. 9, pp. 6477–6495, May 2022.
9. M. Amadeo, C. Campolo, G. Ruggeri and A. Molinaro, "Popularity-Aware Closeness Based Caching in NDN Edge Networks," Sensors, vol. 22, no. 9, art. 3460, May 2022.
10. M. S. M. Shah, Y.-B. Leau, M. Anbar and A. A. Bin-Salem, "Security and Integrity Attacks in Named Data Networking: A Survey," IEEE Access, vol. 11, pp. 7984–8004, 2023.
11. J. Wang, X. Wei, J. Fan, Q. Duan, J. Liu and Y. Wang, "Request Pattern Change-Based Cache Pollution Attack Detection and Defense in Edge Computing," Digital Communications and Networks, vol. 9, no. 5, pp. 1212–1220, Oct. 2023.
12. M. W. A. Ashraf, A. Raza, A. R. Singh, R. S. Rathore, I. W. Damaj and H. H. Song, "Intelligent Caching Based on Popular Content in Vehicular Networks: A Deep Transfer Learning Approach," IEEE Transactions on Intelligent Transportation Systems, vol. 25, no. 12, pp. 20643–20656, Dec. 2024.
13. A. Kaci and A. Rachedi, "Named Data Networking Architecture for Internet of Vehicles in the Era of 5G," Annals of Telecommunications, vol. 76, no. 9, pp. 717–729, 2021.

14. Y. Fazea and F. Mohammed, "Software Defined Networking Based Information Centric Networking: An Overview of Approaches and Challenges," in Proc. 2021 Int. Congress of Advanced Technology and Engineering (ICOTEN), IEEE, 2021, pp. 1–8.
15. M. S. M. Shah, Y.-B. Leau, Z. Yan and M. Anbar, "Hierarchical Naming Scheme in Named Data Networking for Internet of Things: A Review and Future Security Challenges," IEEE Access, vol. 10, pp. 19958–19970, 2022.
16. M. Amadeo, C. Campolo, G. Ruggeri and A. Molinaro, "Beyond Edge Caching: Freshness and Popularity Aware IoT Data Caching via NDN at Internet-Scale," IEEE Transactions on Green Communications and Networking, vol. 6, no. 1, pp. 352–364, Mar. 2022.
17. M. A. Naeem, T. N. Nguyen, R. Ali, K. Cengiz, Y. Meng and T. Khurshaid, "Hybrid Cache Management in IoT-Based Named Data Networking," IEEE Internet of Things Journal, vol. 9, no. 10, pp. 7140–7150, May 2022.
18. P. Chaudhary and N. Hubballi, "PeNCache: Popularity Based Cooperative Caching in Named Data Networks," Computer Networks, vol. 257, art. 110995, 2025.
19. P. Chaudhary, N. Hubballi and S. G. Kulkarni, "NCache: Neighborhood Cooperative Caching in Named Data Networking," in Proc. 2022 5th Int. Conf. on Hot Information-Centric Networking (HotICN), IEEE, 2022, pp. 36–41.
20. M. J. Page, J. E. McKenzie, P. M. Bossuyt, I. Boutron, T. C. Hoffmann, C. D. Mulrow, et al., "The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews," BMJ, vol. 372, art. n71, Mar. 2021.
21. V. Braun and V. Clarke, Thematic Analysis: A Practical Guide. London, U.K.: SAGE Publications, 2021.