

Privacy-Preserving Smart Metering in Smart Grids: A Systematic Survey and Multidimensional Taxonomy

Farid Zaredar, Mobin Aghamirkarimi, Matin Aghamirkarimi, and Morteza Amini

Abstract—Smart meters support billing, grid operation, demand response, analytics, and distributed-energy services, but their fine-grained measurements can also reveal persistent patterns of household and business activity. Evaluating privacy-preserving solutions therefore requires more than identifying the mechanism used; it requires examining what information each observer receives, which result the application releases, and which trust and lifecycle assumptions sustain the claimed protection. This paper presents a systematic survey of privacy-preserving smart metering based on 23 prior surveys and 150 primary studies published from 2010 to the final search cutoff. Searches of IEEE Xplore, Scopus, and the Web of Science Core Collection were completed on 29 August 2026 and supplemented by citation, title, and publication-version searches, backward and forward snowballing, and targeted Google Scholar searches. We develop a multidimensional taxonomy that connects applications and authorized outputs with information flows, privacy mechanisms, adversary and collusion models, guarantees, residual leakage, operational resilience, efficiency, and evidence maturity. The synthesis shows that mechanism labels alone are poor indicators of assurance. Input confidentiality does not control inference from exact or repeated outputs; data locality does not protect model updates or released models, and anonymity may coexist with linkage through timing, location, settlement, or persistent state. Although later work supports richer functions, stronger adversaries, dynamic participation, and hybrid architectures, construction-level assurance, longitudinal output privacy, coalition and state compromise, lifecycle-safe operation, cross-layer composability, and reproducible deployment evidence remain unevenly addressed. The survey converts these findings into an evidence-derived research agenda and a requirement-driven framework for selecting protocol families according to the required output, acceptable leakage, trust boundary, lifecycle conditions, and available evidence.

Index Terms—advanced metering infrastructure, multidimensional taxonomy, privacy-preserving protocols, secure data aggregation, smart grid privacy, smart meters, systematic literature review, differential privacy.

I. INTRODUCTION

SMART meters turn electricity metering from an occasional billing event into a continuous digital sensing process. Their measurements support accurate billing, outage management, demand response, load forecasting, distributed-energy coordination, and customer-facing services. The same measurements, however, form a longitudinal record of activities inside a home or business. Privacy-preserving smart metering

must therefore reconcile two requirements: the grid must obtain the information required for an authorized service, while no party should learn more about an individual than that service requires.

A. Why Privacy in Smart Metering?

Smart metering is becoming a persistent global data layer rather than a special-purpose billing upgrade. The U.S. Energy Information Administration recorded 140.5 million advanced metering infrastructure meters in 2024, representing 83.6% of the meters in its national table [1]. Looking beyond current deployment, a 2026 industry forecast estimates that the worldwide stock of electricity, gas, and water smart meters will increase from 2.1 billion in 2025 to 3.9 billion in 2035. Electricity meters are projected to remain the largest component, increasing from 1.42 billion to 2.4 billion, while residential electricity-smart-meter penetration rises from 55% to 85% of households worldwide [2].

This expansion accompanies a more dynamic and decentralized electricity system. The International Energy Agency estimates that meeting national energy and climate goals would double the need for power-system flexibility between 2022 and 2030, require more than 80 million km of grid to be added or refurbished by 2040, and raise annual grid investment to more than USD 600 billion by 2030, with particular emphasis on digitalizing and modernizing distribution networks [3]. Smart meters sit at the boundary between this infrastructure and individual premises. They observe the consumption, generation, storage, and charging behavior needed for flexibility services. Their privacy significance therefore grows not only with the number of installed devices, but also with the range of operational and commercial decisions made from their data.

Sampling frequency changes the nature of the information collected. At 15-minute intervals, a smart meter records 96 readings per day and 35,040 readings per year; even at 30-minute intervals, it records 17,520. Compared with twelve monthly readings, these frequencies provide 2,920 and 1,460 times more observations, respectively. Data volume alone does not determine privacy loss, but a dense longitudinal trace exposes temporal patterns that a monthly total cannot reveal, including recurring demand changes, daily routines, and similarities across days.

Primary studies show that the resulting inference risk is technically meaningful. Aggregate load traces can expose appliance events, although the feasibility of detection depends on

The authors are with the Department of Computer Engineering, Sharif University of Technology, Tehran, Iran. First author: Farid Zaredar (e-mail: farid.zaredar78@sharif.edu). Corresponding author: Morteza Amini (e-mail: amini@sharif.edu).

the sampling interval, appliance characteristics, and available auxiliary information [4], [5]. Electricity measurements can also support occupancy inference [6]. More recently, a user-classification experiment on unprotected daily profiles reported re-identification rates of 99%, 73%, and 93% on the ECO, NREL, and Pecan Street datasets, respectively [7]. These are study-specific results, not population-wide attack probabilities. Nevertheless, they demonstrate that a load profile can act as a behavioral and identity signature when the adversary has suitable observations and auxiliary data.

These developments make privacy-preserving smart metering a system-level research problem rather than an isolated cryptographic task. Metering data now flows among meters, gateways, utilities, edge and cloud platforms, market operators, energy-service providers, distributed-energy resources, and charging infrastructure. Each application changes the information collected, the parties allowed to receive it, the result that must be disclosed, and the adversaries against which protection is required. The corresponding literature is distributed across cryptography, communications, machine learning, control, and energy markets, with studies adopting different threat models, privacy objectives, trust assumptions, and evaluation methods. A comprehensive survey is therefore needed to reconstruct these designs at the system level, compare their guarantees under consistent criteria, and determine which privacy, correctness, lifecycle, and deployment requirements remain unresolved.

B. Why Another Survey?

Recent reviews confirm the breadth of privacy research in smart grids. They examine AMI security and privacy, privacy-preserving aggregation and authentication, cryptographic and learning-based mechanisms, and protection across layered smart-grid architectures [8]–[12]. These reviews provide useful introductions to the field. However, their dominant unit of comparison is commonly a mechanism category or stated security feature. Studies grouped under the same label—such as homomorphic aggregation, differential privacy, blockchain, or federated learning—may protect different information, reveal different outputs, trust different entities, and offer substantially different levels of technical evidence.

The need for this survey is therefore not merely that the literature has expanded. Protocol selection requires a common, application-aware unit of analysis. We connect 23 prior surveys with a technical reconstruction of 150 primary studies published between 2010 and the final search cutoff in 2026. For each approach, we examine the supported service, information available to each observer, mechanism that actually provides privacy, adversary and collusion boundary, authorized output and residual leakage, lifecycle behavior, computational and communication costs, and evidentiary basis.

This reconstruction distinguishes an assurance claimed by the authors from a property supported by the construction, equations, proofs, implementation, or experiments. It also separates protection of the transmitted input from protection of the released result, repeated outputs, protocol state, metadata, and application-level decisions. The resulting synthesis is intended to help researchers compare genuinely related designs,

recognize incomplete guarantees, identify defensible research gaps, and select protocol families from application and threat-model requirements rather than from technique names alone.

C. Scope and Research Questions

This survey examines privacy-preserving approaches for smart metering published from 2010 to the final search cutoff in 2026. Prior survey papers are analyzed to establish the conceptual development of the field and to identify limitations in existing reviews. Primary research studies are then examined to provide protocol-level and mechanism-level evidence. The scope covers privacy-preserving approaches supporting smart-metering applications, including billing, grid operation, demand response, data aggregation, analytics, value-added services, distributed-energy and e-mobility services, and customer-side privacy control.

The review is guided by the following research questions:

- RQ1** *How have existing surveys organized the field of privacy-preserving smart metering, and which methodological and analytical dimensions remain insufficiently covered?*
- RQ2** *Which smart-metering applications, supported functions, data products, and authorized outputs are addressed by the primary studies, and how has this coverage evolved from 2010 to the final search cutoff in 2026?*
- RQ3** *Which system architectures, communication settings, trust assumptions, attacker capabilities, and collusion boundaries are considered by the proposed approaches?*
- RQ4** *Which cryptographic, statistical, identity-protection, trusted-hardware, learning-based, and customer-side mechanisms are used, either independently or in combination, to preserve smart-meter privacy?*
- RQ5** *Which information and privacy targets are protected, what formal or empirical privacy guarantees are provided, and which residual leakage channels remain insufficiently addressed?*
- RQ6** *To what extent do the proposed approaches support the required functionality, correctness, dynamic participation, fault tolerance, revocation, recovery, malicious-input handling, and other operational or lifecycle requirements?*
- RQ7** *What computation, communication, storage, and scalability costs are reported, and what level of evidence supports the security, privacy, performance, and deployability claims?*
- RQ8** *Which cross-dimensional tradeoffs and research gaps emerge from the collected evidence, and how can system requirements be mapped to suitable families of privacy-preserving approaches?*

RQ1 establishes the relationship between this survey and previous reviews. RQ2–RQ7 guide the extraction and comparison of evidence from the primary studies. RQ8 integrates the resulting findings into an evidence-derived research agenda and a requirement-oriented protocol-selection framework.

D. Contributions and Organization

Against this background, the survey makes five principal contributions.

- **A systematically constructed and critically assessed evidence base.** We analyze 23 prior surveys to establish the development and remaining coverage gaps of the review literature, and reconstruct 150 primary studies published from 2010 to the final search cutoff in 2026. Technical properties are extracted from the primary publications rather than inherited from earlier reviews. The analysis distinguishes author claims from properties supported by definitions, constructions, proofs, implementations, or experiments, thereby preserving the strength and limitations of the available evidence.
- **A unified, multidimensional framework for protocol analysis.** We introduce a taxonomy that jointly captures the application and authorized output; architecture and information flow; privacy and supporting mechanisms; trust, adversary, and collusion assumptions; privacy targets and guarantees; residual leakage; lifecycle and operational resilience; efficiency; and evidence maturity. This framework separates a mechanism from the guarantee it actually provides and prevents approaches with different outputs, observers, or trust boundaries from being treated as directly comparable.
- **Application-centric and cross-dimensional protocol synthesis.** Rather than organizing the literature only by cryptographic or statistical technique, we compare approaches within the services they support, including billing and settlement, grid operation and aggregation, communication and authentication, analytics and data sharing, distributed-energy and e-mobility services, and customer-side privacy control. We then examine relationships across the complete corpus, showing how authorized outputs, residual leakage, trust assumptions, operational resilience, efficiency, and evaluation maturity jointly determine the assurance provided by a design.
- **An evidence-based assessment of progress and remaining limitations.** We trace the evolution of the field from 2010 to the 2026 search cutoff and identify where later work provides substantive advances, such as richer functionality, stronger collusion models, dynamic participation, fault tolerance, verifiability, and learning-based services. We also identify recurring boundaries that remain insufficiently resolved, including inference from exact or repeated outputs, coalition and state exposure, incomplete lifecycle analysis, non-composing cross-layer protections, and inconsistent deployment evidence.
- **A research agenda and requirement-driven selection framework.** We convert the synthesis into an evidence-derived gap map that associates each recurring weakness with the evidence required to address it. We further map application, output, trust, leakage, lifecycle, and deployment requirements to suitable protocol families and their exclusion conditions. Illustrative design scenarios demonstrate how these criteria can guide protocol selection without implying that one mechanism is universally superior.

The remainder of the survey is organized as follows. Section II describes the review protocol, search and selec-

tion process, evidence extraction, quality assessment, coding procedure, and reproducibility controls. Section III introduces the AMI information flows, privacy risks, authorized-output model, adversary assumptions, and principal privacy-preserving building blocks. Section IV evaluates the coverage of earlier surveys and establishes the remaining review gap. Section V presents the multidimensional taxonomy, and Section VI applies it in an application-centric synthesis of the primary studies. Section VII examines cross-dimensional tradeoffs, evidence maturity, and the evolution of the field. Section VIII develops the research agenda, requirement-to-protocol-family mapping, and illustrative design scenarios. Finally, Section IX discusses threats to validity, and Section X concludes the survey.

II. REVIEW METHODOLOGY

This section describes the literature search, study selection, evidence extraction, quality assessment, and synthesis used to construct and analyze the review corpora.

A. Review Protocol and Reporting Strategy

The review protocol was defined before the comparative coding and synthesis of the full corpus. It follows the systematic-literature-review guidance of Kitchenham and Charters [13]. The identification, screening, eligibility, and inclusion stages are reported using PRISMA 2020 [14]. Search sources, queries, dates, limits, and results are documented according to PRISMA-S [15].

The review combines systematic mapping with critical technical analysis to examine whether the reported privacy, security, functionality, and efficiency claims are supported under the stated assumptions.

The studies differ in their functions, datasets, platforms, security parameters, adversary models, system sizes, and evaluation procedures. Consequently, a statistical meta-analysis is not applied to the complete corpus. Numerical results are compared directly only when the supported functions, assumptions, units, and experimental conditions are sufficiently compatible. Otherwise, structured qualitative comparison is used.

The database search was conducted on 29 August 2026 in IEEE Xplore, Scopus, and the Web of Science Core Collection. The publication window extended from January 2010 to the search date; 2026 is therefore a partial year. The main search combined smart-metering terms with general privacy terms. A supplementary technical search combined the smart-metering terms with mechanism-specific terms and was used in Scopus and Web of Science. No venue-ranking, citation-count, or open-access filter was applied. Complete platform-specific search strings, fields, filters, and export counts are reported in Appendix A.

B. Eligibility Criteria

The identified publications were assigned to three evidence streams. The *primary-study corpus* contains original technical studies used to answer RQ2–RQ8. The *prior-survey corpus* contains review articles used primarily to answer RQ1 and

establish the novelty boundary. The *contextual corpus* contains supporting sources that inform the motivation and foundations but are not counted as primary studies.

Primary studies had to satisfy the publication period, smart-metering relevance, privacy objective, technical contribution, and full-text evidence requirements. Detailed inclusion, exclusion, and evidence-reassignment rules are reported in Appendix B, Tables XVIII and XIX. Methodological weakness or incomplete reporting did not independently cause exclusion when the minimum eligibility requirements were satisfied. These limitations instead affected the study's quality and evidence-maturity assessment.

All eligible primary studies remain part of the review corpus. Studies authored by members of the review team are subject to the same eligibility, extraction, quality-assessment, and critical analysis rules as all other studies.

C. Deduplication and Study Screening

We preserved the original database exports and performed all cleaning on separate working copies. Duplicates were resolved using standardized DOI values. Records without a reliable DOI were compared using normalized titles. Authors, publication year, and venue were then used to confirm possible matches. A similar title was not sufficient grounds for automatically merging two records. Related publication versions were handled according to criterion E8 in Table XIX.

Screening proceeded in two stages using Tables XVIII and XIX. Title-and-abstract screening removed clearly ineligible records. When the available metadata did not support a confident decision, the record was retained for full-text assessment. The full text was then examined to determine the smart-metering setting, protected information, technical contribution, authorized output, and supporting analytical or empirical evidence.

The final review corpus contains 150 primary studies and 23 prior surveys. Contextual sources support the background and interpretation but are not included in either total. Record provenance, related publication versions, and study-level inclusion or exclusion decisions are retained in the supplementary screening documentation.

D. Backward and Forward Snowballing

We supplemented the database results with backward and forward snowballing. Backward snowballing examined the reference lists of the 23 included surveys and technically relevant primary studies. Forward snowballing used citation links in Scopus, Web of Science, and targeted Google Scholar checks to locate later extensions, evaluations, or cryptanalyses. Google Scholar was used only as a supplementary citation- and version-tracing source.

Every candidate found through snowballing was processed using the same deduplication and eligibility rules as a database record. All supplementary searches followed the same final cutoff date as the database searches.

As a final coverage check, the retained corpus was matched against the three database exports. Of the 173 included publications, 167 were retrieved from at least one database.

Six additional eligible publications were identified through backward and forward snowballing and supplementary Google Scholar searches.

E. Evidence-Card Extraction and Quality Assessment

A structured evidence card was completed for each included study to support consistent comparison across the corpus. Evidence was extracted from the original full text and, where relevant, its appendices, supplementary material, or a more complete related version. Prior surveys were not used as substitutes for primary publications when coding technical properties. The extraction fields correspond to the taxonomy in Section V; their complete definitions and permitted values are provided in Appendix D.

We distinguished properties explicitly claimed by the authors from those supported by formal analysis, algorithmic inspection, implementation, or experiment. Interpretations derived from the described design were recorded as reviewer inferences rather than author statements. Numerical results were retained with their units, evaluated entity, experimental conditions, and reported variability when available.

Prior surveys were examined using a separate coverage form to support the comparison in Section IV.

Study quality was assessed across six domains: clarity of the objective and system model; completeness of trust and adversary assumptions; connection between the mechanism and claimed privacy guarantee; support for correctness and operational claims; adequacy and transparency of the evaluation; and reproducibility and evidence maturity. The domains were considered separately rather than collapsed into a single score. Quality assessment was used to qualify the strength of individual claims and their weight in the critical synthesis. Unresolved properties remained coded as *unclear*.

F. Taxonomy Construction, Coding, and Synthesis

The initial taxonomy dimensions were derived from the research questions, recurring concepts in prior surveys, and established properties of privacy-preserving systems. They were then refined through examination of the primary-study corpus. Categories were added, divided, merged, or renamed when needed to represent important distinctions among the included approaches. Each category had to have a clear definition, be supported by study evidence, and contribute to meaningful comparison. Terms with equivalent meanings were normalized.

Coding was performed at the study level using the evidence-card fields described in Section II-E and the definitions in Appendix D. A study could receive multiple labels when it supported several applications, combined mechanisms, protected different information types, or considered more than one adversary. When a coding definition changed, the affected studies were recoded.

The synthesis proceeded at three levels. First, descriptive mapping examined the distribution of studies across the taxonomy dimensions. Second, application-centric synthesis compared studies supporting similar services in terms of authorized output, system and trust model, privacy mechanism, guarantee and residual leakage, operational support, cost, and evidence

maturity. Third, cross-dimensional synthesis examined recurring relationships and tradeoffs across these dimensions.

Research gaps were derived from the coded corpus. They reflected important requirements that were repeatedly absent, weakly supported, evaluated only under narrow conditions, or addressed under restrictive assumptions or unresolved tradeoffs. These findings support the research agenda and requirement-to-protocol-family guidance in Section VIII. The guidance identifies suitable protocol families under stated requirements and assumptions.

G. Reproducibility Controls

The appendices document the complete search strings, operational eligibility criteria, included-study registry, coding definitions, and detailed protocol matrix. Database exports and derived records are preserved as part of the review archive, although their redistribution remains subject to database licensing conditions. Copyrighted full texts are not redistributed.

Bibliographic databases and their indexing records change over time, so a later execution of the same query may not return exactly the same records. Reproducibility in this review therefore refers to preserving the search strategy and snapshot, decision rules, provenance, and supporting evidence used to construct the reported corpus. A subsequent search update should be treated as a new review version and reported separately from the original search.

III. PRIVACY FOUNDATIONS AND UNIFIED ANALYSIS MODEL

Privacy-preserving smart-metering approaches cannot be compared only by the cryptographic or statistical mechanisms that they employ. The meaning of a privacy claim also depends on the system architecture, the information being processed, the entities that receive it, the output required by the application, and the adversary against which protection is expected. Two approaches may use the same mechanism while providing different guarantees because they operate under different trust assumptions or release different outputs. Conversely, technically different mechanisms may address the same privacy objective under comparable system conditions.

This section establishes a common foundation for analyzing these differences. It describes the AMI architecture and information flows, privacy leakage, security and privacy objectives, authorized outputs, trust and adversary models, and the main families of privacy-preserving mechanisms. This unified model provides the terminology required for the taxonomy in Section V and the protocol-level synthesis in Section VI.

A. AMI Architecture, Data, and Information Flows

Advanced metering infrastructure (AMI) is a cyber-physical system that connects customer-side devices, communication networks, utility platforms, and external service providers. It supports billing, grid monitoring, demand response, forecasting, market services, customer feedback, and other data-driven applications. The NIST smart-grid conceptual model describes these functions through interacting domains and communication pathways [16]. Because actual implementations differ

among regions, utilities, and applications, we use a logical role-based architecture rather than assuming one fixed physical topology.

The reference architecture contains four logical zones. The *customer-premises zone* includes smart meters, appliances, home or building energy-management systems, local storage, distributed energy resources, and electric vehicles. The *field and aggregation zone* contains gateways, data concentrators, aggregators, and edge or fog nodes that relay, validate, pre-process, or aggregate metering data. The *utility and service zone* includes the head-end system (HES), meter data management system (MDMS), billing platforms, operational control systems, and analytical services. The *external-service zone* represents electricity markets, distribution and transmission system operators, cloud platforms, and authorized third parties. Not every study contains all these entities, and a single entity may perform several logical roles.

AMI communication is commonly described through home or building area networks, neighborhood or field area networks, and wide area networks. These network labels indicate communication scope, but they do not determine whether an entity is trusted. Trust boundaries must therefore be derived from the information available to each entity and its permitted behavior. This treatment is consistent with NISTIR 7628 [17] and prior AMI security and privacy analyses [18], [19].

Figure 1 summarizes the logical entities and information flows used throughout this survey.

Four flow classes are relevant to the privacy analysis. First, *upstream flows* carry individual readings, local aggregates, events, bids, charging records, alarms, or learning updates toward aggregators and utility systems. Second, *downstream flows* carry tariffs, demand-response commands, control signals, credentials, software updates, and customer feedback. Third, *service and market flows* deliver derived results to authorized recipients. These results may include an exact bill, regional aggregate, load forecast, anomaly alert, market decision, trained model, or privatized dataset. Finally, *physical power flows* determine the net load observable at the grid connection. Storage, local generation, flexible appliances, and scheduling can alter this profile before it is measured, without modifying the reported value after measurement.

For the subsequent analysis, entities are coded according to what they receive, learn, compute, and release. Service inputs and outputs are identified separately, since the required function, sampling frequency, attribution level, and recipient influence the appropriate privacy design [20]–[22].

B. Privacy Leakage and Inference Risks

Smart-meter privacy risk arises because metering data form a time-dependent record of energy use rather than a set of independent measurements. An individual reading may disclose little by itself, but a sequence of readings can reveal recurring activities and relationships between energy use and time. Greater sampling frequency generally increases the operational value of the data, but it may also improve an observer's ability to identify appliances and infer customer behavior [4], [5], [23].




Fig. 1. Unified role- and information-flow model for privacy-preserving smart metering. Solid arrows represent upstream metering data, derived data, and authorized outputs, whereas the dashed arrow represents downstream tariffs, commands, credentials, and feedback. Yellow bidirectional arrows represent physical power flow.

For this survey, *privacy leakage* means information about a protected person, device, activity, or measurement that an entity can derive beyond the output authorized for its role. The entity's view may contain message contents, timestamps, identifiers, communication patterns, intermediate values, released results, public system parameters, and auxiliary knowledge.

We distinguish three related forms of leakage. First, *content leakage* concerns the direct exposure of individual readings or insufficiently protected derived values. Second, *association leakage* occurs when a report, session, location, or output can be linked to a customer or device. Third, *inferential leakage* occurs when an observer combines available data with background knowledge to estimate information that was not explicitly released. These forms can occur together.

A major source of inferential leakage is non-intrusive load monitoring (NILM). NILM methods analyze changes and patterns in aggregate electricity demand to estimate the operating states of individual appliances. The resulting information may support further inference about occupancy, daily routines, or other household activities [6], [24]. The success of such inference depends on sampling frequency, measurement resolution, appliance characteristics, observation duration, the inference method, and the adversary's prior knowledge. Therefore, a reduction in data granularity may reduce a particular attack's accuracy without establishing a general privacy guarantee.

Removing a customer's name does not necessarily make a sequence anonymous if reports remain linkable through stable pseudonyms, network addresses, timestamps, locations, or repeated behavioral patterns. This concern is especially important in vehicle-to-grid environments, where charging sessions and locations may allow an observer to follow a mobile participant across service points [21], [25].

Privacy leakage may also arise from correctly produced and intentionally released outputs. Exact aggregates can reveal information about small groups, and differences between overlapping groups or consecutive reporting periods may expose individual contributions. Repeated noisy releases may accumu-

late information, while analytical models, alerts, and service decisions may reveal properties of their inputs [18], [20], [26].

Reported privacy measures and guarantees must be interpreted according to the property they describe. Signal distortion, mutual information, attack accuracy, and differential privacy address different aspects of disclosure and are not interchangeable [24], [26].

C. Security, Privacy, and Authorized Outputs

Security and privacy are related but not interchangeable. Confidentiality prevents unauthorized entities from reading protected data, while integrity, authentication, access control, and availability protect the correctness, origin, permitted use, and accessibility of system resources. These properties are necessary for secure AMI operation, but they do not alone prevent privacy leakage [17], [19]. For example, an encrypted and authenticated channel may securely deliver fine-grained readings to an authorized utility, while still allowing that utility to infer household activities. A privacy-preserving computation may also be unreliable if an attacker can alter its inputs or impersonate a participant.

A meaningful privacy claim must identify the protected information, the individual or group to whom it relates, the entity from which it is protected, and the result that the entity is permitted to obtain. We call this result the *authorized output*. Privacy does not require a system to reveal no information; instead, it limits disclosure beyond what is needed for the intended service. This function-oriented view is central to established smart-meter privacy models [18], [27].

Authorized outputs differ among applications. Billing may require an exact charge and evidence that it was computed correctly, but not necessarily the complete interval-level load profile. Grid operation may require a regional aggregate, state estimate, or demand-response result rather than identifiable household readings. Analytics may require a prediction, trained model, statistic, or service decision, while market functions may require an accepted bid or clearing result. Thus,

the same disclosure may be necessary for one service but excessive for another [20]–[22].

For each study, we therefore compare the output required by the application with the information actually revealed to every participating entity. This includes exposed intermediate values, identifiers, metadata, and additional results that may be derived from repeated observations. Any disclosure beyond the stated authorized output forms part of the residual-leakage assessment. The mechanisms used to protect different disclosure channels are introduced in Section III-E.

D. Trust, Adversary, and Collusion Models

A system model identifies the participating entities and their communication paths, whereas a trust model states which entities are expected to follow the protocol and which information they are permitted to access. An adversary model specifies how these expectations may be violated. A privacy guarantee must therefore be interpreted under the stated trust assumptions and adversary capabilities.

A *trusted* entity is assumed to perform its assigned operations correctly and not misuse the information it receives. An *honest-but-curious*, or semi-honest, entity follows the protocol but may analyze its legitimate view to infer additional information. A *malicious* entity may deviate from the protocol, select dishonest inputs, manipulate messages, or cooperate with other entities. The term *semi-trusted* is used inconsistently in the literature and is therefore accepted only when the study defines its permitted behavior and observations [10], [18]. Trust is an assumption of the analysis, not evidence that an entity cannot be compromised.

Adversaries may be external or internal, passive or active, and static or adaptive. An external adversary may observe network traffic or inject, replay, modify, delay, and suppress messages. An internal adversary controls or compromises a legitimate participant and may obtain its keys, stored data, protocol state, and received messages. Physical or credential exposure may reveal a meter's secrets without giving the adversary complete control over its execution; this case differs from active corruption, in which the adversary can also choose inputs and alter protocol behavior. Static corruption fixes the compromised entities before execution, whereas adaptive corruption allows the adversary to select them during or between protocol runs [24], [28].

Collusion occurs when two or more entities combine their observations, keys, intermediate values, or protocol states. A protocol that protects a reading from an aggregator alone may fail when the aggregator cooperates with the utility, a key authority, or compromised meters. A collusion-resistance claim must therefore identify the permitted coalition, the information shared, the maximum number or fraction of corrupted participants, and any parties that must remain honest. Merely distributing a computation among several entities does not establish collusion resistance [10], [11].

These assumptions and capabilities are recorded for each study. Unspecified capabilities are recorded as *not reported* rather than inferred from the selected privacy mechanism.

E. Privacy-Preserving Mechanisms and Building Blocks

Privacy-preserving smart-metering systems use several technical mechanisms. These mechanisms protect different interfaces and make different assumptions about keys, trusted parties, collusion, and permitted outputs.

Cryptographic computation protects data while a required function is evaluated. Homomorphic encryption permits selected operations, commonly addition, to be performed over ciphertexts and is therefore widely used for meter-data aggregation. Functional encryption allows a key holder to learn a specified function of encrypted readings, while attribute-based encryption restricts decryption to recipients satisfying an access policy. Secret sharing distributes private inputs as shares among parties, while secure multiparty computation supports joint computation over protected inputs. Masking-based aggregation adds values that cancel when the authorized aggregate is computed. Commitments combined with zero-knowledge proofs can demonstrate correctness without disclosing the underlying reading. These mechanisms differ in key ownership, supported functions, verification, failure handling, and resistance to collusion [10], [20], [29].

Statistical protection modifies data or released results to reduce inference. Perturbation, generalization, quantization, and temporal or spatial aggregation reduce the precision available to an observer. Differential privacy provides a formal bound on how much the output distribution can change when an individual's protected contribution changes. In the central model, a trusted curator applies the randomization, whereas in the local model each meter perturbs its data before transmission. The resulting protection depends on the privacy parameters, adjacency definition, sensitivity, and release process. Repeated releases and temporally correlated meter readings require particular care because privacy loss and inference may accumulate over time [24], [26].

Identity- and linkability-protection mechanisms aim to prevent reports or transactions from being associated with a customer. They include changing pseudonyms, anonymous credentials, blind or group signatures, token-based reporting, and anonymous communication channels. Such mechanisms may support authentication without revealing a stable identity, but they do not necessarily hide the reported value. Their effectiveness also depends on observable metadata and recurring consumption patterns [21], [25].

Trusted and hardware-assisted mechanisms place selected privacy or security operations inside a trusted intermediary, tamper-resistant meter, smart card, or protected hardware component. These mechanisms may isolate keys, separate identity from consumption data, or enforce permitted computations. However, they transfer part of the privacy guarantee to the correct operation and compromise resistance of that component. Physically unclonable functions primarily support device identification and key establishment; they do not protect consumption data unless combined with an appropriate privacy mechanism [9], [30].

Blockchain, cloud or edge computing, and federated learning are enabling architectures. Their privacy depends on the mechanisms used and the information recorded, exchanged,

and released.

Customer-side physical mechanisms alter the load observed at the grid connection. Batteries, local renewable generation, flexible appliances, and load scheduling can conceal appliance events without relying solely on communication-layer protection. Their effectiveness is constrained by storage capacity, charging and discharging limits, energy cost, user requirements, and the attack or leakage measure used for evaluation [26], [31].

Hybrid protocols combine encryption or secure aggregation with authentication, output perturbation, anonymity, trusted hardware, or physical load shaping.

IV. EXISTING SURVEYS AND THE REMAINING EVIDENCE GAP

Prior surveys provide the conceptual and technical foundation for the present review, but they differ in scope, organizing principles, and methodological detail. This section examines their development, compares their coverage through a common analytical matrix, and identifies the remaining gap addressed by the present survey.

A. Evolution and Scope of Prior Surveys

The prior-survey corpus contains 23 reviews published between 2013 and the final search cutoff in 2026. These reviews cover applications, trust assumptions, adversaries, privacy mechanisms, physical load shaping, and deployment, but differ in their scope and organizing principles.

The early reviews established several concepts that still shape current work. Erkin et al. [20] linked the metering architecture and trust model to the function computed over protected readings, particularly spatial and temporal aggregation. Zeadally et al. [32] examined privacy across communication layers and compared identity protection, trusted intermediaries, aggregation, billing, and load shaping. Saputro and Akkaya [33] organized AMI solutions according to customer-identity exposure and examined the feasibility of homomorphic aggregation. Finster and Baumgart [21] separated billing from grid operations and compared approaches using sampling frequency, attribution, exactness, and system complexity.

Between 2016 and 2018, the scope expanded toward specialized environments and more explicit threat analysis. Han and Xiao [25] examined identity, location, charging, payment, and participation-related privacy in mobile V2G settings. Zhu et al. [31] focused on rechargeable-battery methods that alter the load visible at the grid connection. Asghar et al. [18] combined application, trust, and privacy perspectives across billing, grid operations, and value-added services, distinguishing cryptographic from statistical privacy. Ferrag et al. [28] organized schemes according to communication and system models and connected them to attacks, countermeasures, protocol phases, and security analysis.

The 2019 and 2020 surveys broadened the comparison in different directions. Desai et al. [30] covered battery load hiding, hardware-based methods, obfuscation, anonymization, and privacy-preserving aggregation. Sultan [22] placed attribution and function at the center of the analysis, covering billing,

operational metering, rewards, and incentive-based demand response. Kumar et al. [19] linked the HAN–NAN–WAN architecture to attacker access, system threats, privacy goals, and countermeasures. Agarkar and Agrawal [34] examined authentication, key management, and privacy-preserving aggregation, including lattice-based designs. Farokhi [26] separated data manipulation from physical demand shaping and compared information-theoretic and statistical privacy measures.

The 2022 reviews connected privacy to further smart-meter services. Ben Romdhane et al. [35] classified studies according to the reason for metering and the required attribution, covering operational, billing, and combined uses. Abreu and Pereira [29] separated changes to physical consumption from changes to reported data and provided a focused treatment of homomorphic encryption. Dong et al. [36] connected smart-meter data analysis and privacy preservation to forecasting and energy management. Abdalzaher et al. [37] covered differential privacy, learning-based methods, aggregation, pseudonyms, and privacy metrics.

Recent reviews have extended this coverage. Kua et al. [24] organized methods around weak and strong adversaries and included communication-side protection and customer-side energy management. Ajiboye et al. [8] reviewed privacy and security across AMI data, components, communication protocols, vulnerabilities, and countermeasures. Bibi et al. [9] adopted a broad software–hardware view, covering cryptographic, statistical, learning-based, and physical methods, together with public datasets. Han et al. [10] focused on privacy-preserving aggregation and classified schemes primarily by their dependence on a trusted authority. Deokar et al. [11] jointly reviewed authentication, key agreement, and aggregation, with attention to multidimensional data, dynamic membership, and meter failure. Hasan et al. [12] placed AMI within a five-layer smart-grid cyber-physical architecture and connected cryptographic mechanisms to attacks, standards, protocol evaluation, and performance criteria.

The following coverage matrix compares these reviews through a common set of analytical and methodological dimensions.

B. Survey Coverage Matrix

To compare the prior surveys consistently, we evaluated them using eight dimensions derived from the research questions and the final taxonomy. Table I reports the resulting coverage matrix. The dimensions, coverage symbols, and interpretation of the present-survey row are defined below the table.

The matrix shows that privacy mechanisms are the most consistently covered dimension. Application and output requirements receive detailed treatment in several surveys, while architecture, trust, and privacy guarantees are treated unevenly. Lifecycle and operational resilience receive less consistent attention, with most reviews providing partial or incidental coverage.

The weakest shared dimensions are evidence assessment and review-method reproducibility. Under the coverage criteria used here, neither dimension receives detailed and structured

TABLE I
COVERAGE OF ANALYTICAL AND METHODOLOGICAL DIMENSIONS IN PRIOR SURVEYS

Survey	Scope		Models and Assurance		Operational Evidence			Review Process
	App./output	Mech.	Arch./trust	Privacy	Lifecycle	Perf./deploy.	Evidence	Method
FOUNDATIONAL REVIEWS: 2013–2015								
Erkin et al. [20] (2013)	D	D	D	P	–	P	–	–
Zeadally et al. [32] (2013)	P	D	P	P	–	P	–	–
Saputro and Akkaya [33] (2014)	D	D	P	P	P	P	–	–
Finster and Baumgart [21] (2015)	D	D	D	P	P	D	–	–
EXPANSION AND SPECIALIZATION: 2016–2020								
Han and Xiao [25] (2016)	D	D	P	D	P	P	–	–
Zhu et al. [31] (2017)	P	D	P	D	P	P	–	–
Asghar et al. [18] (2017)	D	D	D	D	P	P	–	–
Ferrag et al. [28] (2018)	D	D	D	D	P	D	P	P
Desai et al. [30] (2019)	P	D	P	P	P	P	–	–
Sultan [22] (2019)	D	D	P	P	–	P	–	–
Kumar et al. [19] (2019)	D	D	D	D	D	P	P	–
Agarkar and Agrawal [34] (2019)	P	D	P	P	P	P	–	–
Farokhi [26] (2020)	P	D	P	D	P	D	P	–
CONTEMPORARY REVIEWS: 2022–2026								
Ben Romdhane et al. [35] (2022)	D	D	P	P	P	P	–	–
Abreu and Pereira [29] (2022)	P	D	P	P	P	P	–	–
Dong et al. [36] (2022)	D	P	–	P	–	P	–	–
Abdalzaher et al. [37] (2022)	P	D	P	P	–	P	–	–
Kua et al. [24] (2023)	D	D	D	D	P	D	–	–
Ajiboye et al. [8] (2024)	P	D	P	P	P	P	–	–
Bibi et al. [9] (2025)	P	D	P	P	P	P	–	–
Han et al. [10] (2025)	P	D	D	P	D	P	–	P
Deokar et al. [11] (2026)	P	D	P	P	P	P	–	–
Hasan et al. [12] (2026)	P	D	D	P	P	D	P	P
Present survey (2010–2026)[†]	D	D	D	D	D	D	D	D

Coverage symbols: D = detailed and structured treatment; P = partial, selective, or inconsistent treatment; – = absent or incidental treatment. The symbols describe coverage rather than correctness or overall publication quality.

Dimensions: App./output = applications, supported functions, and authorized outputs; Mech. = privacy-preserving mechanisms; Arch./trust = architecture, trust, adversary, and collusion models; Privacy = privacy targets, guarantees, and residual leakage; Lifecycle = dynamic participation, fault tolerance, revocation, and recovery; Perf./deploy. = computation, communication, scalability, evaluation, and deployment; Evidence = assessment of proofs, experiments, implementations, and artifacts; Method = reproducibility of literature search, selection, extraction, and coding.

[†]For the present survey, D denotes a dimension explicitly included in the review protocol, evidence-card template, taxonomy, and synthesis framework. It is not an independent quality rating. The 2010–2026 range denotes the primary-study search period rather than the publication year of the survey.

treatment in any of the prior surveys. Performance and deployment receive broader coverage, but their interpretation still depends on the functions, assumptions, and evaluation conditions of the reviewed studies.

These findings identify the dimensions that require a more integrated comparison. The following subsection defines the remaining gap and explains how the present survey addresses it.

C. Remaining Gap and Positioning of This Survey

The comparison in Table I shows the need for an integrated framework that examines the technical, operational, and evidential properties of privacy-preserving smart-metering approaches together. Previous surveys have made important contributions, but they have generally studied different parts of the problem using different organizing principles. These differences reflect the objectives of the individual reviews, but they make a unified comparison difficult.

The remaining gap concerns how these properties are connected. The required service and authorized output must be examined together with the information revealed, trust assumptions, adversary capabilities, and collusion boundary. Lifecycle and resilience properties must also be considered alongside privacy guarantees and computational cost. The

strength of the supporting evidence is equally important, since reported conclusions can hide differences in evidence strength and experimental conditions.

This survey addresses these gaps through a multidimensional analysis applied consistently to the included primary studies. It connects the application and authorized output with the privacy mechanism, system and adversary models, guarantees and residual leakage, operational behavior, resource costs, and evidence maturity. The analysis distinguishes author claims from supported results and properties inferred from the reported design, following the extraction and assessment rules in Section II-E.

Privacy-preserving approaches are designed for different services, architectures, adversaries, and operational constraints. The survey therefore examines cross-dimensional tradeoffs and maps system requirements to suitable protocol families. This analysis supports a multidimensional taxonomy, an application-centered synthesis, an evidence-derived research-gap map, and a requirement-oriented protocol selection framework.

V. MULTIDIMENSIONAL TAXONOMY

To support consistent comparison, we organize the primary studies using the multidimensional taxonomy shown in Fig. 2.

The taxonomy describes what each approach is intended to achieve, how it provides protection, the conditions under which its claims hold, and the evidence supporting those claims. Its dimensions are applied jointly rather than as independent rankings. They provide the common structure used for the application-centric and cross-dimensional analyses in Sections VI and VII.

A. Design Principles and Coding Rules

The taxonomy follows six design principles. First, the application and its authorized output are identified before the privacy mechanism. Second, the mechanism is separated from the resulting privacy guarantee. Third, trust and adversary assumptions are treated as part of each guarantee. Fourth, lifecycle and operational-resilience properties are treated as separate dimensions. Fifth, efficiency is interpreted within its evaluation context. Sixth, technical properties are separated from their supporting evidence.

These principles are applied through study-level, multi-label coding, following the procedure described in Section II-F. Category definitions, permitted values, and the distinction among *not reported*, *unclear*, and *not applicable* are provided in Appendix D.

B. Application and Supported Functionality

The application dimension identifies the service that an approach supports and the result that legitimate entities are permitted to obtain. We code the specific function rather than using a broad label such as *smart-grid privacy*.

For each study s , the functional profile is represented as

$$\Phi_s = (\mathcal{A}_s, \mathcal{D}_s, \mathcal{F}_s, \mathcal{R}_s, \mathcal{O}_s), \quad (1)$$

where $\mathcal{A}_s$ denotes the application family, $\mathcal{D}_s$ the input data, $\mathcal{F}_s$ the supported function, $\mathcal{R}_s$ the authorized recipient, and $\mathcal{O}_s$ the authorized output. Table II summarizes the application families used in this survey.

Inputs and outputs are coded independently of the application label. Authorized outputs are classified as *exact*, *approximate*, *perturbed*, or *decision-based*. We also record their granularity, frequency, and recipient.

C. Privacy and Supporting Mechanisms

The mechanism dimension identifies how an approach limits access to protected information. We distinguish principal privacy mechanisms from supporting security mechanisms and architectural or computational enablers.

For each study s , the mechanism profile is represented as

$$\mathcal{M}_s = (\mathcal{P}_s, \mathcal{S}_s, \mathcal{E}_s), \quad (2)$$

where $\mathcal{P}_s$ contains the principal privacy mechanisms, $\mathcal{S}_s$ contains supporting security mechanisms, and $\mathcal{E}_s$ contains architectural or computational enablers. Multiple codes are assigned when mechanisms are combined. Table III summarizes the privacy mechanism families and their combinations.

Supporting security mechanisms include authentication, key establishment, access control, digital signatures, message-authentication codes, commitments, zero-knowledge proofs,

physically unclonable functions, and verification procedures. A component is coded as a principal privacy mechanism only when it directly protects against disclosure, identity exposure, linkability, or inference under the stated system and adversary model.

Architectural and computational enablers include cloud, fog, and edge computing, blockchain or distributed ledgers, federated learning, and hierarchical aggregation or aggregation trees. For hybrid designs, each component is coded according to its actual role. The resulting guarantee and residual leakage are assessed separately in Section V-E.

D. Architecture, Trust, Adversary, and Collusion Models

A privacy guarantee is valid only within the system and threat model under which it is established. Approaches that release the same output may provide different protection when they rely on different trusted entities, key holders, communication paths, or assumptions about collusion.

For each study s , the system-and-threat profile is represented as

$$\Psi_s = (\text{Ent}_s, \text{Flow}_s, \text{Key}_s, \text{Trust}_s, \text{Adv}_s, \text{Coll}_s), \quad (3)$$

where the components record the participating entities, information flows, key and control authority, trust assumptions, adversary capabilities, and collusion boundaries, respectively. Table IV summarizes the corresponding coding dimensions.

Entity names are normalized according to their actual roles and the information they can access. An aggregator that only relays ciphertexts is therefore distinguished from one that decrypts individual readings. The same entity may also perform different roles during registration, reporting, aggregation, billing, and revocation.

Trust is recorded for each entity and protocol phase, while collusion is recorded as an explicit boundary rather than a binary supported-or-unsupported property. Protection established only against passive observation is not interpreted as protection against active deviation, participant compromise, physical attacks, or inference from an authorized output. Unspecified adversary capabilities and collusion boundaries are recorded as *not reported*.

This profile provides the assumptions needed to interpret privacy guarantees and residual leakage in the following subsection.

E. Privacy Targets, Guarantees, and Residual Leakage

A privacy mechanism is meaningful only when its protection target and guarantee are stated clearly. Approaches that claim to preserve “smart-meter privacy” may protect different information from different entities.

For each study s , we represent its privacy-assurance profile as

$$\Omega_s = (\text{Target}_s, \text{Subject}_s, \text{Adversary}_s, \text{Guarantee}_s, \text{Evidence}_s, \text{Leakage}_s), \quad (4)$$

where Target_s identifies the protected information, Subject_s identifies the person, device, or group to which that information relates, and Adversary_s identifies the entity




Fig. 2. Overview of the multidimensional taxonomy. Design principles and coding rules guide the joint assessment of six analytical dimensions.

TABLE II
APPLICATION AND FUNCTIONALITY TAXONOMY

Application family	Included functions	Representative authorized outputs
Revenue, billing, and market services	Coarse- and fine-grained billing, dynamic tariffs, settlement, incentives, rewards, and billing verification.	Bill, charge, payment result, settlement value, or evidence supporting a billing dispute.
Grid operations and flexibility	Spatial, temporal, multidimensional, and subset aggregation; load monitoring, state estimation, forecasting, grid optimization, demand response, load control, and flexibility coordination.	Aggregate load, state estimate, forecast, optimization or demand-response decision, flexibility value, or operational control signal.
Privacy-aware communication, authentication, and access	Privacy-aware reporting, anonymous or unlinkable authentication, key establishment, authorization, and controlled access to metering data.	Authentication or authorization decision, session key, access token, or privacy-protected report.
Analytics, data sharing, and value-added services	Data queries, protected data publication and sharing, synthetic-data generation, consumer/prosumer profiling, household-characteristic classification, anomaly or theft detection, consumer services, statistical analysis, and privacy-preserving learning.	Sanitized dataset or time series, synthetic dataset, query result, statistic, profile, cluster, class label, alert, prediction, model, or protected model update.
Prosumer, distributed-energy, and e-mobility services	Microgrid operation, distributed-energy coordination, peer-to-peer trading, electric-vehicle charging, and vehicle-to-grid services.	Trading or charging result, aggregate energy value, allocation, schedule, settlement, or control decision.
Customer-side privacy control	Battery load hiding, renewable-assisted protection, flexible-load scheduling, demand shaping, and local data minimization or obfuscation.	Modified net-load profile, local schedule, or reduced and transformed meter report.

from which it is protected. The remaining components record the stated guarantee, the evidence supporting it, and the information that may still be revealed. Table V summarizes the privacy targets used in the taxonomy.

Privacy guarantees are coded according to their stated form. Cryptographic guarantees may provide computational or information-theoretic protection under defined assumptions, while statistical guarantees may bound disclosure through differential privacy or another formal privacy measure. Identity-oriented guarantees may provide anonymity, pseudonymity, or unlinkability within a specified adversary model. Empirical inference resistance is recorded separately from formally established guarantees.

The scope of each guarantee is recorded explicitly. Semantic security of an encryption scheme protects encrypted inputs under its cryptographic assumptions but does not prevent leakage from the authorized output. Differential privacy applies to the defined neighboring datasets, privacy parameters, release mechanism, and composition setting. A pseudonym does not by itself establish unlinkability, since reports may still be linked through metadata or auxiliary information. A proof of one property is therefore not interpreted as proof of every

privacy objective claimed by the system.

Supporting evidence is classified as a stated claim, formal definition and proof, analytical leakage bound, attack-based evaluation, or empirical privacy measurement. A guarantee is recorded as formally supported only when the claimed property, adversary, assumptions, and proof conclusion correspond to one another. Experimental reduction in attack accuracy supports the tested inference setting but does not establish protection against untested adversaries or auxiliary information.

Residual leakage is recorded even when the principal privacy claim is satisfied. This distinction allows the survey to identify both the protection achieved by an approach and the information that remains available under its stated operating conditions.

F. Lifecycle and Operational Resilience

Changes in participation, keys, or system state can prevent output recovery or weaken privacy even when the underlying cryptographic mechanism remains secure. We therefore treat lifecycle and resilience properties as a separate taxonomy dimension.

TABLE III
TAXONOMY OF PRIVACY MECHANISMS AND HYBRID DESIGNS

Mechanism family or design	Representative mechanisms	Principal protective role
Encryption and protected cryptographic computation	Conventional encryption, signcryption, homomorphic, functional, and attribute-based encryption, and proxy re-encryption.	Protects data from unauthorized access or allows an authorized function or access policy to be applied while limiting disclosure of the underlying readings.
Commitment-based protection and privacy-preserving verification	Hiding commitments, homomorphic commitments, and zero-knowledge proofs.	Allows billing, input validity, or computational correctness to be verified without revealing the protected inputs beyond the intended output.
Secret sharing and multiparty computation	Threshold secret sharing, secure multiparty computation, oblivious transfer, and distributed decryption.	Distributes sensitive information or computation while protecting inputs against coalitions within the stated threshold and adversary assumptions.
Masking and secure aggregation	Additive masks, pairwise masks, perturbation cancellation, random splitting, and keyed linear or matrix transformations.	Hides individual contributions or selected data components while allowing an aggregate or another authorized result to be recovered.
Statistical privacy and perturbation	Differential privacy, randomized perturbation, generalization, quantization, data reduction, spectral filtering, temporal permutation, learned stochastic sanitization, and traffic-pattern obfuscation.	Limits information revealed by released values, patterns, or statistical results, usually with an explicit utility tradeoff.
Identity and linkability protection	Pseudonyms, temporary identifiers, anonymous authentication, unlinkable credentials, blind or group signatures, and anonymous communication channels.	Limits the association of readings, transactions, locations, or sessions with a customer or device.
Trusted or hardware-assisted protection	Trusted intermediaries, tamper-resistant meters, smart cards, trusted execution environments, and secure hardware.	Restricts sensitive processing or identity management to an assumed trusted component or protected device.
Customer-side physical protection	Battery load hiding, renewable-assisted masking, flexible-load scheduling, and demand shaping.	Changes the externally observable load so that it reveals less about underlying household activity.
Hybrid designs	Combinations of cryptographic, statistical, identity-based, hardware-assisted, or physical mechanisms.	Combines mechanisms to address several privacy targets or limitations of an individual mechanism.

For each study s , the lifecycle-and-resilience profile is represented as

$$\Upsilon_s = (\text{Correct}_s, \text{Membership}_s, \text{Fault}_s, \text{Revocation}_s, \text{Recovery}_s, \text{Integrity}_s, \text{Availability}_s). \quad (5)$$

The components record normal-case correctness, membership changes, fault and dropout tolerance, revocation, compromise recovery, integrity protection, and availability. Table VI summarizes the properties examined under each dimension.

For approximate or perturbed outputs, the reported error or utility condition is recorded instead of requiring exact recovery. Correctness and privacy are evaluated separately: a dropout-recovery procedure may restore an aggregate while revealing masks, shares, or individual contributions. Threshold recovery may also introduce a new collusion boundary.

Dynamic participation includes joining and leaving beyond initial registration. Revocation is assessed separately because removing an identifier from an access list does not necessarily invalidate previously issued keys or prevent access to later protocol messages.

Integrity mechanisms are examined according to what they verify. Message authentication can establish origin and detect modification, but it does not show that a compromised meter reported a valid measurement. Resistance to external message modification is therefore distinguished from input validity and resistance to malicious participants.

A lifecycle property is recorded as supported only when the study describes the relevant procedure and the conditions under which it operates. General claims of being dynamic, fault tolerant, or resilient are retained as author claims when the corresponding behavior is not specified.

G. Efficiency, Evaluation, and Evidence Maturity

Efficiency claims are meaningful only when the measured resource, responsible entity, protocol phase, and evaluation conditions are stated. Execution time measured on a desktop computer does not directly establish suitability for a resource-constrained meter. We therefore separate resource requirements from the evidence used to evaluate them.

For each study s , we define its efficiency-and-evidence profile as

$$\mathcal{P}_s^{\text{EE}} = (\text{Comp}_s, \text{Comm}_s, \text{Store}_s, \text{Scale}_s, \text{Eval}_s, \text{Reprod}_s, \text{Maturity}_s). \quad (6)$$

The profile records computation, communication, storage, scalability, evaluation design and comparison validity, reproducibility, and evidence maturity. Table VII summarizes the properties coded under these dimensions.

Resource costs are recorded separately for each entity and protocol phase. Offline preprocessing is distinguished from online reporting, and per-participant costs from system-wide costs. Terms such as *lightweight*, *efficient*, and *low overhead* remain author claims when the study does not provide sufficient analysis or measurement.

The largest evaluated configuration describes the tested range; it does not establish behavior beyond that range. Direct numerical comparisons are made only when the approaches provide compatible functions and guarantees under sufficiently similar security levels, hardware, datasets, system sizes, and measurement units. Otherwise, the comparison is qualitative and the differences in conditions are reported explicitly.

Real smart-meter data may improve practical relevance, but dataset realism alone does not prove privacy, security, or

TABLE IV
TAXONOMY OF SYSTEM, TRUST, ADVERSARY, AND COLLUSION MODELS

Dimension	Recorded properties	Principal analytical question
Entities and architecture	Smart meters, gateways, aggregators, utilities, authorities, service providers, cloud or edge nodes, prosumers, and mobile participants; centralized, hierarchical, distributed, peer-to-peer, or customer-side organization.	Which entities generate, receive, transform, store, or release protected information?
Information flow	Meter readings, identities, keys, masks, ciphertexts, aggregates, queries, model updates, control messages, and final outputs.	What information is visible to each entity during every protocol phase?
Key and control authority	Key generation, distribution, storage, decryption, identity mapping, access control, aggregation, and output recovery.	Which entities can enable, prevent, or independently reconstruct the protected result?
Trust assumptions	Fully trusted, honest-but-curious, untrusted, conditionally trusted, or distributed trust; assumptions may differ across entities and phases.	Which actions must an entity perform correctly, and which information is it assumed not to misuse?
Adversary capabilities	Passive observation, active message manipulation, protocol deviation, insider access, entity compromise, physical access, static or adaptive compromise, and inference from legitimate outputs or auxiliary information.	What can the adversary observe, modify, control, or infer, and when can compromise occur?
Collusion boundary	Colluding entity types, maximum coalition size, threshold assumptions, parties required to remain honest, and excluded coalitions.	Against which coalitions is privacy claimed, and which are excluded?

TABLE V
TAXONOMY OF PRIVACY TARGETS AND REPRESENTATIVE RESIDUAL LEAKAGE

Privacy target	Protected information	Representative residual leakage
Measurements and intermediate data	Individual or multidimensional readings, partial aggregates, masks, ciphertexts, and temporary protocol values.	Exact aggregates, repeated-query differences, compromised keys, or values recoverable through collusion.
Identity and linkability	The association among a customer, meter, reading, transaction, location, or protocol session.	Stable pseudonyms, timing patterns, network identifiers, location information, or identity mappings held by another entity.
Behavior and inference	Occupancy, appliance use, routines, lifestyle attributes, and other properties inferred from temporal consumption patterns.	Long-term trends, correlated measurements, auxiliary information, or inference from authorized outputs.
Billing, charging, and market information	Consumption details, tariffs, bills, payments, charging sessions, trading values, bids, and settlement information.	Information implied by an exact bill, price, schedule, transaction history, or settlement result.
Analytics and learning information	Query inputs, training records, features, model updates, gradients, predictions, and participation in a dataset.	Information inferred from model parameters, updates, query responses, predictions, or repeated interactions.
Participation and communication metadata	Group membership, reporting activity, communication relationships, event timing, and service participation.	Traffic volume, message timing, public-ledger records, communication topology, or observable joining and leaving events.
Physical load	Underlying household demand and appliance activity observable through the grid-facing net-load profile.	Information remaining because of battery limits, energy constraints, long-term averages, forecast errors, or incomplete load shaping.

deployability. Synthetic data may also be suitable for testing a clearly defined property. Reproducibility depends on sufficient algorithmic, parameter, and execution details, rather than code availability alone.

Evidence maturity is represented by the available evidence types rather than a single quality score. A formal proof can support a claim under a stated model but does not establish practical efficiency. A prototype demonstrates implementability but does not replace a privacy proof, and evaluation on meter-class hardware does not by itself demonstrate utility-scale deployment.

VI. APPLICATION-CENTRIC PROTOCOL SYNTHESIS

The multidimensional taxonomy in Section V defines the common properties used to describe the included primary studies. This section applies those dimensions to the main smart-metering applications. An application-centric synthesis is necessary because the same privacy mechanism may support different functions and authorized outputs, while the same application may be implemented under substantially different architectures, trust assumptions, adversaries, and privacy guarantees.

The synthesis emphasizes recurring design patterns, meaningful differences, supported findings, and unresolved limitations. Representative studies illustrate important mechanisms, tradeoffs, and evidence gaps. Approaches supporting several

TABLE VI
TAXONOMY OF LIFECYCLE AND OPERATIONAL-RESILIENCE PROPERTIES

Dimension	Recorded properties	Principal analytical question
Correctness and output recovery	Correctness definition, required participants, recovery procedure, approximation error, and conditions causing failure.	Can the intended output be recovered accurately under the stated normal conditions?
Dynamic participation	Joining, leaving, temporary absence, group changes, and the resulting key, mask, state, or communication updates.	Can membership change without restarting the system or exposing past or future information?
Fault and dropout tolerance	Offline meters, communication failure, missing reports, aggregator failure, tolerated dropout threshold, and degraded output.	Which failures are tolerated, and what happens to correctness and privacy when they occur?
Revocation and key lifecycle	Registration, key issuance, renewal, replacement, expiration, revocation, and redistribution of affected state.	How is access removed, and which information remains available to a revoked participant?
Compromise and recovery	Compromise detection, isolation, rekeying, state repair, replacement, and protection of earlier and later reporting periods.	Can secure operation be restored after a meter, key, or authority is compromised?
Integrity and malicious input	Message authentication, replay protection, range or consistency checks, input validation, aggregate verification, and poisoning resistance.	Can manipulated reports or protocol deviations corrupt the output or privacy guarantee?
Availability and state consistency	Denial-of-service exposure, single points of failure, synchronization, timeout behavior, rollback, and consistency of distributed state.	Can the service continue or recover without inconsistent security and privacy state?

TABLE VII
TAXONOMY OF EFFICIENCY, EVALUATION, AND EVIDENCE MATURITY

Dimension	Recorded properties	Interpretation
Computation	Entity, protocol phase, operation type, asymptotic complexity, operation counts, execution time, memory, and energy use.	Identifies where computation occurs and whether the reported cost is analytical, measured, or only claimed.
Communication	Message size, messages per entity, number of rounds, total transmitted data, reporting frequency, and communication topology.	Distinguishes local savings from costs transferred to aggregators, authorities, or other participants.
Storage and state	Keys, masks, shares, ciphertexts, identity records, historical data, revocation information, and model or protocol state.	Shows whether retained state grows with participants, data dimensions, or reporting periods.
Scalability	Number of meters, aggregators, regions, data dimensions, recipients, rounds, model size, reporting periods, and membership changes.	Identifies the specific form and tested range of system growth.
Evaluation context	Hardware, software, libraries, security parameters, datasets, workloads, network assumptions, system size, repetitions, and variability.	Defines the conditions under which the reported results were obtained.
Comparison validity	Baselines, supported functions, outputs, threat assumptions, security level, platforms, datasets, units, and excluded costs.	Determines whether numerical comparison among approaches is meaningful.
Reproducibility	Algorithmic detail, parameters, datasets, source code, scripts, configurations, baseline implementations, and execution instructions.	Records which parts of the analysis or experiment can be independently repeated.
Evidence maturity	Complexity analysis, formal proof, simulation, software prototype, real-data experiment, testbed, meter-class hardware, pilot, and operational deployment.	Shows which analytical, experimental, and practical claims have direct support.

applications may appear in more than one subsection, but each publication is counted only once in corpus-level totals. The complete study-level coding is retained in the supplementary protocol matrix in Appendix F. Figure 3 outlines the application families and principal topics examined in this synthesis.

Table VIII summarizes the primary application assignments of the 150 included studies. Each publication contributes once to this descriptive count, while secondary application functions remain recorded in the protocol matrix and discussed in the corresponding subsections. These shares therefore describe primary focus within the included corpus rather than the prevalence of all supported application capabilities. Appendix E

provides the complete subcategory assignments.

A. Billing, Tariffs, Incentives, and Settlement

Billing differs from operational aggregation because its result is normally customer-specific and attributable. A flat tariff may require only a billing-period consumption total, whereas time-of-use and real-time tariffs require a price-weighted sum of interval readings. Block tariffs may further require a piecewise function that depends on cumulative consumption. Incentive and market programs introduce other sensitive inputs and outputs, including baselines, bids, rewards, clearing prices, allocations, and settlement values. Billing privacy must therefore specify which inputs remain hidden,




Fig. 3. Application families and principal topics covered by the application-centric protocol synthesis. Studies may contribute to more than one family.

TABLE VIII
DISTRIBUTION OF THE 150 STUDIES BY EXCLUSIVE PRIMARY APPLICATION ASSIGNMENT

Primary application family	Studies	Share
Billing, tariffs, incentives, and settlement	19	12.7%
Grid operations, aggregation, and flexibility	59	39.3%
Communication, authentication, and access	21	14.0%
Analytics, data sharing, and value-added services	24	16.0%
Prosumer, distributed-energy, and e-mobility services	11	7.3%
Customer-side privacy control	16	10.7%
Total	150	100.0%

which result is authorized, who computes it, and whether its correctness can be verified.

Within the 150-study primary corpus, 34 studies implement or propose a billing-, tariff-, incentive-, or settlement-related technical function. Five additional studies intersect with this application through value-added services, prosumer trading, V2G settlement, or foundational privacy architectures and are examined mainly in their corresponding subsections. The analysis below organizes this evidence by computation and disclosure pattern.

Temporal aggregation and billing separation. The simplest

approach replaces a fine-grained customer trace with a billing-period total. Zhang et al. [38] support spatial aggregation for grid operation and temporal aggregation for individual billing. Chen et al. [39] extend fault-tolerant spatial aggregation to recover a customer's temporal sum over a billing cycle. Naeem et al. [40] combine spatial load monitoring with temporal household aggregation and blockchain-supported audit records. Jo et al. [41] likewise support exact spatial aggregation and per-customer temporal aggregation. Abdallah et al. [42] instead reduce periodic household disclosure through forecast-based, event-driven cluster management, while retaining cluster-level billing and individual billing records at the local aggregation layer. An unweighted temporal total is insufficient for a general time-varying tariff unless price-specific subtotals or further information are released. Short billing periods or overlapping totals may also permit inference about the fine-grained trace.

Other architectures separate attributable billing data from higher-frequency operational data. The lattice-based design of Zhang et al. [43] protects reports during transmission and uses pseudonymous identities, but the control center decrypts individual readings and the power supplier can resolve the billing identity. Its protection therefore concerns communication confidentiality and identity separation rather than consumption privacy from the authorized decryptor. Zhu et al. [44] com-

bine conditional anonymous authentication with fog-assisted aggregation and billing. The control center need not receive customers' real identities during routine processing, but the fog infrastructure and identity-recovery procedure become part of the billing privacy boundary.

Local and delegated verifiable billing. Molina-Markham et al. [45] propose an early architecture in which the meter computes the bill locally and uses zero-knowledge verification to demonstrate consistency with committed readings. Fine-grained operational reports are relayed through a trusted neighborhood gateway without their household identifiers. The billing construction remains an architectural sketch rather than a fully specified cryptographic protocol.

Rial and Danezis [46] show that fine-grained readings need not be disclosed merely to verify a bill. Their generic construction combines meter-certified readings, provider-certified tariff policies, commitments, signatures, and non-interactive zero-knowledge proofs. The customer computes the charge locally and reveals the final fee together with evidence of correct computation. Their optimized construction reduces the cost for public linear tariffs. This approach requires trusted metering evidence, secure storage of certified readings, and verification compatible with the tariff. The final fee, billing period, and number of contributing records may still reveal information.

P2SM by Ni et al. [47] also supports a customer-verifiable price-weighted bill while protecting individual readings from a misbehaving collector. Here, the collector computes encrypted bills from temporarily stored encrypted reports using homomorphic proxy re-encryption and authenticators. The operation center remains a trusted decryption authority. Verifiability therefore supports billing correctness without implying that the readings remain solely at the customer.

Encrypted and function-revealing computation. Homomorphic encryption and functional encryption allow an authorized entity to obtain a billing result without directly decrypting every input. PMBFE by Ibrahim et al. [48] uses inner-product functional encryption to reveal a spatial load sum or a customer's tariff-weighted temporal sum. ETDfE by Ibrahim et al. [49] also supports tariff-weighted billing; its principal theft-detection contribution is examined in Section VI-D. PPSO by Xue et al. [50] separates an outsourced service that receives an exact demand aggregate from a trusted utility that computes individual customer bills.

Other schemes broaden the supported functions. MDMS by Alsharif et al. [51] combines multidimensional and multisubset aggregation with delegated dynamic billing through split-vector/matrix transformations and canceling masks. Cancellation across a billing interval also creates a boundary: utility-gateway collusion may reveal a meter's interval aggregate vector beyond its scalar bill. Zhao et al. [52] use lattice-based homomorphic computation for billing and statistical functions. Marandi et al. [53] examine encrypted analytics covering billing, statistics, and classification. Liu et al. [54] reuse protected reports for tariff-based billing and regional leakage detection, while Zeng et al. [55] support aggregate cost, individual cost, load monitoring, data queries, and theft detection.

PPMA by Li et al. [56] proposes a billing extension, although its printed cancellation step leaves a residual mask and does not generally recover the intended temporal total. Liu et al. [57] support controlled arithmetic queries, including billing-oriented computations, over stored encrypted reports. MDA-FLH by Chen et al. [58] also supports aggregate tariff statistics rather than a customer-specific bill; its multidimensional construction is examined in Section VI-B.

Function-revealing and query-oriented constructions provide greater control than unrestricted decryption, but privacy depends on which keys or queries are authorized, who generates them, and which parties may collude. Enough linearly independent outputs over the same readings may recover the original reading vector. Bills, aggregates, intermediate projections, and repeated query results therefore remain part of the privacy analysis.

Masking, obfuscation, and lightweight computation. Masking-based approaches reduce public-key operations but introduce different trust and lifecycle requirements. EDAS by Gope and Sikdar [59] hides readings from the intermediate aggregator using one-time masks and temporary identities. Its trusted service provider nevertheless holds the cancellation information and reconstructs individual readings for billing. EPIC by Alsharif et al. [60] combines canceling masks with end-to-end integrity mechanisms and supports dynamic billing without revealing every interval reading to the collector. Tonyali et al. [61] use canceling obfuscation values to support state estimation and recover an exact billing-period total.

Mask cancellation may depend on complete participation, correction information for missing meters, or a trusted entity that knows the masks. Privacy from an intermediary does not imply privacy from the party holding the cancellation information.

Perturbation with billing preservation or correction. Perturbation-based approaches hide the fine-grained profile while retaining sufficient billing accuracy. Barbosa et al. [62] apply local Laplace perturbation to reduce appliance-usage inference while preserving approximate temporal and spatial utility. Zheng et al. [63] combine distributed noise with time-series permutation, although the relationship between sanitized reports and exact charging requires careful interpretation. Hassan et al. [64] perturb real-time readings in a renewable-energy setting while retaining longer-term totals for billing. Gough et al. [65] release locally perturbed readings and propose periodic correction information for billing. Tran et al. [66] separate perturbed high-frequency cloud publication from an encrypted low-frequency channel carrying the exact accumulated billing total. The correction or exact-billing channel remains a separate privacy boundary.

Zhao et al. [67] combine masking for exact individual billing and theft-detection projections with distributed differential privacy for regional aggregation. Zaredar et al. [68] instead constrain tariff-weighted noise to cancel in the final charge, allowing the utility to compute an exact bill from perturbed interval reports. The construction also supports defined forms of post-report tariff adjustment. Its privacy evidence concerns reconstruction resistance under the evaluated

conditions rather than a general differential-privacy guarantee.

Accurate billing and formal inference privacy are different objectives. Zero-sum or tariff-weighted noise cancellation preserves a chosen billing function, but dependence among noise values may reveal additional information. Differential-privacy noise must be calibrated for the released output; billing from perturbed readings may require approximation, later correction, or a separate exact-billing channel. Repeated releases and alternative tariff evaluations can add further constraints on the protected readings.

Dynamic pricing, incentives, and settlement. Dynamic-pricing protocols may protect the customer's reading, the applicable price, or both. UDP by Liang et al. [69] allows a community gateway to recover exact community demand and privately evaluate information related to an individualized dynamic price. The trusted utility still obtains the information required for final billing.

Incentive-based demand response links metering history, baseline calculation, curtailment, and reward. Gong et al. [70] allow the provider to maintain a longitudinal profile under a pseudonym and later support reward redemption. The provider receives fine-grained pseudonymous readings; the principal guarantee protects the identity-profile link rather than the profile from the provider. Chim et al. [71] hide individual advance power-use commitments during aggregation but reveal a participant's identity and plan when billing reconciliation is required.

Balli et al. [72] use distributed ElGamal operations and zero-knowledge proofs to compute a multi-unit demand-response auction without a trusted bidding intermediary. Each customer learns its own outcome, whereas the utility obtains the aggregate demand vector, clearing information, and the outcomes required to identify winners. Exact market outputs may still expose an individual bid when the participant set is small, participation changes across rounds, or most participants collude. Related value-added-service incentives and prosumer and V2G settlement mechanisms are considered in Sections VI-D and VI-E.

Table IX summarizes the principal design patterns and selected representatives.

Support for post-report tariff revision, block tariffs, billing disputes, refunds, malicious meter inputs, meter failure, membership changes, revocation, key compromise, and repeated billing queries remains uneven. Evaluation ranges from analytical cost estimates to desktop implementations and limited prototypes, with few end-to-end assessments of the complete metering, computation, verification, and dispute process. Billing suitability therefore depends on the required tariff, authorized disclosure, trust assumptions, and evidence supporting the complete process.

B. Grid Operations, Aggregation, and Flexibility

Operational smart-metering services require timely information about demand, load distribution, system state, and available flexibility. The authorized output may be an exact regional sum, a noisy estimate, a multidimensional aggregate, a subset statistic, a forecast input, or a demand-response decision.

Their leakage depends on reporting frequency, aggregation groups, released dimensions, participant sets, and the ability to combine results across time or queries.

Exact spatial and temporal aggregation. Early schemes commonly authorize an operation center to learn the sum of readings collected from a fixed residential area. Li et al. [73] use Paillier encryption and an aggregation tree so intermediate meters combine ciphertexts without observing the underlying readings. Kursawe et al. [74] develop several canceling-mask constructions, including a lightweight pairwise-key design, for aggregate release or comparison with an independent feeder measurement. Fan et al. [75] protect individual reports against internal participants while releasing an exact neighborhood sum. Busom et al. [76] distribute decryption across the participants to reduce reliance on a single decryption authority.

Exact group totals support regional monitoring and demand-supply balancing, but may permit inference through overlapping groups, repeated observations, or subtraction of known participants' values. Small groups and changing membership increase this risk. Zareidar et al. [77] consider exact operational aggregation under explicit aggregator-utility collusion and bounded exposure of meter state, making the remaining unknown mask components and corruption boundary part of the guarantee.

Operational utility is not always reducible to a sum. Tonyali et al. [61] construct perturbations in a subspace that preserves the output of weighted-least-squares state estimation, while the perturbations cancel over the billing period. Correctness and privacy depend on the measurement model, the obfuscation subspace, and trusted distribution of the masking values.

Multidimensional, multisubset, and query-oriented aggregation. EPPA by Lu et al. [78] packs multiple measurement dimensions into one Paillier plaintext and recovers an exact sum for each dimension. MuDA by Chen et al. [79] extends protected computation to averages, variances, and one-way-ANOVA-related quantities, with an enhanced variant that perturbs the released statistics. Shen et al. [80] organize exact per-area and per-dimension aggregates into a cube structure supporting selected-area totals and averages. Mohammadali et al. [81] combine multidimensional aggregation with meter-failure handling, while Li et al. [82] recover per-dimension sums and sum-of-squares information in a fog-assisted architecture. MDA-FLH by Chen et al. [58] extends multidimensional aggregation by allowing distinct linear coefficients for different users and measurement dimensions. The control center recovers weighted results for each dimension rather than only unweighted sums. Correct recovery depends on the encoding bounds and cancellation of the aggregation masks.

Multisubset schemes classify readings into predefined consumption ranges and release the number of customers and aggregate consumption in each range. MDMS by Alsharif et al. [51] produces these statistics within multiple data dimensions. Wang et al. [83] add fault handling and membership operations. Li et al. [84] support aggregate queries over predefined consumption intervals together with a completeness witness. Li et al. [85] release an aggregate only when a group reaches a defined size threshold and support controlled sharing with multiple grid departments.

TABLE IX
PRINCIPAL DESIGN PATTERNS FOR PRIVACY-PRESERVING BILLING, INCENTIVES, AND SETTLEMENT

Design pattern	Authorized result	Principal privacy boundary	Selected representatives
Temporal aggregation or identity separation	Billing-period total or attributable low-frequency record	Hides the fine-grained trace or its identity, but may not support a general time-varying tariff.	Zhang et al. [38]; Chen et al. [39]; Zhang et al. [43]
Local or delegated verifiable billing	Exact bill and evidence of correctness	Readings remain local or encrypted during delegated computation; privacy depends on trusted metering or decryption entities and tariff-compatible verification.	Rial and Danezis [46]; Ni et al. [47]
Encrypted or functional computation	Tariff-weighted sum, aggregate cost, or another authorized function	Privacy depends on key issuance, function authorization, collusion, and leakage from repeated exact outputs.	Ibrahim et al. [48], [49]; Alsharif et al. [51]; Zhao et al. [52]
Masking and obfuscation	Exact bill, temporal total, or protected aggregate	The holder of cancellation information may recover individual values; failure handling may require additional disclosure.	Gope and Sikdar [59]; Alsharif et al. [60]; Tonyali et al. [61]
Perturbation with correction	Approximate bill, exact corrected bill, or exact bill under a cancellation constraint	Billing utility depends on the noise structure, correction channel, repeated observations, and supported tariff changes.	Barbosa et al. [62]; Zhao et al. [67]; Zaregar et al. [68]
Private pricing, incentives, and settlement	Price indication, reward, clearing result, allocation, or settlement value	Identity, profile, bid, and transaction privacy are distinct; market outputs can permit linkage or inference.	Liang et al. [69]; Gong et al. [70]; Balli et al. [72]

These richer outputs improve load characterization but introduce additional leakage. A count of one converts a subset sum into an individual value, and repeated interval queries may reveal narrower groups. Query authorization and minimum group sizes must therefore be maintained throughout the protocol lifecycle. Packing multiple dimensions also requires explicit bounds to prevent coefficient overlap, modular wraparound, and incorrect decoding.

Noisy and differentially private operational outputs. Acs and Castelluccia [86] distribute noise generation across meters so the supplier obtains a differentially private aggregate rather than individual readings. Bao et al. [87] combine homomorphic aggregation, distributed differential privacy, and meter-failure handling. Their basic construction releases the exact sum of functioning meters, whereas the enhanced construction releases a Laplace-perturbed sum. Shi et al. [88] combine private stream aggregation, distributed noise, diverse groupings, and failure localization. DiPrism by Ni et al. [89] adds range-based filtering, dropout compensation, and distributed differences of Gamma noise shares that combine into a Laplace-distributed perturbation.

The operational meaning of these outputs depends on the privacy parameters and grid task. Noise suitable for regional monitoring may obscure rapid local events, and acceptable error depends on group size and temporal resolution. Repeated releases over time or overlapping groups require composition accounting when differential privacy is claimed. Failure handling must preserve both the intended noise distribution and the meaning of the reported population. Replacing missing reports with stale readings, omitting them, and estimating them produce different operational outputs.

Failure handling and dynamic participation. Many aggregation schemes rely on cancellation relations formed over a known participant set. A missing report can prevent decryption or leave an uncanceled mask. PDAFT by Chen et al. [39] distributes decryption among several control-center servers and uses correction material for absent meters. Xue et al. [90] use

pairwise masks, evolving keys, and secret sharing to support failure recovery and membership changes without a separate trusted authority. FPDA by Wu et al. [91] uses threshold-assisted reconstruction to address an earlier scheme whose recovery material could expose a target's mask. DUDA by Liu et al. [92] supports aggregation over a changing active population without complete system reinitialization.

Fault recovery is itself privacy-sensitive. The capability to reconstruct a missing meter's cancellation value may also expose an active target unless recovery requests are constrained and failures verified. Exact aggregates over the same readings from participant sets differing by one meter reveal that meter's value by subtraction. Recovery procedures must therefore specify who authorizes reconstruction, which population the output represents, and whether a minimum aggregation-set size is preserved.

Integrity, verifiability, and malicious inputs. Encryption does not establish that a report is fresh, correctly formed, or based on a truthful measurement. EPIC by Alsharif et al. [60] combines canceling masks with source authentication and end-to-end aggregate integrity. Zhang et al. [93] consider deletion and tampering by an intermediate aggregator, adding verification intended to prevent a modified or incomplete fog-level aggregate from being accepted. Pang et al. [94] use two-party secret sharing and private range checking to reject values outside an accepted interval before aggregation. Li et al. [84] address completeness of outsourced interval queries, and Yu et al. [95] emphasize lightweight verification of the final aggregate.

These properties address different failures. A signature authenticates report origin but does not prove physical measurement correctness. A range check rejects out-of-range values but may accept a plausible false reading. Aggregate verification can detect selected forms of modification or omission without establishing that the original measurements were honest. Communication integrity, aggregation correctness, completeness, input validity, and physical measurement

trustworthiness are therefore assessed separately.

Demand response and flexibility coordination. Demand response links metering information to an operational action. EPPDR by Li et al. [96] releases building-level and system-wide demand sums and returns a response coefficient to participating users. EDAS by Gope and Sikdar [59] provides an authorized group aggregate for demand-response management, while PPSO by Xue et al. [50] allows an outsourced service to obtain the exact aggregate needed for real-time price prediction.

Other programs require customer-specific information. Gong et al. [70] maintain pseudonymous longitudinal profiles for baseline calculation and reward allocation. Balli et al. [72] protect individual demand-response bids while releasing the aggregate bid vector and customer-specific auction outcomes. Aggregation alone is therefore insufficient when the service requires attribution, verification, or individual settlement.

Table X summarizes the principal operational design patterns and selected representatives.

Across the corpus, a recurring limitation is the separation between cryptographic correctness and operational validity. Many schemes prove or demonstrate that protected inputs produce a stated aggregate under their model. Fewer evaluate whether the released information supports a real grid decision, whether repeated outputs remain private, or how privacy behaves under failures, changing groups, compromised meters, and auxiliary information. Operational suitability therefore requires evidence that the protected output supports the intended service under the stated threat and operating conditions.

C. Privacy-Aware Communication, Authentication, and Access

Authentication and privacy are not equivalent. An identity may be hidden from a network observer while remaining visible to the service provider, and an encrypted channel does not limit what its receiving endpoint learns. Authentication and access-control schemes are therefore examined here when they address a defined privacy property, such as anonymity, pseudonymity, unlinkability, credential confidentiality, or controlled disclosure of metering information.

Pseudonymous and anonymous reporting. Efthymiou and Kalogridis [97] separate attributable low-frequency billing records from high-frequency operational reports sent under a different identity. A trusted escrow retains the correspondence, so privacy depends on protecting this mapping and preventing linkage through communication metadata. Diao et al. [98] use a linkable anonymous credential that permits one valid anonymous report per meter and reporting period. Reports from different periods are intended to be unlinkable, while duplicates within a period can be detected. The collector nevertheless receives plaintext measurements, so the scheme protects their association with an identity rather than their contents. Zhu et al. [44] integrate conditional anonymous authentication with randomized signatures, fog-assisted aggregation, and billing. The fog infrastructure and identity-recovery procedure remain privileged linkage points.

Persistent pseudonyms support services that follow one participant over time. Gong et al. [70], for example, retain

a pseudonymous longitudinal profile for demand-response baseline and reward calculation. Behavioral patterns, network addresses, timing, and later settlement may still permit identity linkage. Pseudonym rotation can reduce long-term linkability but may interrupt services requiring historical continuity.

Reporting patterns can reveal behavior even when the readings are encrypted. STDL by Ibrahim et al. [99] addresses occupancy inference from transmission and silence in change-and-transmit reporting. A learned local controller adds redundant reports to disguise this pattern, while homomorphic aggregation protects reading contents. The defense requires local occupancy information and is evaluated against a defined inference attack rather than supported by a general privacy guarantee.

Anonymous authentication and key establishment. Tsai and Lo [100] use identity-based credentials so a trusted anchor is not contacted during online meter-provider authentication. The provider can recover the meter identity, meaning that anonymity concerns network observers rather than the authenticated provider. Odelu et al. [101] explicitly consider exposure of ephemeral session state in a three-message authenticated key-agreement construction. Abbasinezhadmood and Nikooghadam [102] use self-certified elliptic-curve keys so the trusted authority contributes to registration without learning each participant's complete private key.

Lightweight schemes often replace certificate processing or pairings with hashes, symmetric primitives, or limited elliptic-curve operations. Zhang et al. [103] update encrypted meter state between sessions and derive a fresh session key, but the provider recovers the actual meter identity. PALK by Khan et al. [104] combines password-based local verification, pseudonymous identities, and elliptic-curve key agreement. Sadhukhan et al. [105] authenticate smart-grid gateways using long-term and ephemeral elliptic-curve values, although the transmitted long-term certificate does not establish unlinkability. Yu et al. [106] use pseudo-identities and hash/XOR-based exchanges for demand-response communication. Habib et al. [107] instead use hyperelliptic-curve-based signcryption between a smart appliance and its smart meter.

Baksh et al. [108] combine password, biometric, and device-card factors for meter-server authentication. Despite its aggregation-oriented title, the protocol specifies authentication and session-key establishment rather than a protected smart-meter aggregate. Its anonymity claim also requires qualification: the transmitted identity-masking fields allow an observer to recover the identity under the printed equations. These studies must therefore be interpreted according to the protected observer, exposed credentials and session state, and trust placed in the provider or authority.

Device-assisted and PUF-based authentication. Physical access to smart meters makes stored long-term keys an important attack surface. Gope and Sikdar [109] combine a physically unclonable function (PUF), fuzzy extraction, changing shadow identities, and synchronization-recovery values. PMAKE by Gope [110] combines PUF-derived values with evolving shared state and backup pseudonyms. Rostampour et al. [111] use ECC and a PUF-derived credential in a two-message meter-provider exchange.

TABLE X
PRINCIPAL DESIGN PATTERNS FOR PRIVACY-PRESERVING GRID OPERATION, AGGREGATION, AND FLEXIBILITY

Design pattern	Authorized operational output	Principal privacy or correctness boundary	Selected representatives
Exact or function-preserving computation	Regional or temporal sum, or a preserved operational result such as a state estimate	Exact outputs permit differencing and auxiliary-information inference; function-preserving transformations depend on the operational model and protection of their masking parameters.	Li et al. [73]; Kursawe et al. [74]; Tonyali et al. [61]; Zaredar et al. [77]
Multidimensional or multisubset aggregation	Per-dimension sums, range counts, subset totals, or statistical quantities	Richer outputs increase utility but may reveal small subsets and require correct packing, bounds, and query control.	Lu et al. [78]; Alsharif et al. [51]; Wang et al. [83]; Li et al. [84]
Noisy or differentially private aggregation	Perturbed regional total or approximate operational statistic	Accuracy depends on group size and task; differential-privacy claims additionally require composition accounting for repeated releases.	Acs and Castelluccia [86]; Shi et al. [88]; Ni et al. [89]
Fault-tolerant or dynamic aggregation	Aggregate over responding or currently active meters	Recovery material and changing participant sets can expose masks or enable differencing unless carefully controlled.	Chen et al. [39]; Xue et al. [90]; Wu et al. [91]; Liu et al. [92]
Verifiable or malicious-secure aggregation	Accepted aggregate with integrity, validity, or completeness evidence	Authentication, range validity, completeness, and physical measurement correctness are separate properties.	Alsharif et al. [60]; Pang et al. [94]; Yu et al. [95]
Demand response and flexibility	Demand aggregate, response coefficient, baseline, bid vector, or allocation decision	Aggregate privacy may not protect identity, participation history, individual actions, or settlement information.	Li et al. [96]; Gong et al. [70]; Xue et al. [50]; Balli et al. [72]

A PUF can reduce direct storage of a reusable device secret, but anonymity, forward secrecy, and physical-compromise resistance require additional protocol properties. Security depends on challenge management, helper-data protection, response stability, enrollment assumptions, and available recovery values. Message loss or deliberate blocking can also desynchronize stateful pseudonym updates. Recovery must restore communication without introducing fixed identifiers that allow sessions to be linked.

Access control and multi-recipient disclosure. Authentication determines who is communicating, whereas access control determines what an authenticated entity may obtain. Ruj and Nayak [112] combine Paillier aggregation with decentralized ciphertext-policy attribute-based encryption. A substation decrypts an aggregate and then protects the resulting record under an attribute policy. This controls downstream disclosure but does not hide the aggregate from the substation.

EPDA by Alsharif et al. [113] provides different aggregate views to distribution operators, suppliers, and the transmission operator through recipient-specific transformation and decryption keys. Liu et al. [114] instead protect dimension keys with attribute-based encryption so heterogeneous electricity-data users recover only authorized dimensions of an exact multidimensional aggregate.

Wang et al. [115] distribute trust across multiple authorities and aggregate pseudonymous signatures for report authentication. Their proposed tracing and revocation procedure assigns privileged identity-recovery capabilities to the authorization centers. However, the transmitted pseudonym pairs expose a stable cross-session value, contradicting the claimed unlinkability under the printed construction. Gai et al. [116] combine traceable group signatures, authorization, and audit records in a permissioned-blockchain and edge-computing architecture. Identity protection comes from the signature and authorization mechanisms, while tracing and ledger records

introduce privileged access and potential linkage.

Access policies must remain effective through key issuance, policy updates, revocation, and authority compromise or collusion. Previously distributed keys and ciphertexts also remain relevant. Correct authorization does not prevent inference from an exact output, particularly when aggregates cover small groups or fine-grained dimensions.

Table XI summarizes the principal patterns and selected representatives.

Many studies provide formal or symbolic authentication analysis and cryptographic cost estimates, but fewer evaluate large-scale concurrent authentication, long-term pseudonym management, recovery after message loss, authority compromise, or traffic-analysis resistance. Authentication evidence must be distinguished from evidence supporting anonymity or unlinkability. V2G-specific authentication and location privacy are examined in Section VI-E.

D. Analytics, Data Sharing, and Value-Added Services

Smart-meter analytics includes load forecasting, customer profiling, electricity-theft detection, occupancy and appliance inference, and other value-added services. These applications may release a reusable dataset, a trained model, an individual prediction, or a classification result. Their privacy boundaries are therefore different. Keeping raw readings local may still expose gradients; encrypted inference may reveal the prediction; and a sanitizer that defeats one classifier may remain vulnerable to another. The protected input, intermediate representation, authorized output, recipient, and reuse policy must therefore be identified.

Within the current 150-study primary corpus, 33 studies directly implement a privacy-preserving release, learning, analytical-computation, data-sharing, or value-added-service function. Two additional access-control studies are cross-coded here because they govern downstream disclosure of

TABLE XI
PRINCIPAL PATTERNS FOR PRIVACY-AWARE COMMUNICATION, AUTHENTICATION, AND ACCESS CONTROL

Design pattern	Authorized result	Principal privacy boundary	Selected representatives
Privacy-aware reporting and identity protection	Validated report or longitudinal profile under a protected identifier, or reports with controlled transmission patterns	Identity protection may leave measurement contents visible; reporting-pattern protection is attack-dependent. Escrow, tracing information, traffic metadata, or behavior may still expose identity or household activity.	Efthymiou and Kalogridis [97]; Diao et al. [98]; Gong et al. [70]; Zhu et al. [44]; Ibrahim et al. [99]
Anonymous authentication and key establishment	Authenticated peer and fresh session key	An identity hidden on the public channel may remain visible to the provider, and an encrypted channel does not restrict endpoint access to metering data.	Tsai and Lo [100]; Odelu et al. [101]; Zhang et al. [103]; Yu et al. [106]
PUF-assisted or multifactor authentication	Device authentication and session-key establishment	Privacy depends on enrollment, helper-data protection, challenge management, factor compromise, response stability, synchronization recovery, and the identity-masking construction.	Gope and Sikdar [109]; Gope [110]; Rostampour et al. [111]; Baksh et al. [108]
Multi-recipient and fine-grained access	Recipient-specific aggregate, record, or measurement dimension	Key issuance, authority collusion, policy updates, revocation, and leakage from the authorized exact output remain relevant.	Ruj and Nayak [112]; Alsharif et al. [113]; Liu et al. [114]
Conditional traceability and governance	Pseudonymous message with provision for authorized identity recovery	Tracing and governance authorities have privileged identity access; public pseudonym fields or audit records may still permit long-term linkage.	Gai et al. [116]; Wang et al. [115]

aggregate metering records. The synthesis below organizes this evidence by the artifact released or computed rather than by cryptographic primitive.

Sanitized, minimized, and synthetic data release. Early studies reduce or transform data before disclosure. Eibl and Engel [5] show empirically that coarser temporal resolution suppresses particular appliance-event features, although it does not provide a formal guarantee against every inference method. Sankar et al. [117] formulate release as an information-theoretic privacy–utility problem using mutual-information leakage and reconstruction distortion. Barbosa et al. [62] add local Laplace noise while preserving useful temporal and population-level sums, but do not establish a complete privacy budget for the repeatedly released trajectory.

Later mechanisms combine perturbation with application-specific utility. Hassan et al. [64] release perturbed readings from renewable-energy-integrated homes while retaining separate billing totals; their calibration is differential-privacy inspired rather than a complete longitudinal guarantee. Gough et al. [65] evaluate locally perturbed reports through power flow, billing, and privacy-cost allocation, but specify their sampling and sensitivity assumptions only partially. Zheng et al. [63] distribute noise across meters and shuffle each time series. Their formal differential-privacy basis applies to the aggregate, whereas resistance of individual sanitized reports to non-intrusive load monitoring is primarily empirical.

Learned release mechanisms optimize against an inference adversary. Shateri et al. [118], [119] train causal recurrent releasers using directed information to represent leakage from private temporal attributes such as occupancy. Their demonstrated inference–distortion trade-off depends on the modeled adversary and an entropy surrogate, and the releaser assumes access to the private attribute it conceals. Ou et al. [120] derive local differential-privacy bounds for a singular-spectrum appliance classification measure. However, the released object

is a complete sanitized waveform, and repeated release is not analyzed.

Tran et al. [66] combine independent meter noise with privately assigned canceling central noise to improve regional forecasting utility while separating exact billing from high-frequency publishing. Their experimental noise settings do not consistently satisfy the stated differential-privacy condition. Zaman et al. [7] perturb traces before training a conditional generative model and release synthetic profiles for forecasting and occupancy-related analysis. These outputs inherit differential privacy through post-processing only if the input mechanism provides the stated guarantee. Its calibration and longitudinal accounting are not sufficiently specified to establish that premise.

Federated and distributed learning. Federated learning reduces raw-data transfer, but does not by itself protect updates or outputs. Lin et al. [121] train a federated CNN–LSTM to infer household characteristics across electricity retailers. Raw profiles remain local, but exact gradients reach the parameter server and the model deliberately infers sensitive attributes. Similarly, Manzoor et al. [122] reduce communication for heterogeneous load forecasting through single-layer aggregation, while leaving the selected layer updates visible to the server.

Fernández et al. [123] separately evaluate central differential privacy and secure aggregation for residential forecasting. The former bounds the influence of a client contribution at an accuracy cost; the latter hides individual updates while revealing their exact aggregate. Different experimental configurations prevent a direct claim of superiority. Badr et al. [124] use a functional-encryption-derived construction to aggregate net-energy forecasting models, while FedDetect by Wen et al. [125] combines local perturbation with encrypted aggregation of theft-detection updates. Their protection depends on key authorities, collusion, participation, encoding, and the information retained in the global model.

Zhang et al. [126] combine multi-key approximate homomorphic encryption with a trusted execution environment holding client secrets, but evaluate the proposed smart-grid forecasting framework on image classification rather than load data. He et al. [127] combine hierarchical federated forecasting with masking-derived first-layer features and regional noise. Reused masks expose temporal differences, and authorized linear projections remain visible, limiting trace confidentiality.

Encrypted statistics, queries, and private inference. Several studies protect predefined statistics rather than release reusable datasets. MuDA by Chen et al. [79] uses BGN-based computation for averages, variance-related quantities, and ANOVA-oriented statistics, with a noisy variant intended to resist differencing. EPDAS by Kamil et al. [128] targets subset-specific sums, averages, and variances. Zhao et al. [52] use lattice-based somewhat homomorphic encryption for sums, bills, ANOVA, maximum, and minimum operations, while Marandi et al. [53] use CKKS-based approximate computation for statistics and nonlinear classification. Input protection remains conditional on decryption trust and correct encoding and parameter selection; released results may still reveal sensitive information.

Query-oriented designs permit later selection of data or functions. Liu et al. [57] store encrypted individual reports and support arithmetic queries through homomorphic and interactive operations. The service provider’s universal trapdoor and requester–cloud noncollusion are central assumptions, and overlapping exact results can reveal individual contributions. Li et al. [84] provide exact total and interval-specific queries with a completeness witness, but expose interval membership to infrastructure parties and do not prevent singleton or differencing leakage. Zeng et al. [55] extend functional computation to monitoring, cost queries, individual results, and theft-detection projections; function authorization and mask cancellation under changing participation remain important boundaries.

Electricity-theft detection illustrates why the computing entity matters. Yao et al. [129] encrypt reports in transit, but a trusted detector decrypts individual readings before applying its CNN. Ibrahim et al. [49] issue functional keys for first-layer neural projections, while Zhao et al. [67] release exact mask-corrected first-layer projections. Both restrict access relative to raw decryption, but enough independent projections of the same input with known coefficients may permit reconstruction. SecTCN by Wu et al. [130] protects forecasting inference in an honest-but-curious cloud, while relying on a trusted edge for decryption and selected nonlinear operations; plaintext training remains outside its protection.

Controlled sharing and recipient-specific disclosure. The access-control designs discussed in Section VI-C govern downstream disclosure: Ruj and Nayak [112] apply decentralized attribute policies to aggregate records, while EPDA by Alsharif et al. [113] provides recipient-specific aggregate views. Li et al. [85] authorize sharing after a measurement group reaches a threshold and use threshold-assisted proxy re-encryption. Liu et al. [114] protect dimension keys with attribute-based encryption so recipients recover only authorized aggregate dimensions. These policies restrict access, but

small groups, fine-grained dimensions, repeated outputs, and authority collusion may still permit inference.

Incentivized value-added-service collection. LWIPSM by Zaredar and Amini [131] lets customers join finite reporting programs and submit temporally aggregated, optionally locally perturbed readings for a reward. Blind-signed hash-chain credentials authenticate successive reports, while pseudonyms and an anonymous overlay separate enrollment from reporting. The construction provides per-report local differential privacy when noise is enabled and accounts for session composition. Reports within one program remain intentionally linkable, and protection depends on trusted meters and the adopted anonymity and blind-signature primitives. The evaluation establishes protocol feasibility, but does not quantify downstream model accuracy under different reporting granularities and noise settings.

Table XII summarizes the principal patterns.

Across the corpus, analytical utility is evaluated more thoroughly than privacy under adaptive attack. Forecasting error, classification accuracy, and runtime are commonly reported, whereas reconstruction, membership and property inference, repeated disclosure, and defenses against malicious participants receive uneven attention. Comparisons also depend on datasets, partitions, prediction horizons, model architectures, privacy units, and attacker knowledge. Reported accuracy and efficiency should therefore be interpreted alongside the protected artifact, authorized output, and evaluated adversary.

E. Prosumer, Distributed-Energy, and E-Mobility Services

Distributed generation, shared storage, energy trading, and electric mobility turn customers into producers, sellers, storage participants, or mobile loads. The resulting privacy targets extend beyond household consumption to renewable availability, net-energy balance, bids, trading counterparts, storage schedules, battery state, charging location, and settlement history.

The synthesis examines 12 studies addressing prosumer, distributed-energy, or e-mobility functions. Several also contribute to analytics, authentication, billing, or grid operation. Customer-local load shaping, including the information-theoretic construction of Giacomini et al. [132], is examined separately in Section VI-F.

Renewable metering and prosumer analytics. Hassan et al. [64] combine grid, solar, and wind measurements locally, perturb the real-time combined-energy report, and retain exact longer-term grid and source-specific totals. Their longitudinal privacy is not supported by a formal differential-privacy composition analysis, and the exact totals remain authorized leakage. Verma et al. [133] use homomorphic signcryption and canceling masks to recover an exact regional consumption and production sum. Protection depends on trusted mask generation and infrastructure noncollusion; overlapping aggregates of the same readings can permit participant-set differencing. Algebraic analysis of the reported construction also identifies limitations in its integrity and aggregator–control-office collusion claims.

Badr et al. [124] protect federated updates for short-term net-energy forecasting through a functional-encryption-derived

TABLE XII
PRINCIPAL DESIGN PATTERNS FOR PRIVACY-PRESERVING ANALYTICS, DATA SHARING, AND VALUE-ADDED SERVICES

Design pattern	Authorized artifact or result	Principal privacy boundary	Selected representatives
Sanitized or inference-aware release	Coarsened or transformed time series for later analysis	Protection depends on granularity, leakage definition, modeled adversary, and robustness against alternative or repeated inference.	Eibl and Engel [5]; Sankar et al. [117]; Shateri et al. [118], [119]
Differentially private or synthetic publishing	Perturbed aggregate, reusable trace, or synthetic dataset	Differential-privacy guarantees require a defined privacy unit and adjacency, correct sensitivity and noise calibration, and composition accounting. Synthesis inherits protection only from a valid private input mechanism.	Barbosa et al. [62]; Ou et al. [120]; Zaman et al. [7]
Federated raw-data locality	Shared model trained without centralizing raw records	Individual updates, participation, the global model, and deliberately inferred sensitive attributes may remain visible.	Lin et al. [121]; Manzoor et al. [122]
Protected learning and inference	Protected aggregate update, class, forecast, or prediction	Privacy depends on noise and key authorities, collusion and dropout, trusted inference endpoints, and leakage from projections or final models.	Fernández et al. [123]; Wen et al. [125]; Badr et al. [124]; Wu et al. [130]
Encrypted statistics or function queries	Statistic, interval result, or authorized arithmetic function	Exact outputs and overlapping queries may permit differencing; function authorization, decoding bounds, and decryption trust remain critical.	Chen et al. [79]; Liu et al. [57]; Marandi et al. [53]; Zeng et al. [55]
Controlled sharing and recipient access	Threshold-, attribute-, dimension-, or recipient-authorized aggregate	Correct authorization does not prevent inference from small groups, fine-grained dimensions, repeated outputs, or authority collusion.	Ruj and Nayak [112]; Alsharif et al. [113]; Li et al. [85]; Liu et al. [114]
Incentivized private collection	Reusable privacy-processed dataset and participant reward	Privacy depends on local transformation, anonymous authorization and routing, session linkability, and trusted metering.	Zareadar and Amini [131]

aggregation mechanism. The utility receives the aggregate update and global model rather than individual updates. Protection nevertheless depends on the key-distribution center, collusion and participation assumptions, numeric encoding, and information revealed by the model. Sharma et al. [134] use collaborative autoencoder training and clustering for consumption and rooftop-PV profiling. Raw profiles remain local, but shared models are not protected through differential privacy or secure aggregation, and the recipients of the final clusters are not specified precisely.

Shared-storage coordination. Lee et al. [135] coordinate HVAC demand and shared energy storage through federated deep reinforcement learning. Local building measurements are not directly centralized during training, but the global server eventually receives fully trained local models, infers building energy schedules, and computes the storage-dispatch policy. The construction therefore preserves measurement locality more clearly than it protects models or schedules from the coordinator.

Prosumer trading and governance. Son et al. [136] combine function-hiding inner-product encryption, one-time identifiers, and bid commitments for peer-to-peer trading on a private blockchain. The smart contract learns comparison outcomes and matches buyers and sellers without learning exact bids. After matching, however, the distribution system operator learns participant identities, bid openings, transferred energy, and account updates for settlement; the matched peer also receives the bid opening. Protection assumes tamper-resistant meters holding a shared function-hiding-encryption secret key. Repeated trades and persistent blockchain accounts may still link one-time identifiers.

Gai et al. [116] use group signatures, edge authorization, smart contracts, and a permissioned blockchain for traceable energy-transaction and resource governance. Their principal privacy target is participant identity rather than the energy value. A privileged super-node can authorize and trace participants, while stable addresses and immutable transaction records remain potential linkage points. Blockchain provides auditability; identity protection comes from the group-signature and authorization mechanisms.

V2G authentication, aggregation, and settlement. Abdallah and Shen [137] use changing self-generated pseudonyms and lightweight cryptography for EV charging and discharging. A local aggregator learns the active session's pseudonym, energy request or offer, price information, and session key, whereas the control center retains aggregate power, price, and settlement information. Persistent per-EV records are limited, but network metadata and the aggregator's session view may expose location and participation. EV authorization also partly relies on economic ordering, such as payment before charging, rather than an authority-issued credential.

Liang et al. [138] combine randomized anonymous credentials, PUF-assisted device checks, and additive masking. Charging stations exchange mask shares so an aggregation gateway can recover an exact regional electricity total without directly receiving each station's value. Protection depends on noncollusion and complete mask cancellation; dropout can disrupt recovery, and overlapping aggregates of the same readings may permit differencing. The evaluation emphasizes authentication rather than the complete aggregation subsystem.

Other mobility-related studies principally address identity and authentication. Vallent et al. [139] propose pseudonymous

certificateless signatures, batch verification, and conditional tracing. Their aggregation concerns signatures rather than energy measurements, and persistent public values weaken unlinkability. Yu et al. [140] use password, biometric, smart-card, and PUF-derived information for session-key establishment among the user, charging station, and service provider, without computing a protected energy aggregate or bill. Prateek et al. [141] combine quantum-assisted exchanges with pseudo-identities and encrypted charging-location information. The trusted utility can resolve identities, the charging station obtains the location needed to serve the request, and the evaluation demonstrates ideal superdense-coding circuits rather than an adversarial quantum-network deployment. These designs address observer-dependent identity privacy rather than complete charging-behavior or metering-data protection.

Table XIII summarizes the application patterns and their principal disclosure boundaries.

Evaluation remains fragmented. Trading work includes a constrained-device and private-blockchain prototype, whereas storage and prosumer-analytics studies emphasize prediction, clustering, or control utility. Mobility studies mainly report cryptographic operation costs, symbolic analysis, network simulation, or idealized quantum circuits. End-to-end evaluation connecting authentication, metering, physical energy transfer, aggregation, billing, and dispute resolution remains limited, particularly under collusion, malicious participation, membership changes, and repeated market activity.

F. Customer-Side Privacy Control

Customer-side mechanisms intervene before a fine-grained household trace leaves the premises. They may reduce reporting resolution, transform the numerical report, or physically reshape the demand observed by the meter using a battery, renewable source, or flexible load. Physical load shaping is distinctive: the utility receives the customer's truthful grid exchange after local control rather than an artificially modified meter value. Privacy depends on the control policy, resource constraints, attacker knowledge, and information retained in the net-load trajectory.

The synthesis examines 18 studies addressing customer-side privacy control, including multifunctional studies assigned elsewhere in the primary-focus statistics. Other locally executed perturbation, publishing, and value-added-service mechanisms are discussed in the billing and analytics subsections.

Local data minimization and transformed reporting. These mechanisms change what is reported rather than the physical demand. Eibl and Engel [5] evaluate temporal coarsening through interval averages, maxima, and sampled readings, showing how reporting granularity affects appliance-event detectability. Coarser data reduce behavioral detail but also limit operational timeliness and accuracy. Overlapping summaries or additional high-resolution measurements may weaken the protection.

Sankar et al. [117] formalize release privacy through mutual-information leakage and rate-distortion analysis. Their reconstructed or perturbed load sequence preserves utility through a distortion constraint, while the privacy result depends on the stochastic load model and leakage metric.

Abdallah and Shen [42] instead use customer-side forecasts and event-triggered updates to reduce continuous household reporting. A building-area coordinator manages demand adjustments and billing, but still receives exceptional updates and individual billing information.

Battery-assisted appliance-signature hiding. Early physical approaches use rechargeable storage to suppress appliance transitions in the grid-facing load. Kalogridis et al. [23] use a power router and a best-effort battery policy, while McLaughlin et al. [4] introduce non-intrusive load leveling to reduce the effectiveness of non-intrusive load monitoring. Sun et al. [142] extend household load hiding through Q-learning-based scheduling of battery, EV, and thermal-storage resources.

Reduced load variation or appliance-event visibility does not establish a formal privacy guarantee. Leakage may remain in ramps, battery saturation, timing, long-term energy balance, and correlations with prices, weather, or appliance models. Evaluation against one NILM algorithm or a limited set of traces does not establish resistance to stronger or retrained inference models.

Information- and estimation-theoretic load shaping. Formal approaches optimize a leakage measure rather than visual flatness. Tan et al. [143] study energy harvesting and rechargeable storage using information-theoretic privacy. Giaconni et al. [132] examine renewable generation and finite storage through mutual-information leakage between household demand and visible grid load. Li et al. [144] formulate battery control as a belief-state Markov decision process and derive information-theoretic policies, including a tractable solution for an idealized demand model.

Farokhi and Sandberg [145] instead use Fisher information and the Cramér-Rao relation to derive lower bounds on the error of unbiased demand estimators under regularity assumptions. This differs from mutual information: the latter measures statistical dependence, whereas Fisher information characterizes local estimation precision. Their privacy interpretations therefore depend on different statistical assumptions and observation models.

Storage is not the only physical resource. Chin et al. [146] compare flexible thermal loads with conventional storage for user-demand shaping. Thermal flexibility can reduce reliance on dedicated batteries, but introduces temperature dynamics, comfort constraints, weather dependence, and rebound effects. The reported comparison also assumes accurate physical models and knowledge of demand over the prediction horizon.

Cost-aware and differentially private control. Battery privacy competes with electricity price, capacity, charge/discharge limits, conversion loss, and degradation. Yang et al. [147], [148] formulate online energy-management policies that combine load-shaping objectives with time-of-use cost. Natgunanathan et al. [149] shape demand around a progressive target while exploiting price variation, and Hossain et al. [150] use heuristic optimization to balance mutual-information leakage and electricity cost. Reported improvements depend on tariffs, forecasts, initial battery state, and the privacy metric; battery aging and replacement cost are evaluated unevenly.

TABLE XIII
PRIVACY PATTERNS FOR PROSUMER, DISTRIBUTED-ENERGY, AND E-MOBILITY SERVICES

Application pattern	Authorized result	Principal privacy boundary	Corpus studies
Renewable and prosumer metering	Perturbed real-time trace, exact long-term total, or regional consumption/production sum	Source-specific totals and overlapping exact aggregates may reveal generation availability or individual contributions; masking depends on trust and collusion assumptions.	Hassan et al. [64]; Verma et al. [133]
Net-energy analytics and profiling	Aggregate model update, global forecast model, representation, or behavioral cluster	Local raw data do not prevent leakage from updates, models, latent features, clusters, or underspecified output recipients.	Badr et al. [124]; Sharma et al. [134]
Shared-storage coordination	Building model, inferred demand schedule, and storage-dispatch policy	The coordinator may infer building schedules from shared or complete models even when the original measurements remain local.	Lee et al. [135]
Prosumer trading and governance	Price comparison, trading match, bid opening, settlement, or auditable allocation	Order and match leakage, operator access to exact transactions, privileged tracing, and persistent ledger records permit linkage.	Son et al. [136]; Gai et al. [116]
E-mobility identity, charging, and aggregation	Session key, charging authorization, regional charging total, or settlement record	Location metadata, persistent public identifiers, tracing authorities, exact totals, mask collusion, and participant dropout remain relevant.	Abdallah and Shen [137]; Vallent et al. [139]; Liang et al. [138]; Yu et al. [140]; Prateek et al. [141]

Other studies seek differential privacy through physically feasible battery actions. Zhao et al. [151] combine battery-based load hiding with a contextual control policy, while Zhang et al. [152] use truncated Laplace perturbations for privacy and time-of-use cost management. Liu et al. [153] alternate intermittent perturbation with deep reinforcement learning for battery control. However, their reported load-dependent truncation produces differing output supports for neighboring loads, which does not establish the claimed pure differential-privacy guarantee.

Physical feasibility must be included in the privacy analysis. The released trajectory is constrained by energy conservation, battery state, and power limits. A longitudinal differential-privacy claim therefore requires consistent adjacency, sensitivity, noise calibration, and composition accounting, including the effects of saturation and data-dependent control decisions. A per-slot proof or simulated privacy metric alone does not establish protection of the complete trajectory.

Table XIV summarizes these 18 studies.

Evaluation is dominated by trace-driven simulation, analytical optimization, and leakage or load-flatness metrics. Few studies test several inference attacks, evaluate physical meter and storage hardware, or jointly model degradation, forecast error, conversion loss, user comfort, tariff changes, and adversarial knowledge. Reported privacy values are therefore comparable only when the protected variable, observation horizon, attacker model, resource constraints, and leakage metric are aligned.

VII. CROSS-DIMENSIONAL CRITICAL SYNTHESIS AND EVOLUTION

The preceding synthesis showed that neither applications nor mechanism names alone provide a sufficient basis for comparing privacy-preserving smart-metering systems. This section therefore examines the relationships among mechanisms, authorized outputs, residual leakage, trust assumptions, operational resilience, efficiency, and evidence maturity across the primary-study corpus. It also traces how these relationships

have evolved over time. The complete study-level evidence is retained in Appendix F; the discussion below uses representative studies to explain cross-dimensional findings rather than repeat the application-specific descriptions.

A. Mechanisms, Authorized Outputs, and Residual Leakage

The corpus reveals a recurring separation between protecting the computation and limiting what its result discloses. Mechanisms protect different interfaces, so their guarantees are comparable only after specifying the output, its recipient, its precision and frequency, and the information available from other protocol executions.

Exact computational outputs. Homomorphic encryption, functional encryption, secret sharing, and canceling masks support exact regional sums, tariff-weighted charges, multidimensional statistics, or selected linear functions without routinely revealing every input [48], [51], [73], [74], [94]. Local verifiable billing instead allows the customer to retain the readings and reveal an exact fee with evidence of correct computation [46]. In each case, the authorized result remains visible.

Consequently, input confidentiality should not be interpreted as output privacy. Exact aggregates over the same readings and participant sets that differ by one meter can reveal that meter's contribution through subtraction. Small subsets may likewise expose individual contributions, while sufficiently many independent linear outputs with known coefficients can recover the underlying reading vector. Functional encryption restricts decryption to issued functions, but privacy still depends on which function keys are issued and whether their outputs can be combined. Access control restricts recipients; it does not remove the information contained in an authorized exact result.

Perturbed and transformed outputs. Distributed-noise schemes can provide approximate regional aggregates, whereas local mechanisms release perturbed readings or transformed profiles [62], [86], [89]. Perturbation alone does not establish differential privacy. Where a differential-privacy guarantee is provided, it applies to the stated protected unit,

TABLE XIV
PRINCIPAL PATTERNS FOR CUSTOMER-SIDE SMART-METER PRIVACY CONTROL

Control pattern	Grid-visible result	Principal privacy boundary	Corpus studies
Data minimization or transformed release	Coarsened, perturbed, reconstructed, or event-triggered report	Privacy depends on granularity, distortion, reporting triggers, repeated releases, and auxiliary high-resolution information.	Eibl and Engel [5]; Sankar et al. [117]; Abdallah and Shen [42]
Heuristic appliance-signature hiding	Truthful physical grid draw after storage or flexible-load control	Load flattening or suppression of selected transitions does not establish resistance to stronger NILM models or long-term inference.	Kalogridis et al. [23]; McLaughlin et al. [4]; Sun et al. [142]
Formal leakage-minimizing load shaping	Grid draw optimized under an information- or estimation-theoretic objective and resource constraints	Guarantees depend on the leakage metric, statistical assumptions, attacker observations, resource dynamics, and comfort or energy constraints.	Tan et al. [143]; Farokhi and Sandberg [145]; Giacon et al. [132]; Li et al. [144]; Chin et al. [146]
Cost-aware or DP-oriented battery control	Battery-shaped grid draw under privacy, tariff, and physical control objectives	DP claims require valid calibration and temporal accounting under physical constraints; cost results depend on tariffs, forecasts, battery state, losses, and aging.	Yang et al. [147], [148]; Zhao et al. [151]; Zhang et al. [152]; Hossain et al. [150]; Liu et al. [153]; Natgunanathan et al. [149]

adjacency relation, sensitivity bound, privacy parameters, and release sequence. Repeated releases require composition analysis, while meter failure or selective participation may change the noise distribution and the population represented by the result.

Exact utility obtained through structured perturbation requires separate interpretation. Tariff-weighted or zero-sum noise can preserve a bill or aggregate by introducing dependencies among perturbations [68]. Exact recovery of that function is not itself a differential-privacy guarantee, and the recovered result remains part of the observer’s information. The relevant comparison therefore concerns whether the application requires an exact result, which error is acceptable, and what additional information is exposed through correction or repeated outputs.

Customer-side physical control changes the observation itself. Batteries, renewable generation, and flexible loads reshape the demand measured at the grid connection [132]. The resulting profile may still reveal energy balance, resource saturation, control actions, or correlations with tariffs and weather. Physical shaping can therefore reduce behavioral leakage under resource constraints without establishing independence between the observable load and household demand.

Identity, learning, and decision outputs. Pseudonyms and anonymous credentials primarily protect the association between a report and its source. They do not necessarily hide its value, and persistent pseudonyms may intentionally preserve a longitudinal profile for baseline calculation or settlement [70], [98]. Timing, network identifiers, behavioral regularity, and later account resolution may still restore linkage.

A similar distinction applies to analytics and market services. Federated learning can keep raw readings local while exposing updates, models, predictions, or control decisions [124], [135]. Encrypted trading can hide exact bids from the blockchain while revealing comparison outcomes and matches; matched bid openings are subsequently disclosed to the settlement authority and matched peer [136]. Persistent ledger records may create an additional linkage surface [116]. These disclosures must be included in the privacy analysis even when they are required by the application.

Mechanism composition is useful when its components address complementary disclosure channels. Its effectiveness nevertheless depends on compatible trust assumptions and the information released by the complete protocol. Protocol selection should therefore begin with the minimum result required by the service and specify its recipient, granularity, precision, frequency, retention, and permitted secondary use. These requirements determine which mechanisms are suitable and which residual leakage remains.

B. Trust Assumptions, Guarantees, and Operational Resilience

Across the corpus, privacy guarantees are inseparable from the placement of trust and the conditions under which the protocol operates. A design may remove one intermediary while retaining a key authority, identity resolver, trusted meter, model coordinator, or quorum of non-colluding participants. The relevant questions are which entities receive sensitive information, which actions they must perform honestly, which coalitions are excluded, and whether these assumptions remain valid during failures and lifecycle changes.

Trust placement and collusion boundaries. Architectures producing the same output may expose substantially different information. In EDAS, the aggregator processes masked reports, but the trusted service provider reconstructs individual readings for billing [59]. Functional encryption limits decryption to issued functions, yet the key issuer determines which functions can be evaluated [48]. Threshold and distributed-decryption schemes divide authority among several parties, introducing assumptions about quorum availability, share protection, and tolerated collusion [39], [76].

Privileged roles also remain in pseudonymous settlement, federated learning, and permissioned-ledger governance through identity resolution, key management, model coordination, or tracing [70], [116], [124]. These designs may reduce routine disclosure without eliminating privileged access.

Collusion must therefore be defined by entity type and protocol phase. An aggregator and utility may hold complementary ciphertexts, masks, keys, or intermediate values.

Participants knowing every other contribution can also recover the remaining input from an exact aggregate. Masking and threshold constructions depend on explicit coalition limits [90], whereas other designs distribute unknown mask components to resist an aggregator–utility coalition [77]. A collusion-resistance claim must identify the coalition composition, tolerated size, available protocol state, and excluded combinations.

Adversary capability and guarantee scope. A guarantee against an honest-but-curious participant assumes correct protocol execution. It does not cover fabricated inputs, false failure declarations, state manipulation, or other malicious deviations unless explicitly modeled. Authentication and integrity can detect unauthorized modification of reports [60]; private range checking constrains submitted values [94]; and query-completeness verification detects omitted records [84]. These properties do not establish that a compromised meter reports its true physical consumption: a plausible but false measurement may satisfy all three checks.

Compromise claims also require precise interpretation. FKLm-PDA separates secret components and claims resilience to bounded key leakage under a non-collusion assumption; this claim should not be extended to complete device-state exposure or collusion between holders of complementary components [154]. Likewise, key evolution does not establish forward security if retained recovery information allows earlier protections to be reversed. Proofs and attack evaluations must therefore be interpreted within their stated adversary, leakage, and auxiliary-information models.

Failures, recovery, and dynamic participation. Operational resilience can change the privacy boundary. An absent meter may leave an uncanceled mask or an insufficient decryption quorum. Recovery restores availability through masks, shares, or correction values, but may expose an active target if failure declarations can be abused. PDAFT distributes recovery across control-center servers [39]; Xue et al. support recovery and membership changes without a separate trusted authority [90]; and FPDA uses threshold-assisted dropout recovery [91]. More recent schemes support a changing active population without complete system reinitialization [92].

Correct recovery does not by itself preserve privacy. Changes in participation can create the comparable exact outputs discussed in Section VII-A, while recovery information and previously issued keys may remain available after failure or revocation. Dynamic operation must therefore specify how joining, leaving, temporary absence, and revocation affect cryptographic state, coalition thresholds, and authorized outputs. Authenticated failure reports, justified recovery thresholds, and sufficient uncompromised contributions are necessary where the privacy argument depends on these conditions.

Resilience should not be treated solely as an availability feature when failure handling changes what an entity can observe or reconstruct. A lifecycle privacy claim must specify which operational events it covers and demonstrate that the required trust, collusion, and output restrictions remain valid during those events.

C. Efficiency, Scalability, and Evidence Maturity

Efficiency in privacy-preserving metering is not a single runtime value. Protocols redistribute computation, communication, storage, and trust among meters, gateways, utilities, authorities, and customers. Their costs also depend on the authorized function: an exact sum, a multidimensional statistic, a verifiable bill, and a machine-learning update are not equivalent workloads. Efficiency claims are therefore meaningful only when the function, protected parties, security parameters, system scale, and evaluation platform are stated together.

Cost placement and functional scope. Homomorphic aggregation commonly assigns encryption to meters, ciphertext combination to intermediate nodes, and decryption to the recipient. Li et al. [73] analytically reduce routed packet transmissions and collector-side decryptions through in-network Paillier aggregation, but provide no implementation or meter-level energy measurement. Packing techniques can amortize one ciphertext across several dimensions, as in EPPA [78], provided that input bounds, modulus selection, and decoding are valid. Richer statistical functions introduce additional costs that cannot be compared directly with those of a protocol releasing only one sum [52].

Masking and symmetric constructions can reduce online public-key work at the meter, but introduce key state, synchronization, recovery traffic, or participant-management overhead. Xue et al. [90] report lower normal-reporting computation than Paillier-based alternatives, while communication grows with users and reporting periods. Enrollment, recovery, revocation, and membership change are not benchmarked end to end. A “lightweight meter” claim must therefore be separated from a lightweight system claim.

Scalability and comparison validity. The relevant scaling variables differ across applications. Aggregation and query protocols depend on meters, dimensions, recipients, requested subsets, and reporting or query frequency. Billing depends on tariff intervals and permitted revisions. Federated learning additionally depends on model size, communication rounds, client availability, and dropout, as illustrated by the forecasting scheme of Badr et al. [124].

Many evaluations vary meters or dimensions while leaving network conditions, failures, membership changes, and setup costs unexamined. Asymptotic complexity does not reveal constants, memory pressure, message fragmentation, or device energy consumption; linear runtime over a small tested range does not establish deployment-scale performance. Direct numerical comparisons require comparable functions and threat models, compatible security strength and platforms, and consistent workload and cost accounting.

Evidence maturity and reproducibility. The corpus spans analytical operation counts, desktop benchmarks, trace-driven simulations, device emulation, and prototypes. These forms of evidence answer different questions. Rial and Danezis [46] provide implementation evidence for verifiable billing, including meter modifications and measured customer and provider operations. Zareadar et al. [68] combine ARM-based emulation, a year-long metering dataset, communication-path assumptions, and reconstruction experiments. Such evaluation

extends beyond operation counts, but remains distinct from an operational AMI field trial.

Application utility must also be distinguished from privacy evidence. Lee et al. [135] evaluate shared-storage management and cost outcomes through federated deep reinforcement learning, but do not formally define privacy or test inference attacks against updates and models. Similarly, ideal circuit simulation in the quantum-assisted V2G protocol of Prateek et al. [141] demonstrates the intended operations without establishing resilience under noisy hardware or an adversarial quantum channel. Real datasets and realistic application models do not independently establish privacy.

Favorable benchmarks must likewise be interpreted together with correctness and reproducibility. FKLM-PDA reports low meter-side latency and compact reports under a desktop configuration with a stated 80-bit security level, but incomplete fog-level cost accounting limits performance interpretation. Internal equation inconsistencies and formal leakage details unavailable in the examined material separately limit verification of its correctness and privacy claims [154]. Public code, executable parameters, datasets, test vectors, repeated trials, and uncertainty reporting remain uneven across the corpus.

The principal evaluation gap is therefore end-to-end. Few studies jointly evaluate routine reporting, lifecycle events, failures, malicious inputs, device resources, network conditions, output accuracy, and privacy under attack. Evidence maturity should be judged claim by claim: a protocol may have a formal confidentiality proof but limited deployment evidence, or an application prototype without a formal inference guarantee. Each conclusion should therefore be tied to the evaluated function, actor-specific costs, scale, security parameters, operational conditions, and reproducibility evidence.

D. Evolution and Convergence of the Field

Across the 150 primary studies, the field has not evolved through a simple replacement of older mechanisms by newer ones. Instead, several research branches emerged in parallel and were progressively combined. The following periods organize these overlapping developments rather than define strict boundaries between them.

Foundational diversification (2010–2013). Early work established several enduring design choices. Homomorphic aggregation protected readings from intermediate nodes while releasing an exact regional sum [73]; canceling masks offered an alternative under different participation and key-sharing assumptions [74]; and multidimensional packing expanded aggregation beyond one scalar measurement [78]. Differentially private aggregation accepted controlled output error to limit inference [86]. Other studies moved computation toward the customer through verifiable local billing [46] or reshaped the observable load using renewable generation and storage [143]. These studies already distinguished protection during computation, control of released information, and alteration of the measured signal.

Expansion toward applications and resilience (2014–2018). Subsequent work extended privacy mechanisms beyond static neighborhood aggregation. Demand-response protocols incorporated response delivery and key evolution [96]; anonymous

credentials supported identity-protected reporting [98]; and fault-tolerant schemes introduced failure localization, correction information, or distributed recovery [39], [88]. The application scope broadened to include dynamic pricing, state estimation, incentive calculation, and V2G settlement. EPIC combined masking with report authentication and aggregate integrity [60], while distributed cryptographic operations and zero-knowledge proofs supported private demand-response auctions [72]. Many constructions nevertheless retained fixed participant sets, trusted authorities, or honest-but-curious assumptions.

Architectural and functional integration (2019–2022). Later studies integrated privacy mechanisms with fog, edge, cloud, blockchain, and federated-learning architectures. Multidimensional and multisubset aggregation was combined with delegated billing and several authorized operational views [51]. Robust masking schemes added membership changes, recovery, and evolving state [90], while PUF-assisted authentication reduced reliance on permanently stored device secrets [109]. Functional encryption supported selected inner products for monitoring and billing [48]. In parallel, blockchain-supported trading protected selected bid information while retaining settlement functions [136], and federated learning kept raw measurements local during forecasting or control [124], [135]. These architectures expanded the service boundary while introducing additional trust dependencies and potential leakage through persistent records, model updates, and outputs.

Recent emphasis on richer guarantees (2023–2026). Recent studies explicitly address input validation, query completeness, and dynamic participation through private range checking [94], verification of outsourced aggregates [84], and aggregation over a changing active population [92]. Lattice-based and post-quantum-oriented constructions address encrypted statistics and classification, multidimensional aggregation, and authenticated pseudonymous reporting [43], [53], [155]. Other studies combine identity and consumption protection in mobile charging [138] or support bounded post-report tariff revision while preserving exact billing [68]. These examples extend functionality and the scope of claimed guarantees, but do not by themselves establish mature end-to-end security or deployment readiness.

Three forms of convergence are visible in the included studies: greater attention to authorized outputs, recurring combinations of complementary mechanisms, and explicit treatment of lifecycle events. These developments connect functions that were previously examined separately, but also increase the number of assumptions and disclosure channels that must be evaluated together.

The cross-study evidence suggests that functional breadth has advanced faster than assurance and deployment maturity. Newer schemes are not automatically more private, robust, or practical than earlier focused constructions. Formal models, source-level correctness, realistic parameterization, privacy-attack evaluation, lifecycle experiments, and field-scale validation remain uneven. The resulting research priorities concern explicit output control, compatible trust and collusion assumptions, privacy-preserving lifecycle operation, and reproducible end-to-end evidence.

VIII. RESEARCH AGENDA AND PROTOCOL-SELECTION GUIDELINES

This section converts the review findings into a research agenda grounded in recurrent weaknesses of the primary literature and selection guidance that links protocol families to application, trust, leakage, lifecycle, and deployment requirements. Illustrative scenarios show how these requirements guide the selection and combination of mechanisms.

A. Evidence-Derived Gap Map

Earlier surveys called for lighter computation, less centralized trust, better privacy–utility trade-offs, fault tolerance, and real-world evaluation [9], [10], [18], [20], [21]. These priorities remain relevant. Our analysis of the 150 primary studies identifies a more specific pattern: protection is often established for one component or protocol phase, while the complete construction, released outputs, evolving state, or deployment conditions remain insufficiently examined. We treat a gap as evidence-derived when it recurs across distinct mechanism or application families, persists in later work, and materially affects the claimed privacy service. Table XV summarizes these gaps and the evidence needed to address them.

These gaps do not imply that every study is weak on every dimension. The corpus contains focused advances in verifiable billing, threshold aggregation, differential privacy, recovery, input validation, and query completeness [46], [76], [84], [86], [90], [94]. The unresolved challenge is to integrate the requirements relevant to a service and demonstrate their joint assurance under consistent assumptions. This challenge guides the research priorities and selection rules that follow.

B. High-Priority Research Directions

The gap map identifies six priorities for strengthening assurance, controlling disclosure, and evaluating complete privacy-preserving services.

1) *Specify end-to-end privacy contracts and enforce release control.* Future protocols should translate the privacy contract into enforceable rules for the complete observer view. Exact-output systems need controls on query overlap, group size, function-key issuance, and repeated releases, together with analysis of the inference that remains possible. Differentially private systems should connect their privacy unit, adjacency, sensitivity, and composition rules to the actual release process. Learning systems should distinguish protection of records, client updates, aggregates, and released models [48], [51], [84], [123]. The authorized service and its remaining inference risk should be evaluated within the same model.

2) *Make assurance construction-complete and implementation-aligned.* Future work should carry consistent variables, arithmetic domains, adversary assumptions, and failure conditions from the protocol specification through its proofs, implementation, and experiments. This alignment should cover cancellation, packing, distributed decryption, verification, and recovery rather than only the underlying

primitive [155]–[158]. Automated algebra checks, property-based testing, proof-assistant models for critical subprotocols, and negative test vectors can make the alignment auditable.

3) *Treat coalitions, state exposure, and lifecycle events jointly.* Protocols should be modeled as state machines with explicit conditions for privacy, correctness, and availability across enrollment, reporting, recovery, rekeying, and departure. The model should distinguish passive observation, secret-state exposure, and malicious control, and specify whether corruption is static or adaptive. Evaluation should exercise permitted role combinations and retained recovery state, and distinguish exact recovery from omission or estimation of missing inputs [76], [90], [91], [161]. Results should state the corruption thresholds and privacy, correctness, or error guarantees remaining after each supported transition.

4) *Establish compositional privacy for analytics and emerging services.* Research on forecasting and profiling should evaluate leakage from updates, representations, and final models, alongside poisoning and dropout behavior. Trading and mobility systems should analyze order and match leakage, ledger linkability, settlement disclosure, location, and identity–energy associations. Transmission patterns should also be included when they reveal occupancy or participation [99], [124], [134]–[136], [138]. The resulting argument should explain how component guarantees compose and which observers can combine the remaining views.

5) *Develop cryptographically agile post-quantum migration profiles.* The long service life of metering infrastructure makes migration relevant. Research should specify the security notions, parameter sets, applicable failure probabilities, and plaintext or noise bounds, together with key, ciphertext, memory, and communication costs. Migration profiles should cover authentication, aggregation, revocation, and recovery, including coexistence with deployed cryptographic components [43], [53], [155], [162]. Comparisons should use comparable security strengths and account for migration state and bandwidth overhead.

6) *Build reproducible, service-level evaluation infrastructure.* Benchmarks should connect primitive costs to application utility and protocol behavior under latency, loss, membership change, and malicious inputs. They should separate setup, routine operation, and recovery costs by participant role. Analytics and perturbation experiments should evaluate operational accuracy and privacy-attack success under consistent conditions. Reusable artifacts should include code, parameters, dataset provenance, hardware and network settings, random seeds, and uncertainty across repeated runs [53], [94], [123], [126].

C. Requirement-to-Protocol-Family Mapping

Protocol selection should be requirement-driven rather than label-driven. A candidate must provide the required output, protect the specified information from the relevant observers, and remain valid under the permitted coalitions and lifecycle events. Table XVI maps common requirement profiles to candidate families and the conditions governing their use.

The table identifies candidate families rather than complete protocols. Integrity, freshness, authorization, malicious-input

TABLE XV
EVIDENCE-DERIVED RESEARCH GAPS AND EVIDENCE NEEDED TO ADDRESS THEM

Unresolved gap	Evidence from the primary corpus	Evidence needed to address the gap
Construction-level assurance	Several constructions contain inconsistencies or incomplete assurance for the equations, key flow, arithmetic domains, or parameter bounds actually used [155]–[159].	An executable specification with explicit domains and packing or decoding bounds; correctness and security arguments covering the actual transcript and corruption model; and adversarial test vectors.
Privacy across outputs and time	Input encryption alone does not bound inference from exact, small-group, adaptive, or overlapping outputs. Function-key policies, query membership, model states, and repeated releases extend the observer’s view [48], [51], [84], [123].	A privacy contract specifying the subject, observer, complete view, authorized outputs, residual leakage, and time horizon; analysis of combined releases; and inference evaluation with relevant auxiliary information.
Coalition and state compromise	Privacy depends on role-specific non-collusion or trusted components. Archived ciphertexts, recovery material, and exposed state may reveal information beyond that considered in the online analysis [44], [91], [126], [156].	Explicit coalition and state-exposure models, including corruption timing; quantified thresholds; erasure and rotation rules where required; and security arguments covering the permitted combinations.
Lifecycle-safe operation	Recovery and membership changes may alter the trust boundary. Reported fault tolerance may produce an aggregate of responding meters or an estimate, rather than the exact full-population result [90], [158], [160], [161].	A state-transition model with privacy, correctness, and availability conditions; explicit missing-data semantics; adversarial dropout and membership-change tests; and measured recovery, rekeying, and revocation costs.
Cross-layer composability	Guarantees for individual components do not automatically protect the combined views of identity, readings, traffic, models, or transactions in hybrid systems [99], [134]–[136], [138].	End-to-end data-flow and observer analysis; a composition argument matched to the claimed privacy notion; and attack evaluation across the relevant disclosure channels.
Reproducible deployment evidence	Evaluations often isolate selected operations while omitting setup, decoding, recovery, network, energy, or attack costs. Differences in security levels, workloads, and platforms further limit comparison [53], [94], [123], [155].	Comparable workloads and explicit security parameters; actor-specific time, communication, memory, and energy measurements; failure and adversarial tests; uncertainty reporting; and reusable code, parameters, and traces.

validation, and availability must be addressed according to the adversary model.

Selection should proceed by elimination. First fix the required output and acceptable error. Input transformations remain admissible when they preserve the required function, but methods that perturb the final result require an explicit error tolerance. Next exclude candidates with incompatible trust, coalition, metadata, or recovery assumptions. Only then compare computation, communication, storage, latency, and evidence maturity.

Hybrid designs are appropriate when requirements span families, such as exact aggregation followed by a differentially private release or anonymous authentication followed by encrypted settlement. The complete design must be reassessed from the observers’ combined views. The following scenarios illustrate this selection process.

D. Illustrative Design Scenarios

The following scenarios illustrate how application requirements lead to different mechanism combinations. They are derived from the corpus; the cited studies support individual design choices rather than establish the complete composite architectures. In each case, the authorized output is fixed first, followed by the observer, trust, lifecycle, and evidence requirements.

Scenario 1: verifiable time-of-use billing. Suppose a supplier requires an exact, verifiable bill under a published tariff without routine access to the underlying consumption profile. Local verifiable computation, homomorphic encryption, or functional encryption can support this requirement. Commitments and zero-knowledge proofs can bind the charge to certified readings and the tariff; range proofs can separately constrain admissible inputs [46], [48], [94]. Structured perturbation is also admissible when it preserves the tariff-weighted total, although exact billing alone does not establish privacy or verifiability [68]. Selection depends on key ownership, authorization of tariff changes, and dispute procedures. The bill, payment history, and exceptional disclosures remain part of the observer’s view, so retention and access to supporting interval-level evidence must be specified.

Scenario 2: fault-tolerant feeder aggregation. Consider a service requiring exact aggregation of received feeder reports during temporary meter dropout, with an honest-but-curious gateway and some potentially malicious meters. Masking, threshold encryption, or multiparty computation is admissible only when its coalition and recovery assumptions match this model [76], [90]. Recovery of cryptographic masks does not recover an unreported measurement: an exact full-population total requires additional measurement information, while omission or estimation changes the result’s meaning. Authentica-

TABLE XVI
REQUIREMENT-DRIVEN SELECTION OF PRIVACY-PRESERVING PROTOCOL FAMILIES

Dominant requirement	Candidate family and controls	What the selection can provide	Boundary that must be accepted or resolved
Exact bill, query, or policy-limited function	Local verifiable computation, homomorphic or functional encryption; attribute-based access where appropriate; commitments and proofs [46], [48], [84]	Computes or verifies an exact authorized result while limiting disclosure of the underlying readings.	The result remains visible. Key ownership, function-key policies, query overlap, arithmetic bounds, and revocation define the trust and leakage boundary.
Frequent exact aggregation with constrained meters	Additive or pairwise masking, secure aggregation, or hierarchical cancellation with dropout recovery [90], [161]	Can reduce online meter computation and reveal an exact sum over the defined contributing set under the stated setup and coalition model.	Mask cancellation, participant changes, recovery, and missing-data semantics must be explicit. Small groups and comparable overlapping outputs may expose contributions.
Exact aggregation under stronger coalitions or malicious inputs	Threshold encryption, secret sharing, secure multiparty computation, or distributed decryption with authentication and range or consistency proofs [76], [94]	Can distribute authority or computation and provide input privacy and verifiability against a defined coalition or active adversary.	Guarantees depend on corruption thresholds and the complete transcript. Rounds, shares, recovery, and quorum availability add costs; valid inputs need not be truthful measurements.
Repeated or externally released statistics	Central, local, or distributed differential privacy, optionally combined with cryptographic aggregation [86], [123]	Bounds changes in output distributions under the stated adjacency relation, including repeated releases when composition is accounted for.	Adjacency, sensitivity, clipping, budget allocation, and utility loss must match the service. Cryptographic protection does not replace release accounting.
Collaborative forecasting, profiling, or control	Federated learning with secure aggregation or encrypted computation; differential privacy and compatible poisoning defenses where required [124], [126]	Keeps raw records local and can restrict access to individual updates while producing a shared model or control policy.	Update and model leakage, participation, dropout, numeric encoding, and the coordinator's view require analysis. Poisoning defenses must remain compatible with protected updates.
Anonymous or traceable participation in mobility and markets	Anonymous credentials, identity-protecting signature constructions, rotating pseudonyms, commitments, or encrypted comparison [136], [138], [139]	Can authenticate eligibility, support conditional accountability, or match protected bids while restricting public identity or value disclosure.	Tracing, stable public values, timing, location, match outcomes, and settlement may restore linkage. Ledger persistence introduces additional disclosure.
Inference reduction before data disclosure	Battery- or renewable-assisted load shaping, flexible-load scheduling, local feature extraction, or data minimization [132]	Reduces sensitive structure in the observable load or limits the information leaving the premises.	Physical shaping depends on capacity, power, forecasts, comfort, and cost. Data transformations require their own leakage evaluation; both must preserve required service utility.
Long-lived confidentiality or post-quantum migration	Lattice-based encryption, aggregation, signatures, or key establishment with explicit security analysis and a migration profile [53], [155], [162]	Can protect the relevant cryptographic component against quantum-capable adversaries under the stated assumptions and parameters.	Security, applicable failure probabilities, arithmetic bounds, key and ciphertext sizes, authentication, revocation, and migration state require evaluation. Output leakage remains separate.

tion establishes report origin, and range or consistency proofs can constrain inputs, but neither establishes truthful physical consumption [94]. Evaluation should therefore distinguish recovery correctness, missing-data error, and malicious-input handling. Release controls must also account for comparable aggregates whose participant sets expose individual contributions through differencing.

Scenario 3: collaborative short-term forecasting. Assume households train a forecasting model without sending raw profiles to the utility. Federated learning satisfies this locality requirement, while secure aggregation, homomorphic encryption, or multiparty computation can restrict the coordinator's access to individual updates [124], [126]. A household-level differential-privacy guarantee additionally requires appropriate adjacency, clipping, noise calibration, and composition across released training outputs [123]; household participation and individual-record protection are different privacy units. If malicious clients are included, poisoning defenses must remain

compatible with protected updates. Evaluation should jointly examine dropout, numeric encoding, convergence, forecast error, inference attacks, and privacy-budget consumption under the stated observer model.

Scenario 4: accountable EV charging and local energy trading. Suppose an EV must prove eligibility and submit a charging or trading request without exposing its long-term identity to public participants, while an authorized entity retains settlement and dispute-resolution capability. Anonymous credentials, group signatures, or appropriately managed pseudonyms can restrict public identity disclosure. Hiding commitments and encrypted comparison can protect bid information during matching, with specified openings for settlement [136], [138], [139]. Ledger use should be justified by its auditability requirements and evaluated for persistent linkage. The station still observes physical presence and timing, while settlement or tracing authorities may learn identity and transaction details. Anonymity must therefore be specified

separately for public observers, peers, stations, operators, and tracing authorities, including their permitted coalitions.

A defensible design decision should record the selected mechanisms, decisive assumptions, rejected alternatives, residual leakage, and evidence still required before deployment. Each composite design must be assessed against its complete requirements rather than inferred to be secure from the individual components.

IX. THREATS TO VALIDITY

The findings of this survey depend on the completeness of the search, the consistency of study selection and coding, and the comparability of evidence reported by different studies. The review protocol and reproducibility controls reduce these risks but cannot eliminate them.

A. Search, Selection, and Publication Bias

The database search, snowballing, and supplementary Google Scholar searches described in Section II do not guarantee complete retrieval. Relevant studies may use unexpected terminology, be indexed incompletely, or appear outside the selected sources. Citation searching may also favor older, highly cited, or well-connected publications over recent or isolated contributions.

The eligibility rules introduce further boundaries. Relevant work may be absent when a technically sufficient full text or reliable analysis in English was unavailable. Publication bias may also underrepresent negative results and failed implementations. We reduced selection bias by retaining uncertain records for full-text assessment, applying the same eligibility rules to all discovery routes, and retaining eligible studies despite weak evaluation or unfavorable findings. These measures cannot account for relevant work that was never retrieved or published.

The search cutoff was 29 August 2026, so 2026 represents only a partial publication year. The findings therefore reflect the literature retrieved by that cutoff and should not be interpreted as covering all subsequently published or indexed work.

B. Classification, Extraction, and Researcher Bias

Evidence extraction and taxonomy construction require interpretive judgment. Terminology is inconsistent, and studies may leave trust assumptions, adversary capabilities, authorized outputs, or lifecycle behavior incomplete. For example, “privacy” may refer to input confidentiality or resistance to behavioral inference. Some coding decisions therefore depend on our interpretation of the complete construction rather than explicit author statements.

We reduced this risk through structured evidence cards and operational coding rules described in Sections II-E and II-F. Assessments were linked to supporting source locations, with author claims, supported properties, and reviewer-derived inferences distinguished. Missing evidence was recorded as *not reported*, while ambiguous or conflicting evidence was recorded as *unclear*. Multi-label coding preserved multiple applications, mechanisms, and privacy targets; exclusive primary

application assignments were used only for descriptive corpus counts. Affected studies were reconsidered when coding definitions changed.

These controls improve consistency but do not remove researcher bias. The taxonomy reflects analytical choices, and critical interpretations may differ from those intended by the original authors. The review does not report a corpus-wide inter-rater agreement statistic; traceability should therefore not be interpreted as independent validation of every coding decision. Studies authored by review-team members were assessed under the same rules, although potential bias remains. Supporting evidence and explicit critical interpretations allow readers to examine or challenge the resulting classifications.

C. Performance Comparability and Temporal Validity

Performance evidence is highly heterogeneous across the corpus. Studies differ in functions, adversary models, cryptographic parameters, datasets, system sizes, platforms, and cost boundaries. Measurements may also exclude setup, key management, recovery, communication, or device resources. Lower reported runtime or communication therefore does not establish greater efficiency or scalability for the complete system.

We did not perform a corpus-wide statistical meta-analysis or construct a universal performance ranking. Numerical results were compared directly only when functions, units, assumptions, and experimental conditions were sufficiently compatible. Otherwise, structured qualitative comparison preserved the original evaluation context, as described in Sections II-A and VII-C. Unreported optimizations, implementation differences, and limited uncertainty reporting may still affect comparisons. Performance conclusions therefore apply to the evaluated configurations rather than establish deployment-wide predictions.

Technical conclusions also have a temporal boundary. Cryptographic parameters, attack methods, standards, implementations, and deployment requirements evolve. Older evaluations may no longer represent relevant operating conditions, while recent proposals may lack independent analysis or deployment testing. Assessments of concrete security, performance, and evidence maturity reflect the material available at the review cutoff and should be revisited as stronger analyses and operational results emerge.

X. CONCLUSION

This survey examined privacy-preserving smart metering as an end-to-end system problem rather than as a catalogue of cryptographic and statistical techniques. Through a systematic analysis of 23 prior surveys and 150 primary studies published from 2010 to the search cutoff of 29 August 2026, it connected the service being provided with its authorized output, observable information, privacy mechanism, trust and collusion assumptions, residual leakage, lifecycle behavior, cost, and supporting evidence. This multidimensional view makes explicit why approaches using the same mechanism may provide different assurances and why protocols implementing different

outputs or adversary models should not be compared solely by runtime or communication cost.

The strongest finding is that protecting meter inputs or one protocol phase does not establish privacy for the complete service. Exact and repeated aggregates may permit differencing; function keys and query policies determine what encrypted computation ultimately reveals; federated execution can expose updates, models, or inferred schedules; and anonymous participation may remain linkable through metadata, location, transactions, or settlement records. Privacy also depends on whether the stated trust and corruption boundaries survive dropout, recovery, revocation, replacement, and later compromise. Across the review period, the literature progressed from comparatively narrow aggregation and billing constructions toward dynamic, verifiable, learning-based, market-oriented, and post-quantum designs. However, functionality has generally advanced faster than end-to-end assurance and deployment evidence. We found focused solutions to individual requirements, but no recurring protocol pattern that simultaneously resolves construction-level correctness and security, output and longitudinal leakage, coalition and state exposure, lifecycle transitions, cross-layer composition, and reproducible evaluation.

These findings suggest a disciplined direction for both research and protocol selection. A design should begin by specifying the minimum output required by the application, the parties permitted to receive it, the observations and auxiliary information available to each adversary, and the leakage acceptable over the system's lifetime. Candidate families with incompatible output, trust, coalition, or lifecycle assumptions should then be excluded before their performance is compared. Future work should support its claims with implementation-aligned specifications and proofs, explicit release and state-exposure models, adversarial lifecycle testing, and reproducible meter-edge-cloud evaluation. Progress should ultimately be measured not by the number of mechanisms combined or features claimed, but by whether the resulting system delivers the required service while making its disclosure, trust, failure, and evidentiary boundaries precise and defensible.

REFERENCES

- [1] U.S. Energy Information Administration, "Advanced metering count by technology type, 2015 through 2024," Electric Power Annual, Table 10.05; Forms EIA-861 and EIA-861S, 2025, 2024 final data; accessed September 17, 2026. [Online]. Available: https://www.eia.gov/electricity/annual/html/epa_10_05.html
- [2] Transforma Insights, "Smart meter installations projected to reach 3.9 billion globally by 2035, generating USD46 billion in annual revenue," Transforma Insights forecast summary, mar 2026, forecast covering electricity, gas, and water smart meters; published March 20, 2026; accessed September 17, 2026. [Online]. Available: <https://transformainsights.com/news/smart-meter-installations-forecast>
- [3] International Energy Agency, "Electricity grids and secure energy transitions," International Energy Agency, Paris, France, Tech. Rep., 2023, accessed September 17, 2026. [Online]. Available: <https://www.iea.org/reports/electricity-grids-and-secure-energy-transitions>
- [4] S. McLaughlin, P. McDaniel, and W. Aiello, "Protecting consumer privacy from electric load monitoring," in *18th ACM Conference on Computer and Communications Security (CCS'11)*, 2011, pp. 87–98.
- [5] G. Eibl and D. Engel, "Influence of data granularity on smart meter privacy," *IEEE Transactions on Smart Grid*, vol. 6, no. 2, pp. 930–939, 2015.
- [6] D. Chen, S. Kalra, D. Irwin, P. Shenoy, and J. Albrecht, "Preventing occupancy detection from smart meters," *IEEE Transactions on Smart Grid*, vol. 6, no. 5, pp. 2426–2434, 2015.
- [7] Z. Zaman, P. Gauravaram, S. Jha, and W. Hu, "Privacy-preserving smart metering through data synthesis with differential privacy," in *Proceedings of the 12th ACM International Conference on Systems for Energy-Efficient Buildings, Cities, and Transportation (BuildSys '25)*, 2025, pp. 1–11.
- [8] P. O. Ajiboye, K. O.-B. O. Agyekum, and E. A. Frimpong, "Privacy and security of advanced metering infrastructure (AMI) data and network: A comprehensive review," *Journal of Engineering and Applied Science*, vol. 71, no. 1, p. 91, 2024.
- [9] H. Bibi, M. Abolhasan, J. Lipman, M. Abdollahi, and W. Ni, "A comprehensive survey on privacy-preserving technologies for smart grids," *Computers & Electrical Engineering*, vol. 124, p. 110371, 2025.
- [10] Y. Han, Y. Wang, L. Wu, H. Feng, X. Wu, and R. Li, "Survey of privacy-preserving data aggregation schemes in smart grid," *Journal of King Saud University - Computer and Information Sciences*, vol. 37, p. 263, 2025.
- [11] P. Deokar, S. Arora, and A. A. Agarkar, "Privacy preserving data aggregation and authentication schemes in smart grid: An overview," *Information Security Journal: A Global Perspective*, vol. 35, no. 4, pp. 616–631, 2026.
- [12] M. K. Hasan, M. M. Hasan, N. Al-Qirim, S. N. H. S. Abdullah, S. Islam, and M. A. Razzaque, "Review on data privacy and security for IoT-based multifunctional layers of cyber-physical systems in smart grids," *Journal on Information Security*, vol. 2026, no. 1, p. 9, 2026.
- [13] B. Kitchenham and S. Charters, "Guidelines for performing systematic literature reviews in software engineering," Evidence-Based Software Engineering, Tech. Rep. EBSE-2007-01, 2007.
- [14] 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, p. n71, 2021.
- [15] M. L. Rethlefsen, S. Kirtley, S. Waffenschmidt, A. P. Ayala, D. Moher, M. J. Page, J. B. Koffel, and PRISMA-S Group, "PRISMA-S: An extension to the PRISMA statement for reporting literature searches in systematic reviews," *Systematic Reviews*, vol. 10, no. 1, p. 39, 2021.
- [16] A. Gopstein, C. Nguyen, C. O'Fallon, N. Hastings, and D. A. Wollman, "NIST framework and roadmap for smart grid interoperability standards, release 4.0," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Special Publication 1108r4, 2021.
- [17] V. Y. Pillitteri and T. L. Brewer, "Guidelines for smart grid cybersecurity," National Institute of Standards and Technology, Gaithersburg, MD, USA, NIST Interagency/Internal Report 7628 Revision 1, 2014.
- [18] M. R. Asghar, G. Dán, D. Miorandi, and I. Chlamtac, "Smart meter data privacy: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 4, pp. 2820–2835, 2017.
- [19] P. Kumar, Y. Lin, G. Bai, A. Paverd, J. S. Dong, and A. Martin, "Smart grid metering networks: A survey on security, privacy and open research issues," *IEEE Communications Surveys & Tutorials*, vol. 21, no. 3, pp. 2886–2927, 2019.
- [20] Z. Erkin, J. R. Troncoso-Pastoriza, R. I. L. Lagendijk, and F. Perez-Gonzalez, "Privacy-preserving data aggregation in smart metering systems: An overview," *IEEE Signal Processing Magazine*, vol. 30, no. 2, pp. 75–86, 2013.
- [21] S. Finster and I. Baumgart, "Privacy-aware smart metering: A survey," *IEEE Communications Surveys & Tutorials*, vol. 17, no. 2, pp. 1088–1101, 2015.
- [22] S. Sultan, "Privacy-preserving metering in smart grid for billing, operational metering, and incentive-based schemes: A survey," *Computers & Security*, vol. 84, pp. 148–165, 2019.
- [23] G. Kalogridis, C. Efthymiou, S. Z. Denic, T. A. Lewis, and R. Cepeda, "Privacy for smart meters: Towards undetectable appliance load signatures," in *First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010, pp. 232–237.
- [24] J. Kua, M. B. Hossain, I. Natgunanathan, and Y. Xiang, "Privacy preservation in smart meters: Current status, challenges and future directions," *Sensors*, vol. 23, no. 7, p. 3697, 2023.
- [25] W. Han and Y. Xiao, "Privacy preservation for V2G networks in smart grid: A survey," *Computer Communications*, vol. 91–92, pp. 17–28, 2016.
- [26] F. Farokhi, "Review of results on smart-meter privacy by data manipulation, demand shaping, and load scheduling," *IET Smart Grid*, vol. 3, no. 5, pp. 605–613, 2020.
- [27] J.-M. Bohli, C. Sorge, and O. Ugus, "A privacy model for smart metering," in *IEEE International Conference on Communications Workshops*

- (ICC Workshops), *First IEEE International Workshop on Smart Grid Communications, Cape Town, South Africa*, 2010, pp. 1–5.
- [28] M. A. Ferrag, L. A. Maglaras, H. Janicke, J. Jiang, and L. Shu, “A systematic review of data protection and privacy preservation schemes for smart grid communications,” *Sustainable Cities and Society*, vol. 38, pp. 806–835, 2018.
 - [29] Z. Abreu and L. Pereira, “Privacy protection in smart meters using homomorphic encryption: An overview,” *WIREs Data Mining and Knowledge Discovery*, vol. 12, no. 4, p. e1469, 2022.
 - [30] S. Desai, R. Alhadad, N. Chilamkurti, and A. Mahmood, “A survey of privacy preserving schemes in IoE enabled smart grid advanced metering infrastructure,” *Cluster Computing*, vol. 22, no. 1, pp. 43–69, 2019.
 - [31] L. Zhu, Z. Zhang, Z. Qin, J. Weng, and K. Ren, “Privacy protection using a rechargeable battery for energy consumption in smart grids,” *IEEE Network*, vol. 31, no. 1, pp. 59–63, 2017.
 - [32] S. Zeadally, A.-S. K. Pathan, C. Alcaraz, and M. Badra, “Towards privacy protection in smart grid,” *Wireless Personal Communications*, vol. 73, no. 1, pp. 23–50, 2013.
 - [33] N. Saputro and K. Akkaya, “On preserving user privacy in smart grid advanced metering infrastructure applications,” *Security and Communication Networks*, vol. 7, no. 1, pp. 206–220, 2014.
 - [34] A. Agarkar and H. Agrawal, “A review and vision on authentication and privacy preservation schemes in smart grid network,” *Security and Privacy*, vol. 2, no. 4, p. e62, 2019.
 - [35] R. B. Romdhane, H. Hammami, M. Hamdi, and T.-H. Kim, “Privacy-preserving approaches for smart metering: A survey,” in *Proceedings of the 2022 Asia Conference on Electrical, Power and Computer Engineering*, 2022, pp. 1–6.
 - [36] R. Dong, S. Hao, T. H. Yang, Z. Tang, Y. Yan, and J. Chen, “Recent advances in smart meter: Data analysis, privacy preservation and applications,” in *Big Data and Security*, ser. Communications in Computer and Information Science, vol. 1563. Springer, 2022, pp. 105–114.
 - [37] M. S. Abdalzaher, M. M. Fouda, and M. I. Ibrahim, “Data privacy preservation and security in smart metering systems,” *Energies*, vol. 15, no. 19, p. 7419, 2022.
 - [38] L. Zhang, J. Zhang, and Y. H. Hu, “A privacy-preserving distributed smart metering temporal and spatial aggregation scheme,” *IEEE Access*, vol. 7, pp. 28 372–28 382, 2019.
 - [39] L. Chen, R. Lu, and Z. Cao, “PDAFT: A privacy-preserving data aggregation scheme with fault tolerance for smart grid communications,” *Peer-to-Peer Networking and Applications*, vol. 8, no. 6, pp. 1122–1132, 2015.
 - [40] N. Naeem, F. Khan, S. Tahir, I. Makhdoom, M. Tariq, and S. A. Shah, “Blockchain and homomorphic-based data aggregation for load monitoring and billing in smart grids,” *International Journal of Information Security*, vol. 25, no. 3, p. 82, 2026.
 - [41] H. J. Jo, I. S. Kim, and D. H. Lee, “Efficient and privacy-preserving metering protocols for smart grid systems,” *IEEE Transactions on Smart Grid*, vol. 7, no. 3, pp. 1732–1742, 2016.
 - [42] A. Abdallah and X. S. Shen, “Lightweight security and privacy preserving scheme for smart grid customer-side networks,” *IEEE Transactions on Smart Grid*, vol. 8, no. 3, pp. 1064–1074, 2017.
 - [43] S. Zhang, T. Li, Y. Sun, and B. Wang, “A lattice-based anti-quantum privacy-preserving scheme for smart meter,” *International Journal of Electrical Power & Energy Systems*, vol. 166, p. 110502, 2025.
 - [44] L. Zhu, M. Li, Z. Zhang, C. Xu, R. Zhang, X. Du, and N. Guizani, “Privacy-preserving authentication and data aggregation for fog-based smart grid,” *IEEE Communications Magazine*, vol. 57, no. 6, pp. 80–85, 2019.
 - [45] A. Molina-Markham, P. Shenoy, K. Fu, E. Cecchet, and D. Irwin, “Private memoirs of a smart meter,” in *2nd ACM Workshop on Embedded Sensing Systems for Energy-Efficiency in Building (BuildSys ’10)*, 2010, pp. 61–66.
 - [46] A. Rial and G. Danezis, “Privacy-preserving smart metering,” in *10th ACM Workshop on Privacy in the Electronic Society (WPES’11)*, 2011, pp. 49–60.
 - [47] J. Ni, K. Zhang, X. Lin, and X. S. Shen, “Balancing security and efficiency for smart metering against misbehaving collectors,” *IEEE Transactions on Smart Grid*, vol. 10, no. 2, pp. 1225–1236, 2019.
 - [48] M. I. Ibrahim, M. M. Badr, M. M. Fouda, M. Mahmoud, W. Alasmay, and Z. M. Fadlullah, “PMBFE: Efficient and privacy-preserving monitoring and billing using functional encryption for AMI networks,” in *IEEE International Symposium on Networks, Computers and Communications (ISNCC)*, 2020, pp. 1–7.
 - [49] M. I. Ibrahim, M. Nabil, M. M. Fouda, M. M. E. A. Mahmoud, W. Alasmay, and F. Alsolami, “Efficient privacy-preserving electricity theft detection with dynamic billing and load monitoring for AMI networks,” *IEEE Internet of Things Journal*, vol. 8, no. 2, pp. 1243–1258, 2021.
 - [50] K. Xue, Q. Yang, S. Li, D. S. L. Wei, M. Peng, I. Memon, and P. Hong, “PPSO: A privacy-preserving service outsourcing scheme for real-time pricing demand response in smart grid,” *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 2486–2496, 2019.
 - [51] A. Alsharif, M. Nabil, A. Sherif, M. Mahmoud, and M. Song, “MDMS: Efficient and privacy-preserving multidimension and multisubset data collection for AMI networks,” *IEEE Internet of Things Journal*, vol. 6, no. 6, pp. 10 363–10 374, 2019.
 - [52] S. Zhao, F. Li, H. Li, R. Lu, S. Ren, H. Bao, J.-H. Lin, and S. Han, “Smart and practical privacy-preserving data aggregation for fog-based smart grids,” *IEEE Transactions on Information Forensics and Security*, vol. 16, pp. 521–536, 2021.
 - [53] A. Marandi, P. G. M. R. Alves, D. F. Aranha, and R. H. Jacobsen, “Lattice-based homomorphic encryption for privacy-preserving smart meter data analytics,” *The Computer Journal*, vol. 67, no. 5, pp. 1687–1698, 2024.
 - [54] W. Liu, Q. Li, J. Yuan, and J. Shi, “A privacy-preserving billing computation and leakage detection framework for smart grid,” *IEEE Internet of Things Journal*, vol. 12, no. 5, pp. 5961–5969, 2025.
 - [55] Z. Zeng, Z. Yu, Y. Chen, L. Li, Y. Liu, and H. Liu, “A multi-functional and privacy-preserving data aggregation scheme for smart grid,” *Journal of Systems Architecture*, vol. 170, p. 103599, 2026.
 - [56] S. Li, K. Xue, Q. Yang, and P. Hong, “Ppma: privacy-preserving multisubset data aggregation in smart grid,” *IEEE Transactions on Industrial Informatics*, vol. 14, no. 2, pp. 462–471, 2018.
 - [57] J.-N. Liu, J. Weng, A. Yang, Y. Chen, and X. Lin, “Enabling efficient and privacy-preserving aggregation communication and function query for fog computing-based smart grid,” *IEEE Transactions on Smart Grid*, vol. 11, no. 1, pp. 247–257, 2020.
 - [58] D. Chen, T. Zhou, W. Liu, R. Li, L. Wu, and X. Yang, “MDA-FLH: Multidimensional data aggregation scheme with fine-grained linear homomorphism for smart grid,” *IEEE Internet of Things Journal*, vol. 11, no. 2, pp. 3524–3538, 2024.
 - [59] P. Gope and B. Sikdar, “An efficient data aggregation scheme for privacy-friendly dynamic pricing-based billing and demand-response management in smart grids,” *IEEE Internet of Things Journal*, vol. 5, no. 4, pp. 3126–3135, 2018.
 - [60] A. Alsharif, M. Nabil, S. Tonyali, H. Mohammed, M. Mahmoud, and K. Akkaya, “EPIC: Efficient privacy-preserving scheme with EtoE data integrity and authenticity for AMI networks,” *IEEE Internet of Things Journal*, vol. 6, no. 2, pp. 3309–3321, 2019.
 - [61] S. Tonyali, O. Cakmak, K. Akkaya, M. M. E. A. Mahmoud, and I. Guvenc, “Secure data obfuscation scheme to enable privacy-preserving state estimation in smart grid AMI networks,” *IEEE Internet of Things Journal*, vol. 3, no. 5, pp. 709–719, 2016.
 - [62] P. Barbosa, A. Brito, and H. Almeida, “A technique to provide differential privacy for appliance usage in smart metering,” *Information Sciences*, vol. 370–371, pp. 355–367, 2016.
 - [63] Z. Zheng, T. Wang, A. K. Bashir, M. Alazab, S. Mumtaz, and X. Wang, “A decentralized mechanism based on differential privacy for privacy-preserving computation in smart grid,” *IEEE Transactions on Computers*, vol. 71, no. 11, pp. 2915–2926, 2022.
 - [64] M. U. Hassan, M. H. Rehmani, R. Kotagiri, J. Zhang, and J. Chen, “Differential privacy for renewable energy resources based smart metering,” *Journal of Parallel and Distributed Computing*, vol. 131, pp. 69–80, 2019.
 - [65] M. B. Gough, S. F. Santos, T. AlSkaif, M. S. Javadi, R. Castro, and J. P. S. Catalão, “Preserving privacy of smart meter data in a smart grid environment,” *IEEE Transactions on Industrial Informatics*, vol. 18, no. 1, pp. 707–718, 2022.
 - [66] H.-Y. Tran, J. Hu, and H. R. Pota, “Smart meter data obfuscation with a hybrid privacy-preserving data publishing scheme without a trusted third party,” *IEEE Internet of Things Journal*, vol. 9, no. 17, pp. 16 080–16 095, 2022.
 - [67] Z. Zhao, G. Liu, and Y. Liu, “Practical privacy-preserving electricity theft detection for smart grid,” *IEEE Transactions on Smart Grid*, vol. 15, no. 4, pp. 4104–4114, 2024.
 - [68] F. Zaredar and M. Amini, “A lightweight privacy-preserving smart metering billing protocol with dynamic tariff policy adjustment,” arXiv preprint, arXiv:2508.14815, primary category cs.CR, 2025. [Online]. Available: <https://arxiv.org/abs/2508.14815>
 - [69] X. Liang, X. Li, R. Lu, X. Lin, and X. Shen, “UDP: Usage-based dynamic pricing with privacy preservation for smart grid,” *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 141–150, 2013.

- [70] Y. Gong, Y. Cai, Y. Guo, and Y. Fang, "A privacy-preserving scheme for incentive-based demand response in the smart grid," *IEEE Transactions on Smart Grid*, vol. 7, no. 3, pp. 1304–1313, 2016.
- [71] T. W. Chim, S.-M. Yiu, V. O. Li, L. C. Hui, and J. Zhong, "PRGA: Privacy-preserving recording & gateway-assisted authentication of power usage information for smart grid," *IEEE Transactions on Dependable and Secure Computing*, vol. 12, no. 1, pp. 85–97, 2015.
- [72] M. F. Balli, S. Uludag, A. A. Selcuk, and B. Tavli, "Distributed multi-unit privacy assured bidding (PAB) for smart grid demand response programs," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4119–4127, 2018.
- [73] F. Li, B. Luo, and P. Liu, "Secure information aggregation for smart grids using homomorphic encryption," in *First IEEE International Conference on Smart Grid Communications (SmartGridComm)*, 2010, pp. 327–332.
- [74] K. Kursawe, G. Danezis, and M. Kohlweiss, "Privacy-friendly aggregation for the smart-grid," in *Privacy Enhancing Technologies (PETS 2011)*, ser. Lecture Notes in Computer Science, vol. 6794. Springer, 2011, pp. 175–191.
- [75] C.-I. Fan, S.-Y. Huang, and Y.-L. Lai, "Privacy enhanced data aggregation scheme against internal attackers in smart grid," *IEEE Transactions on Industrial Informatics*, vol. 10, no. 1, pp. 666–675, 2014.
- [76] N. Busom, R. Petric, F. Seb , C. Sorce, and M. Valls, "Efficient smart metering based on homomorphic encryption," *Computer Communications*, vol. 82, pp. 95–101, 2016.
- [77] F. Zaregar and M. Amini, "A collusion-resistance privacy-preserving smart metering protocol for operational utility," 2025. [Online]. Available: <https://arxiv.org/abs/2508.14744>
- [78] R. Lu, X. Liang, X. Li, X. Lin, and X. S. Shen, "EPPA: An efficient and privacy-preserving aggregation scheme for secure smart grid communications," *IEEE Transactions on Parallel and Distributed Systems*, vol. 23, no. 9, pp. 1621–1631, 2012.
- [79] L. Chen, R. Lu, Z. Cao, K. AlHarbi, and X. Lin, "MuDA: Multifunctional data aggregation in privacy-preserving smart grid communications," *Peer-to-Peer Networking and Applications*, vol. 8, no. 5, pp. 777–792, 2015.
- [80] H. Shen, M. Zhang, and J. Shen, "Efficient privacy-preserving cube-data aggregation scheme for smart grids," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 6, pp. 1369–1381, 2017.
- [81] A. Mohammadali and M. S. Haghighi, "A privacy-preserving homomorphic scheme with multiple dimensions and fault tolerance for metering data aggregation in smart grid," *IEEE Transactions on Smart Grid*, vol. 12, no. 6, pp. 5212–5220, 2021.
- [82] Z. Li, L. Deng, X. Xiao, N. Wang, L. Liu, and S. Li, "Certificateless privacy-preserving multidimensional data aggregation scheme with fault-tolerance for fog-based smart grids," *Journal of Systems Architecture*, vol. 175, p. 103773, 2026.
- [83] X. Wang, Y. Liu, and K.-K. R. Choo, "Fault-tolerant multisubset aggregation scheme for smart grid," *IEEE Transactions on Industrial Informatics*, vol. 17, no. 6, pp. 4065–4072, 2021.
- [84] X. Li, M. Zhao, Y. Ding, C. Yang, H. Wang, H. Liang, and Y. Wang, "Privacy-preserving data aggregation achieving completeness of data queries in smart grid," *International Journal of Network Management*, vol. 34, no. 6, p. e2273, 2024.
- [85] G. Li, L. Zhou, J. Lian, S. Liu, L. Yang, Y. Zhong, and Q. Li, "Efficient sharing of energy consumption data: A privacy-preserving threshold aggregation approach," *IEEE Internet of Things Journal*, vol. 12, no. 13, pp. 23 988–24 002, 2025.
- [86] G.  cs and C. Castelluccia, "I have a DREAM! (Differentially private smart metering)," in *Information Hiding (IH 2011)*, ser. Lecture Notes in Computer Science, vol. 6958. Springer, 2011, pp. 118–132.
- [87] H. Bao and R. Lu, "A new differentially private data aggregation with fault tolerance for smart grid communications," *IEEE Internet of Things Journal*, vol. 2, no. 3, pp. 248–258, 2015.
- [88] Z. Shi, R. Sun, R. Lu, L. Chen, J. Chen, and X. S. Shen, "Diverse grouping-based aggregation protocol with error detection for smart grid communications," *IEEE Transactions on Smart Grid*, vol. 6, no. 6, pp. 2856–2868, 2015.
- [89] J. Ni, K. Zhang, K. Alharbi, X. Lin, N. Zhang, and X. S. Shen, "Differentially private smart metering with fault tolerance and range-based filtering," *IEEE Transactions on Smart Grid*, vol. 8, no. 5, pp. 2483–2493, 2017.
- [90] K. Xue, B. Zhu, Q. Yang, D. S. L. Wei, and M. Guizani, "An efficient and robust data aggregation scheme without a trusted authority for smart grid," *IEEE Internet of Things Journal*, vol. 7, no. 3, pp. 1949–1959, 2020.
- [91] L. Wu, M. Xu, S. Fu, Y. Luo, and Y. Wei, "FPDA: Fault-tolerant and privacy-enhanced data aggregation scheme in fog-assisted smart grid," *IEEE Internet of Things Journal*, vol. 9, no. 7, pp. 5254–5265, 2022.
- [92] Q. Liu, C. Lin, Y. Yin, H. Wang, Z. Cai, and Y. Zhao, "A dynamic and secure smart grid data aggregation scheme," *International Journal of Cognitive Computing in Engineering*, vol. 7, pp. 412–421, 2026.
- [93] J. Zhang and C. Dong, "Privacy-preserving data aggregation scheme against deletion and tampering attacks from aggregators," *Journal of King Saud University – Computer and Information Sciences*, vol. 35, no. 4, pp. 100–111, 2023.
- [94] H. Pang, K. He, Y. Fu, J.-N. Liu, X. Liu, and W. Tan, "Enabling efficient and malicious secure data aggregation in smart grid with false data detection," *IEEE Transactions on Smart Grid*, vol. 15, no. 2, pp. 2203–2213, 2024.
- [95] P. Yu, W. Huang, and Z. Li, "A secure, lightweight, and verifiable data aggregation scheme for smart grids," *Peer-to-Peer Networking and Applications*, vol. 18, no. 3, p. 135, 2025.
- [96] H. Li, X. Lin, H. Yang, X. Liang, R. Lu, and X. S. Shen, "EPPDR: An efficient privacy-preserving demand response scheme with adaptive key evolution in smart grid," *IEEE Transactions on Parallel and Distributed Systems*, vol. 25, no. 8, pp. 2053–2064, 2014.
- [97] C. Efthymiou and G. Kalogridis, "Smart grid privacy via anonymization of smart metering data," in *First IEEE International Conference on Smart Grid Communications (SmartGridComm 2010)*, 2010, pp. 238–243.
- [98] F. Diao, F. Zhang, and X. Cheng, "A privacy-preserving smart metering scheme using linkable anonymous credential," *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 461–467, 2015.
- [99] M. I. Ibrahim, M. Mahmoud, M. M. Fouda, F. Alsolami, W. Alasmari, and X. Shen, "Privacy preserving and efficient data collection scheme for AMI networks using deep learning," *IEEE Internet of Things Journal*, vol. 8, no. 23, pp. 17 131–17 146, 2021.
- [100] J.-L. Tsai and N.-W. Lo, "Secure anonymous key distribution scheme for smart grid," *IEEE Transactions on Smart Grid*, vol. 7, no. 2, pp. 906–914, 2016.
- [101] V. Odelu, A. K. Das, M. Wazid, and M. Conti, "Provably secure authenticated key agreement scheme for smart grid," *IEEE Transactions on Smart Grid*, vol. 9, no. 3, pp. 1900–1910, 2018.
- [102] D. Abbasinezhad-Mood and M. Nikooghadam, "An anonymous ECC-Based self-certified key distribution scheme for the smart grid," *IEEE Transactions on Industrial Electronics*, vol. 65, no. 10, pp. 7996–8004, 2018.
- [103] L. Zhang, L. Zhao, S. Yin, C.-H. Chi, R. Liu, and Y. Zhang, "A lightweight authentication scheme with privacy protection for smart grid communications," *Future Generation Computer Systems*, vol. 100, pp. 770–778, 2019.
- [104] A. A. Khan, V. Kumar, M. Ahmad, S. Rana, and D. Mishra, "PALK: Password-based anonymous lightweight key agreement framework for smart grid," *International Journal of Electrical Power & Energy Systems*, vol. 121, p. 106121, 2020.
- [105] D. Sadhukhan, S. Ray, M. S. Obaidat, and M. Dasgupta, "A secure and privacy preserving lightweight authentication scheme for smart-grid communication using elliptic curve cryptography," *Journal of Systems Architecture*, vol. 114, p. 101938, 2021.
- [106] S. Yu, K. Park, J. Lee, Y. Park, Y. Park, S. Lee, and B. Chung, "Privacy-preserving lightweight authentication protocol for demand response management in smart grid environment," *Applied Sciences*, vol. 10, no. 5, p. 1758, 2020.
- [107] U. Habib, M. Bano, J. Iqbal, and S. S. Ullah, "A novel efficient and secure privacy-preserving scheme for smart grid based on HECC," *IEEE Access*, vol. 13, pp. 49 290–49 308, 2025.
- [108] R. Baksh, S. Ito, and M. Ahmad, "A comprehensive and secure scheme for privacy-preserving smart meter data aggregation in the smart grid," *Sustainable Energy, Grids and Networks*, vol. 39, p. 101461, 2024.
- [109] P. Gope and B. Sikdar, "Privacy-aware authenticated key agreement scheme for secure smart grid communication," *IEEE Transactions on Smart Grid*, vol. 10, no. 4, pp. 3953–3962, 2019.
- [110] P. Gope, "PMAKE: Privacy-aware multi-factor authenticated key establishment scheme for advanced metering infrastructure in smart grid," *Computer Communications*, vol. 152, pp. 338–344, 2020.
- [111] S. Rostampour, N. Bagheri, B. Ghavami, Y. Bendavid, S. Kumari, H. Martin, and C. Camara, "Using a privacy-enhanced authentication process to secure IoT-based smart grid infrastructures," *The Journal of Supercomputing. The paper was published online in*, vol. 80, pp. 1668–1693, 2024.

- [112] S. Ruj and A. Nayak, "A decentralized security framework for data aggregation and access control in smart grids," *IEEE Transactions on Smart Grid*, vol. 4, no. 1, pp. 196–205, 2013.
- [113] A. Alsharif, M. Nabil, M. M. E. A. Mahmoud, and M. Abdallah, "EPDA: Efficient and privacy-preserving data collection and access control scheme for multi-recipient AMI networks," *IEEE Access*, vol. 7, pp. 27 829–27 845, 2019.
- [114] H. Liu, Y. Peng, Y. Liu, and Z. Zeng, "Privacy-preserving multidimensional data aggregation for diverse electricity data users," *Journal of Systems Architecture*, vol. 161, p. 103363, 2025.
- [115] S. Wang, Y. Liu, W. Chai, F. Feng, Y. Liu, and Z. Liu, "A multi-authority aggregate signature scheme for privacy-preserving and efficient message authentication in the smart grid," *Cybersecurity*, vol. 9, no. 1, p. 90, 2026.
- [116] K. Gai, Y. Wu, L. Zhu, L. Xu, and Y. Zhang, "Permissioned blockchain and edge computing empowered privacy-preserving smart grid networks," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 7992–8004, 2019.
- [117] L. Sankar, S. R. Rajagopalan, S. Mohajer, and H. V. Poor, "Smart meter privacy: A theoretical framework," *IEEE Transactions on Smart Grid*, vol. 4, no. 2, pp. 837–846, 2013.
- [118] M. Shateri, F. Messina, P. Piantanida, and F. Labeau, "Deep directed information-based learning for privacy-preserving smart meter data release," in *IEEE SmartGridComm*, 2019, pp. 1–7.
- [119] —, "Real-time privacy-preserving data release for smart meters," *IEEE Transactions on Smart Grid*, vol. 11, no. 6, pp. 5174–5183, 2020.
- [120] L. Ou, Z. Qin, S. Liao, T. Li, and D. Zhang, "Singular spectrum analysis for local differential privacy of classifications in the smart grid," *IEEE Internet of Things Journal*, vol. 7, no. 6, pp. 5246–5255, 2020.
- [121] J. Lin, J. Ma, and J. Zhu, "Privacy-preserving household characteristic identification with federated learning method," *IEEE Transactions on Smart Grid*, vol. 13, no. 2, pp. 1088–1099, 2022.
- [122] H. U. Manzoor, A. Jafri, and A. Zoha, "Adaptive single-layer aggregation framework for energy-efficient and privacy-preserving load forecasting in heterogeneous federated smart grids," *Internet of Things*, vol. 28, p. 101376, 2024.
- [123] J. D. Fernández, S. P. Menci, C. M. Lee, A. Rieger, and G. Fridgen, "Privacy-preserving federated learning for residential short-term load forecasting," *Applied Energy*, vol. 326, p. 119915, 2022.
- [124] M. M. Badr, M. I. Ibrahim, M. Mahmoud, W. Alasmay, M. M. Fouda, K. H. Almotairi, and Z. M. Fadlullah, "Privacy-preserving federated-learning-based net-energy forecasting," in *IEEE SoutheastCon*, 2022, pp. 133–139.
- [125] M. Wen, R. Xie, K. Lu, L. Wang, and K. Zhang, "FedDetect: A novel privacy-preserving federated learning framework for energy theft detection in smart grid," *IEEE Internet of Things Journal*, vol. 9, no. 8, pp. 6069–6080, 2022.
- [126] P. Zhang, Y. Zhang, Z. Wang, Q. Zhang, and Z. Gu, "Privacy-preserving federated learning scheme for distributed smart grid based on multi-key homomorphic encryption," *Peer-to-Peer Networking and Applications*, vol. 18, no. 6, p. 293, 2025.
- [127] Z. He, H. Zhang, and Y. Liu, "Privacy-preserving electricity load forecasting using hierarchical federated learning," *Journal of Information and Intelligence*, vol. 4, no. 4, pp. 350–363, 2026.
- [128] I. A. Kamil and S. O. Ogunloyin, "EPDAS: Efficient privacy-preserving data analysis scheme for smart grid network," *Journal of King Saud University – Computer and Information Sciences*, vol. 33, no. 2, pp. 208–217, 2021.
- [129] D. Yao, M. Wen, X. Liang, Z. Fu, K. Zhang, and B. Yang, "Energy theft detection with energy privacy preservation in the smart grid," *IEEE Internet of Things Journal*, vol. 6, no. 5, pp. 7659–7669, 2019.
- [130] L. Wu, S. Fu, Y. Luo, and M. Xu, "SecTCN: Privacy-preserving short-term residential electrical load forecasting," *IEEE Transactions on Industrial Informatics*, vol. 20, no. 2, pp. 2508–2518, 2024.
- [131] F. Zaredar and M. Amini, "A lightweight incentive-based privacy-preserving smart metering protocol for value-added services," arXiv submission manuscript, cs.CR, 2026. [Online]. Available: <https://arxiv.org/abs/2508.14703>
- [132] G. Giacon, D. Gündüz, and H. V. Poor, "Smart meter privacy with renewable energy and an energy storage device," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 1, pp. 129–142, 2018.
- [133] G. K. Verma, A. Tiwari, M. Wadhwa, and N. Kumar, "Efficient and secure homomorphic data aggregation for smart meter-to-grid communication," *Computer Networks*, vol. 285, p. 112422, 2026.
- [134] A. Sharma, A. L. Sangal, and K. C. Sharma, "Privacy-preserving prosumer profiling using smart meter data," *Sustainable Energy, Grids and Networks*, vol. 43, p. 101784, 2025.
- [135] S. Lee, L. Xie, and D.-H. Choi, "Privacy-preserving energy management of a shared energy storage system for smart buildings: A federated deep reinforcement learning approach," *Sensors*, vol. 21, no. 14, p. 4898, 2021.
- [136] Y.-B. Son, J.-H. Im, H.-Y. Kwon, S.-Y. Jeon, and M.-K. Lee, "Privacy-preserving peer-to-peer energy trading in blockchain-enabled smart grids using functional encryption," *Energies*, vol. 13, no. 6, p. 1321, 2020.
- [137] A. Abdallah and X. S. Shen, "Lightweight authentication and privacy-preserving scheme for V2G connections," *IEEE Transactions on Vehicular Technology*, vol. 66, no. 3, pp. 2615–2629, 2017.
- [138] Y. Liang, Y. Liu, X. Zhang, and G. Liu, "Physically secure and privacy-preserving charging authentication framework with data aggregation in vehicle-to-grid networks," *IEEE Transactions on Intelligent Transportation Systems*, vol. 25, no. 11, pp. 18 831–18 846, 2024.
- [139] T. F. Vallent, D. Hanyurwimfura, and C. Mikeka, "Efficient certificate-less aggregate signature scheme with conditional privacy-preservation for vehicular ad hoc networks enhanced smart grid system," *Sensors*, vol. 21, p. 2900, 2021.
- [140] S. Yu and K. Park, "PUF-Based robust and anonymous authentication and key establishment scheme for V2G networks," *IEEE Internet of Things Journal*, vol. 11, no. 9, pp. 15 450–15 464, 2024.
- [141] K. Prateek, S. Maity, and N. Saxena, "QSKA: A quantum secured privacy-preserving mutual authentication scheme for energy internet-based vehicle-to-grid communication," *IEEE Transactions on Network and Service Management*, vol. 21, no. 6, pp. 6810–6826, 2024.
- [142] Y. Sun, L. Lampe, and V. W. S. Wong, "Smart meter privacy: Exploiting the potential of household energy storage units," *IEEE Internet of Things Journal*, vol. 5, no. 1, pp. 69–78, 2018.
- [143] O. Tan, D. Gündüz, and H. V. Poor, "Increasing smart meter privacy through energy harvesting and storage devices," *IEEE Journal on Selected Areas in Communications*, vol. 31, no. 7, pp. 1331–1341, 2013.
- [144] S. Li, A. Khisti, and A. Mahajan, "Information-theoretic privacy for smart metering systems with a rechargeable battery," *IEEE Transactions on Information Theory*, vol. 64, no. 5, pp. 3679–3695, 2018.
- [145] F. Farokhi and H. Sandberg, "Fisher information as a measure of privacy: Preserving privacy of households with smart meters using batteries," *IEEE Transactions on Smart Grid*, vol. 9, no. 5, pp. 4726–4734, 2018.
- [146] J.-X. Chin, K. Baker, and G. Hug, "Consumer privacy protection using flexible thermal loads: Theoretical limits and practical considerations," *Applied Energy*, vol. 281, p. 116075, 2021.
- [147] L. Yang, X. Chen, J. Zhang, and H. V. Poor, "Optimal privacy-preserving energy management for smart meters," in *IEEE INFOCOM*, 2014, pp. 513–521.
- [148] —, "Cost-effective and privacy-preserving energy management for smart meters," *IEEE Transactions on Smart Grid*, vol. 6, no. 1, pp. 486–495, 2015.
- [149] I. Natgunanathan, M. B. Hossain, Y. Xiang, L. Gao, D. Peng, and J. Li, "Progressive average-based smart meter privacy enhancement using rechargeable batteries," *IEEE Internet of Things Journal*, vol. 6, no. 6, pp. 9816–9828, 2019.
- [150] M. B. Hossain, I. Natgunanathan, Y. Xiang, L.-X. Yang, and G. Huang, "Enhanced smart meter privacy protection using rechargeable batteries," *IEEE Internet of Things Journal*, vol. 6, no. 4, pp. 7079–7092, 2019.
- [151] J. Zhao, T. Jung, Y. Wang, and X. Li, "Achieving differential privacy of data disclosure in the smart grid," in *IEEE INFOCOM*, 2014, pp. 504–512.
- [152] Z. Zhang, Z. Qin, L. Zhu, J. Weng, and K. Ren, "Cost-friendly differential privacy for smart meters: Exploiting the dual roles of the noise," *IEEE Transactions on Smart Grid*, vol. 8, no. 2, pp. 619–626, 2017.
- [153] X. Liu, H. Wang, G. Chen, B. Zhou, and A. ur Rehman, "Intermittently differential privacy in smart meters via rechargeable batteries," *Electric Power Systems Research*, vol. 199, p. 107410, 2021.
- [154] L. Wang, C. Zhao, Z. Liu, B. Wang, K. Zhang, M. He, and W. Li, "Fault-tolerant and key-leakage resilient lightweight multidimensional privacy-preserving data aggregation scheme in smart grid," *IEEE Transactions on Dependable and Secure Computing*, vol. 23, no. 2, pp. 4057–4070, 2026.

- [155] K. Fan, Y. Ren, Y. Bai, G. Wei, K. Zhang, H. Li, and Y. Yang, "Fault-tolerant and collusion-resistant lattice-based multidimensional privacy-preserving data aggregation in edge-based smart grid," *IEEE Internet of Things Journal*, vol. 11, no. 6, pp. 9487–9504, 2024.
- [156] Y. Liu, W. Guo, C.-I. Fan, L. Chang, and C. Cheng, "A practical privacy-preserving data aggregation (3pda) scheme for smart grid," *IEEE Transactions on Industrial Informatics*, vol. 15, no. 3, pp. 1767–1774, 2019.
- [157] X. Zuo, L. Li, H. Peng, S. Luo, and Y. Yang, "Privacy-preserving multidimensional data aggregation scheme without trusted authority in smart grid," *IEEE Systems Journal*, vol. 15, no. 1, pp. 395–406, 2021.
- [158] H. Liu, Y. Peng, Z. Yu, Y. Liu, L. Li, and Z. Zeng, "FPMDA: Fault-tolerant and privacy-enhanced multidimensional data aggregation without TA," *IEEE Internet of Things Journal*, vol. 12, no. 23, pp. 50547–50558, 2025.
- [159] S. Zhao, "Smart grids data aggregation method on paillier homomorphic encryption," *Applied Artificial Intelligence*, vol. 38, no. 1, p. 2327901, 2024.
- [160] Y. Chen, J.-F. Martínez-Ortega, L. López, H. Yu, and Z. Yang, "A dynamic membership group-based multiple-data aggregation scheme for smart grid," *IEEE Internet of Things Journal*, vol. 8, no. 15, pp. 12360–12374, 2021.
- [161] Y. Chang, J. Li, N. Lu, W. Shi, Z. Su, and W. Meng, "Practical privacy-preserving scheme with fault tolerance for smart grids," *IEEE Internet of Things Journal*, vol. 11, no. 2, pp. 1990–2005, 2024.
- [162] J. Qian, Z. Cao, X. Dong, J. Shen, Z. Liu, and Y. Ye, "Two secure and efficient lightweight data aggregation schemes for smart grid," *IEEE Transactions on Smart Grid*, vol. 12, no. 3, pp. 2625–2637, 2021.
- [163] A. Abdallah and X. S. Shen, "A lightweight lattice-based homomorphic privacy-preserving data aggregation scheme for smart grid," *IEEE Transactions on Smart Grid*, vol. 9, no. 1, pp. 396–405, 2018.
- [164] M. Aurangzeb, Y. Wang, S. Iqbal, A. Naveed, Z. Ahmed, M. Alenezi, and M. Shouran, "Enhancing cybersecurity in smart grids: Deep black box adversarial attacks and quantum voting ensemble models for blockchain privacy-preserving storage," *Energy Reports*, vol. 11, pp. 2493–2515, 2024.
- [165] O. R. M. Boudia, S. M. Senouci, and M. Feham, "Elliptic curve-based secure multidimensional aggregation for smart grid communications," *IEEE Sensors Journal*, vol. 17, no. 23, pp. 7750–7757, 2017.
- [166] Y. Chen, J.-F. Martínez-Ortega, P. Castillejo, and L. López, "A homomorphic-based multiple data aggregation scheme for smart grid," *IEEE Sensors Journal*, vol. 19, no. 10, pp. 3921–3929, 2019.
- [167] F. Fioretto, T. W. K. Mak, and P. V. Hentenryck, "Differential privacy for power grid obfuscation," *IEEE Transactions on Smart Grid*, vol. 11, no. 2, pp. 1356–1366, 2020.
- [168] N. Gai, K. Xue, P. He, B. Zhu, J. Liu, and D. He, "An efficient data aggregation scheme with local differential privacy in smart grid," in *16th International Conference on Mobility, Sensing and Networking (MSN)*, 2020, pp. 73–80.
- [169] H. M. Khan, A. Khan, F. Jabeen, and A. U. Rahman, "Privacy preserving data aggregation with fault tolerance in fog-enabled smart grids," *Sustainable Cities and Society*, vol. 64, p. 102522, 2021.
- [170] Z. Liu, Z. Cao, X. Dong, X. Zhao, T. Liu, H. Bao, and J. Shen, "EPMDA-FED: Efficient and privacy-preserving multidimensional data aggregation scheme with fast error detection in smart grid," *IEEE Internet of Things Journal*, vol. 9, no. 9, pp. 6922–6933, 2022.
- [171] J. Liu, H. Wang, J. Bao, R. Sun, X. Du, and M. Guizani, "RPMDDA: Robust and privacy-enhanced multidimensional data aggregation scheme for fog-assisted smart grids," *IEEE Internet of Things Journal*, vol. 11, no. 9, pp. 16021–16032, 2024.
- [172] L. Lyu, K. Nandakumar, B. Rubinstein, J. Jin, J. Bedo, and M. Palaniswami, "PPFA: Privacy preserving fog-enabled aggregation in smart grid," *IEEE Transactions on Industrial Informatics*, vol. 14, no. 8, pp. 3733–3744, 2018.
- [173] O. R. Merad-Boudia and S. M. Senouci, "An efficient and secure multidimensional data aggregation for fog-computing-based smart grid," *IEEE Internet of Things Journal*, vol. 8, no. 8, pp. 6143–6153, 2021.
- [174] A. Saleem, A. Khan, S. U. R. Malik, H. Pervaiz, H. Malik, M. Alam, and A. Jindal, "FESDA: Fog-enabled secure data aggregation in smart grid IoT network," *IEEE Internet of Things Journal*, vol. 7, no. 7, pp. 6132–6142, 2020.
- [175] P. Singh, M. Masud, M. S. Hossain, and A. Kaur, "Blockchain and homomorphic encryption-based privacy-preserving data aggregation model in smart grid," *Computers & Electrical Engineering*, vol. 93, p. 107209, 2021.
- [176] L. Wu, S. Fu, Y. Luo, H. Yan, H. Shi, and M. Xu, "A robust and lightweight privacy-preserving data aggregation scheme for smart grid," *IEEE Transactions on Dependable and Secure Computing*, vol. 21, no. 1, pp. 270–283, 2024.
- [177] X. Xiao, L. Deng, Z. Li, N. Wang, Y. Qu, and S. Li, "Privacy-preserving multi-dimensional and multi-subset data aggregation scheme against joint attacks in smart grids," *Sustainable Energy, Grids and Networks*, vol. 46, p. 102316, 2026.
- [178] Z. Zeng, X. Wang, Y. Liu, and L. Chang, "MSDA: multi-subset data aggregation scheme without trusted third party," *Frontiers of Computer Science*, vol. 16, no. 1, p. 161808, 2022.
- [179] J. Zhang, W. Zhang, X. Wei, and H. Liu, "EPri-MDAS: An efficient privacy-preserving multiple data aggregation scheme without trusted authority for fog-based smart grid," *High-Confidence Computing*, vol. 4, no. 4, p. 100226, 2024.
- [180] X. Zhang, X. Deng, H. Zhang, Y. Zhang, and J. Xue, "Fine-grained encrypted data aggregation mechanism with fault tolerance in edge-assisted smart grids," *Journal of Information Security and Applications*, vol. 86, p. 103888, 2024.
- [181] J. Zhang and J. Wei, "PFDAM: Privacy-preserving fine-grained data aggregation scheme supporting multifunctionality in smart grid," *IEEE Internet of Things Journal*, vol. 11, no. 15, pp. 25520–25533, 2024.
- [182] L. Zhang, Y. Yang, J. Luo, W. Wang, and J. Gu, "DBIA-DA: dual-blockchain and ISCP-assisted data aggregation for fog-enabled smart grid," *EURASIP Journal on Wireless Communications and Networking*, vol. 2025, no. 1, p. 79, 2025.

APPENDIX A COMPLETE SEARCH STRINGS

The final database search was executed on 29 August 2026 in IEEE Xplore, Scopus, and the Web of Science Core Collection. The publication window began in January 2010 and ended on the search date; 2026 is therefore a partial year. The main query combined smart-metering and general privacy concepts. A supplementary technical query was also executed in Scopus and Web of Science to retrieve studies that described a privacy mechanism without using the general privacy vocabulary consistently. No venue-ranking, citation-count, or open-access filter was applied.

In the preserved export filenames, Q1 identifies the main search and Q2 the supplementary technical search; these labels are distinct from research questions RQ1–RQ8. The following strings reproduce the database-specific Boolean queries used in the review workflow. Their line breaks and indentation have been adjusted only for readability and do not alter the search logic. The corresponding result sets were retained as the original database exports reported in Table XVII.

A. IEEE Xplore

The main query was applied to *All Metadata*. The same Boolean expression was executed over three non-overlapping year intervals to remain within the practical export limit: 2010–2017, 2018–2022, and 2023–2026.

```
("smart meter" OR "smart meters" OR
"smart metering" OR
"advanced metering infrastructure")
AND
(privacy OR "privacy preserving" OR
"privacy-preserving" OR "privacy aware" OR
"privacy-aware" OR anonymization OR anonymity OR
unlinkability OR pseudonymity OR obfuscation OR
"information leakage" OR "data protection")
```

Only the year boundary was applied at retrieval. Language, publication type, and technical relevance were assessed during screening rather than encoded in the recorded IEEE query.

B. Scopus

Both searches used the title, abstract, and keyword field. Results were limited to 2010–2026, English, and the document types *article*, *conference paper*, or *review*.

Main search (Q1).

```
TITLE-ABS-KEY(
("smart meter" OR "smart meters" OR
"smart metering" OR
"advanced metering infrastructure")
AND
(privacy OR "privacy preserving" OR
"privacy-preserving" OR "privacy aware" OR
"privacy-aware" OR anonymization OR anonymity OR
unlinkability OR pseudonymity OR obfuscation OR
"information leakage" OR "data protection")
)
```

Supplementary technical search (Q2).

```
TITLE-ABS-KEY(
("smart meter" OR "smart meters" OR
"smart metering" OR
"advanced metering infrastructure")
AND
("homomorphic encryption" OR
"differential privacy" OR "secret sharing" OR
"secure multiparty computation" OR
"secure multi-party computation" OR
"zero-knowledge proof" OR
"anonymous credential" OR "blind signature" OR
"functional encryption" OR
"attribute-based encryption" OR
"secure aggregation" OR "data masking" OR
"load shaping" OR "battery-based privacy" OR
blockchain OR "federated learning" OR
"trusted execution environment" OR
"lattice-based" OR "post-quantum")
)
```

C. Web of Science Core Collection

Both searches used the *Topic* field, covering titles, abstracts, author keywords, and Keywords Plus. The timespan was 2010–2026 and the language was English. The preserved records are classified as *Article* or *Review*, including compound classifications such as *Article; Proceedings Paper* and *Article; Early Access*.

Main search (Q1).

```
TS(("smart meter" OR "smart meters" OR
"smart metering" OR
"advanced metering infrastructure")
AND
(privacy OR "privacy preserving" OR
"privacy-preserving" OR "privacy aware" OR
"privacy-aware" OR anonymization OR anonymity OR
unlinkability OR pseudonymity OR obfuscation OR
"information leakage" OR "data protection"))
```

Supplementary technical search (Q2).

```
TS(("smart meter" OR "smart meters" OR
"smart metering" OR
"advanced metering infrastructure")
AND
("homomorphic encryption" OR
"differential privacy" OR "secret sharing" OR
"secure multiparty computation" OR
"secure multi-party computation" OR
"zero-knowledge proof" OR
"anonymous credential" OR "blind signature" OR
"functional encryption" OR
"attribute-based encryption" OR
"secure aggregation" OR "data masking" OR
"load shaping" OR "battery-based privacy" OR
blockchain OR "federated learning" OR
"trusted execution environment" OR
"lattice-based" OR "post-quantum"))
```

D. Preserved Exports and Record Counts

Table XVII reports the final exports. Counts refer to raw export rows before within- or cross-database deduplication.

The preliminary file `IEEE_Q1_pilot_top1000_2026-08-29.csv` contains 1,000 calibration records, 974 of which also appear in the final segmented IEEE exports. It was retained for audit purposes but excluded from the identification count. The main and supplementary searches intentionally overlap. After within-source consolidation, the exports contained 2,371 IEEE, 4,140 Scopus, and 2,162 Web of Science records. DOI-first and normalized-title matching across sources removed approximately 4,878 duplicate rows and produced approximately 5,565 unique candidates; the latter remains approximate because DOI-less records and related publication versions required manual resolution.

APPENDIX B

DETAILED ELIGIBILITY CRITERIA

Tables XVIII and XIX define the operational rules used to construct the primary-study, prior-survey, and contextual corpora.

APPENDIX C

INCLUDED PRIMARY STUDIES

This appendix fixes the membership of the primary-study corpus and assigns a stable identifier to each included publication. The identifiers P001–P150 are ordered alphabetically by resolved BibTeX key; they indicate neither study quality nor chronological priority. Each identifier maps to exactly one bibliographic record and is used consistently in the supplementary protocol matrix. Prior surveys and contextual references are excluded from this registry.

P001–P025. P001 [102]; P002 [137]; P003 [163]; P004 [42]; P005 [86]; P006 [60]; P007 [113]; P008 [51]; P009 [164]; P010 [124]; P011 [108]; P012 [72]; P013 [87]; P014 [62]; P015 [27]; P016 [165]; P017 [76]; P018 [161]; P019 [79]; P020 [39]; P021 [166]; P022 [160]; P023 [58]; P024 [71]; P025 [146].

P026–P050. P026 [98]; P027 [97]; P028 [5]; P029 [75]; P030 [155]; P031 [145]; P032 [123]; P033 [167]; P034 [116]; P035 [168]; P036 [132]; P037 [70]; P038 [59]; P039 [109]; P040 [110]; P041 [65]; P042 [107]; P043 [64]; P044 [127]; P045 [150]; P046 [48]; P047 [49]; P048 [99]; P049 [41]; P050 [23].

P051–P075. P051 [128]; P052 [104]; P053 [169]; P054 [74]; P055 [135]; P056 [73]; P057 [96]; P058 [56]; P059 [144]; P060 [84]; P061 [85]; P062 [82]; P063 [69]; P064 [138]; P065 [121]; P066 [156]; P067 [57]; P068 [170]; P069 [153]; P070 [171]; P071 [54]; P072 [114]; P073 [158]; P074 [92]; P075 [78].

P076–P100. P076 [172]; P077 [122]; P078 [53]; P079 [4]; P080 [173]; P081 [81]; P082 [45]; P083 [40]; P084 [149]; P085 [47]; P086 [89]; P087 [101]; P088 [120]; P089 [94]; P090 [141]; P091 [162]; P092 [46]; P093 [111]; P094 [112]; P095 [105]; P096 [174]; P097 [117]; P098 [134]; P099 [118]; P100 [119].

P101–P125. P101 [80]; P102 [88]; P103 [175]; P104 [136]; P105 [142]; P106 [143]; P107 [61]; P108 [66]; P109 [100]; P110 [139]; P111 [133]; P112 [83]; P113 [154]; P114 [115]; P115 [125]; P116 [91]; P117 [176]; P118 [130]; P119 [177];

TABLE XVII
FINAL DATABASE EXPORTS AND VERIFIED RAW RECORD COUNTS

Preserved export	Search field and filters	Range	Rows
IEEE_Q1A_2010-2017_2026-08-29.csv	All Metadata; year partition	2010-2017	686
IEEE_Q1B_2018-2022_2026-08-29.csv	All Metadata; year partition	2018-2022	824
IEEE_Q1C_2023-2026_2026-08-29.csv	All Metadata; year partition	2023-2026	861
Scopus_Q1_main_2010-2026_2026-08-29.csv	TITLE-ABS-KEY; English; article, conference paper, or review	2010-2026	1,924
Scopus_Q2_supplementary_technical_2010-2026_2026-08-29.csv	TITLE-ABS-KEY; same filters as Scopus Q1	2010-2026	3,319
WoS_Q1_main_2010-2026_2026-08-29_final.xls	Topic; English; Article or Review records	2010-2026	956
WoS_Q2_supplementary_technical_2010-2026_2026-08-29_part1_1-1000.xls	Topic; same settings as Web of Science Q1	1-1,000	1,000
WoS_Q2_supplementary_technical_2010-2026_2026-08-29_part2_1001-1873.xls	Topic; same settings as Web of Science Q1	1,001-1,873	873
Total raw export rows			10,443

TABLE XVIII
OPERATIONAL INCLUSION CRITERIA FOR THE REVIEW CORPORA

ID	Criterion	Operational inclusion rule
Primary-study corpus		
Original technical studies used to answer RQ2-RQ8		
I1	Availability and date	A complete study had to be publicly available between January 2010 and the final search cutoff in 2026, inclusively. Complete online-first articles available by the cutoff were eligible. When applicable, both the online-first date and the final publication date were recorded.
I2	Smart-metering relevance	Smart meters, advanced metering infrastructure (AMI), or an equivalent metering setting had to form a central part of the system, application, data source, or evaluation. Eligible settings included consumers, prosumers, smart buildings, distributed energy resources, electric vehicles, and vehicle-to-grid services.
I3	Privacy objective	Privacy preservation had to be an explicit technical objective. Protected information could include meter readings, identity, linkability, appliance use, occupancy, behavioral attributes, billing or charging information, learning data, or information inferred from released measurements or the observable net-load profile.
I4	Technical contribution	The study had to present, extend, analyze, implement, evaluate, or cryptanalyze a privacy-preserving protocol, mechanism, architecture, privacy model, leakage measure, or customer-side protection method. A complete communication protocol was not required when the study made a direct contribution to privacy modeling, measurement, or evaluation.
I5	Evidence sufficiency	The full text had to provide enough information to identify the study objective, smart-metering context, principal mechanism, intended output, and form of analytical or empirical evidence. Missing coding properties were recorded as <i>not reported</i> rather than causing automatic exclusion.
I6	Publication type	Eligible sources included peer-reviewed journal articles, full conference papers, technically complete book chapters, and publicly accessible preprints containing a complete technical description. Publication type and peer-review status were recorded as evidence-maturity properties. Venue rank and citation count were not independent inclusion requirements.
Prior-survey corpus		
Secondary studies used to answer RQ1 and identify the remaining survey gap		
S1	Survey relevance	A survey or review was included when it substantially examined privacy in smart metering, AMI, privacy-preserving data aggregation, or a closely related smart-grid setting. A broader smart-grid security review was retained only when it contained an identifiable and technically relevant smart-meter privacy component.

P120 [50]; P121 [90]; P122 [148]; P123 [147]; P124 [129]; P125 [106].

P126-P150. P126 [140]; P127 [95]; P128 [7]; P129 [68]; P130 [77]; P131 [131]; P132 [178]; P133 [55]; P134 [152]; P135 [103]; P136 [38]; P137 [93]; P138 [179]; P139 [180]; P140 [181]; P141 [43]; P142 [182]; P143 [126]; P144 [151]; P145 [52]; P146 [67]; P147 [159]; P148 [63]; P149 [44]; P150 [157].

APPENDIX D DETAILED CODING DEFINITIONS

This appendix specifies the operational rules used to convert the evidence cards into the study-level taxonomy and supplementary protocol matrix. Coding was performed on the most complete eligible publication version available at the search cutoff. Related preprints, conference papers, early-access records, and journal versions were reconciled to avoid extracting the same technical contribution more than once. A related publication was retained separately only when it

TABLE XIX
OPERATIONAL EXCLUSION AND EVIDENCE-REASSIGNMENT CRITERIA

ID	Criterion and resulting decision
E1	Outside the temporal boundary. The publication fell outside the defined period or became publicly available only after the final search cutoff. Such a study was excluded from the primary-study corpus.
E2	General cybersecurity without a privacy contribution. The study addressed smart-grid cybersecurity but made no meaningful technical contribution to privacy preservation, privacy measurement, or privacy evaluation.
E3	Security-only mechanism. The study addressed authentication, key agreement, access control, intrusion detection, electricity-theft detection, or data integrity without protecting a privacy property such as identity, unlinkability, metering-data confidentiality, or resistance to behavioral inference.
E4	Generic method without a smart-metering instantiation. A generic cryptographic, blockchain, Internet-of-Things, artificial intelligence, or machine-learning method merely mentioned smart grids and was not adapted to or evaluated using a smart-metering system, service, or data source.
E5	Threat or inference study without a protection contribution. The study demonstrated a privacy threat or inference attack but did not provide a privacy mechanism, formal privacy model, leakage measure, or direct evaluation of a privacy-preserving approach. Important threat studies could be reassigned to the contextual corpus.
E6	Secondary or non-research publication. Surveys, tutorials, editorials, and other secondary publications were not counted as primary studies. Relevant surveys were assigned to the prior-survey corpus.
E7	Insufficient full-text evidence. Only an abstract, poster, presentation, or other technically incomplete version was available; the full text did not satisfy I5; or reliable full-text analysis in English was not possible.
E8	Duplicate or superseded version. Related versions of the same study were linked, and the most complete eligible version was used as the principal source. A later extension was counted separately only when it introduced a substantive new mechanism, function, guarantee, adversary model, or evaluation. A preprint and its peer-reviewed version were treated as one study rather than as two independent studies.
E9	Retraction or withdrawal. A publication formally retracted or withdrawn by the final search cutoff was excluded.

introduced a substantively different construction, function, threat model, or evaluation.

Evidence and missingness codes. Each coded property was assigned the most specific evidence status supported by the source. Table XX defines the source-level codes used in the evidence cards and coding registry. An author statement and its supporting evidence were recorded separately; therefore, an explicitly stated property was not treated as established solely because it appeared in the paper.

These source-level codes are retained in the evidence cards and coding registry but are not reproduced in Table XXIII. For the compact matrix, the evidence associated with each assessment dimension is synthesized as *D*, *P*, or *–*. This synthesis considers the relevant definitions, equations, proofs, implementations, and experiments together; it is not an automatic one-to-one conversion between the two coding systems.

Operational codebook. Table XXI defines the principal extraction fields and the rules used to assign their values. Multi-label coding was permitted where a study genuinely supported several applications, mechanisms, privacy targets, adversaries, or evidence types.

Cross-cutting rules. Absence of evidence was not coded as evidence of absence. When terminology differed across studies, labels were normalized only when their operational meaning was equivalent. A mechanism, guarantee, or lifecycle property was coded at the narrowest scope supported by the source. Numerical comparisons were treated as directly comparable only when the studies implemented compatible functions and used sufficiently similar adversary assumptions, security parameters, platforms, datasets, system sizes, and units. Later cryptanalysis or evaluation was recorded as related evidence and did not replace the claims or evidence originally

presented by the assessed study.

APPENDIX E PRIMARY APPLICATION ASSIGNMENTS

Table XXII records one primary application family and one subcategory for each included study. The assignment follows the principal function or technical contribution; secondary functions remain recorded separately in the protocol matrix. For multifunctional studies, the primary assignment is a descriptive coding decision and does not restrict their treatment in the synthesis. All percentages use the complete 150-study corpus as the denominator and are rounded to one decimal place.

APPENDIX F SUPPLEMENTARY PROTOCOL MATRIX

Table XXIII gives a compact, citation-linked view of all 150 primary studies. Controlled codes describe protocol scope, while the evidence symbols report the support available for each assessment dimension; they are not an overall quality score.

Within the protocol-scope columns, *App.* identifies the application or service addressed, *Mech.* identifies the privacy-preserving mechanism, and *Out.* records the result authorized for release. Within the assurance columns, *Adv.* covers the adversary and collusion model, *Priv.* the privacy definition or analysis, *Sec.* the security analysis, and *Corr.* functional correctness. The operational columns report evidence for lifecycle and robustness support (*Life.*), performance (*Perf.*), implementation or empirical evaluation (*Eval.*), and reproducibility (*Repr.*). The symbols and codes used in Table XXIII are defined in Appendix F-A.

TABLE XX
EVIDENCE AND MISSINGNESS CODES USED DURING STUDY-LEVEL CODING

Code	Status	Operational definition
S	Supported	The property is directly supported by an appropriate definition, construction, algorithm, equation, proof, implementation, experiment, or reported result. Support is limited to the stated assumptions, adversary, system configuration, and evaluated conditions.
C	Claimed only	The authors explicitly state the property, but the publication does not provide sufficient technical or empirical evidence to establish it at the claimed scope. The claim is retained without being upgraded to supported.
I	Reviewer inference	The property follows reasonably from the reported construction or results but is not stated explicitly by the authors. The inference is recorded separately from an author claim and linked to its supporting location.
NR	Not reported	The publication does not provide the information needed to code the property after examination of the relevant text, algorithms, equations, tables, figures, appendices, and available supplementary material.
UNC	Unclear	The available evidence is incomplete, ambiguous, or internally inconsistent, so no defensible value can be assigned. The source of the uncertainty is retained in the evidence card.
NA	Not applicable	The property is outside the logical scope of the design. This code is not used merely because the study omits a relevant requirement.

TABLE XXI
OPERATIONAL DEFINITIONS FOR THE PRINCIPAL STUDY-LEVEL CODING FIELDS

Coding field	Recorded values	Decision rule
Bibliographic identity and version	Title, authors, year, venue, DOI or other identifier, publication status, discovery source, and relationship to other versions.	Versions were linked using DOI, normalized title, authorship, year, venue, and technical content. Bibliographic variants were not counted as separate studies unless they contained a substantively distinct contribution.
Application and service	Billing and settlement; grid operation and aggregation; communication and authentication; analytics and data sharing; distributed-energy and e-mobility services; customer-side privacy control.	The application was assigned from the function actually implemented or evaluated, not from background motivation or a list of possible uses. Multiple values were permitted when distinct services were technically supported.
Input, function, recipient, and authorized output	Input information; computed function; permitted recipient; exact, approximate, perturbed, or decision output; release granularity and frequency.	The authorized output was the information intentionally made available by the complete protocol, including bills, aggregates, query answers, models, predictions, matches, schedules, session keys, and settlement records. Intermediate confidentiality was not coded as output privacy.
Architecture and information flow	Centralized, hierarchical, distributed, peer-to-peer, customer-side, cloud, fog, edge, ledger-assisted, or federated organization; participating entities and exchanged information.	Entity roles were coded from the data and authority they possess rather than their names. Information visible during registration, reporting, aggregation, recovery, billing, and revocation was recorded separately when these phases differed.
Privacy and supporting mechanisms	Encryption, secure computation, masking, secret sharing, differential privacy, anonymization, credentials, trusted hardware, learning-based protection, and customer-side transformation or load shaping.	A component was coded as a privacy mechanism only when it directly restricted disclosure, attribution, or inference. Authentication, blockchain, fog placement, or federated execution was coded as a supporting enabler unless its specific privacy role was established.
Trust, adversary, and collusion	Trusted, honest-but-curious, untrusted, or malicious entities; passive and active capabilities; insider, compromise, physical, inference, and auxiliary-information capabilities; tolerated and excluded coalitions.	Trust was assigned by entity and protocol phase. A stronger adversary or collusion guarantee was not inferred from general terms such as “secure,” “untrusted,” or “collusion resistant.” Thresholds and excluded role combinations were recorded explicitly.
Privacy target and subject	Readings and intermediate values; identity and linkability; behavior and occupancy; billing, charging, or market information; analytics and model information; participation metadata; underlying physical demand.	The coding identified what information is protected, whose information it is, and from which observer. Protection against one observer or for one information type was not generalized to other entities or privacy targets.
Guarantee and supporting evidence	Computational or information-theoretic confidentiality; differential privacy; anonymity, pseudonymity, or unlinkability; analytical leakage bounds; empirical inference resistance; information-theoretic or load-shaping metrics.	The guarantee was coded at the scope established by its definition and evidence. A proof for one primitive, an informal security analysis, or an experiment against one attack was not interpreted as an end-to-end guarantee.
Residual leakage	Authorized outputs, repeated or overlapping releases, metadata, timing, location, persistent identifiers, model updates, public-ledger records, collusion views, recovery state, and physical constraints.	Information remaining available under successful protocol execution was recorded even when the principal privacy property was satisfied. Required application output was treated as authorized leakage rather than as information eliminated by the mechanism.
Correctness and lifecycle	Normal-case output recovery; joining and leaving; dropout and communication failure; revocation; key renewal; compromise recovery; replacement; state synchronization.	A lifecycle property was supported only when the corresponding procedure, required participants, tolerated threshold, state change, and effect on privacy and correctness were described. Returning a result after omitting or estimating missing inputs was distinguished from recovering the originally intended result.
Integrity and malicious behavior	Authentication, freshness, replay protection, range or consistency proofs, aggregate verification, query completeness, robust aggregation, and poisoning resistance.	Each control was coded according to what it verifies. Message integrity was not treated as proof that a compromised meter submitted a truthful physical measurement, and input-range validation was not treated as full malicious-participant resistance.
Efficiency and scalability	Computation, communication, storage, memory, energy, rounds, latency, preprocessing, decoding, recovery cost, and scaling variables.	Costs were recorded by entity and protocol phase with their units, security parameters, system size, and measurement platform. Terms such as “lightweight” or “scalable” remained claims when no sufficient analysis or measurement was reported.
Evaluation and evidence maturity	Complexity analysis, proof, simulation, real-data experiment, software prototype, network or hardware testbed, meter-class implementation, pilot, and operational deployment; available code, data, parameters, and scripts.	Only evidence actually presented was coded. A desktop prototype did not establish meter-class feasibility, simulation did not establish deployment, and availability of a public dataset did not imply reproducibility when implementation details or parameters were missing.
Quality and limitations	Clarity of objective and system model; threat-model completeness; mechanism–guarantee alignment; correctness and operational support; evaluation adequacy; reproducibility and maturity.	The six domains were assessed separately. They were not combined into a single quality score, and weakness in one domain could not be offset by strength in another. Quality qualified the interpretation of individual claims but did not exclude an otherwise eligible study.

TABLE XXII
COMPLETE CITATION-LINKED INDEX OF EXCLUSIVE PRIMARY APPLICATION ASSIGNMENTS FOR THE 150 INCLUDED STUDIES

Primary subcategory	number	Share	Included studies
Billing, tariffs, incentives, and settlement			
B1 Temporal aggregation and billing separation	4	2.7%	[38], [40]–[42]
B2 Local and delegated verifiable billing	3	2.0%	[45]–[47]
B3 Encrypted or functional computation	4	2.7%	[48], [52], [54], [55]
B4 Masking, obfuscation, and lightweight computation	1	0.7%	[59]
B5 Perturbation with billing preservation or correction	4	2.7%	[62], [63], [65], [68]
B6 Dynamic pricing, incentives, and settlement	3	2.0%	[50], [69], [70]
Grid operations, aggregation, and flexibility			
G1 Exact spatial and temporal aggregation	11	7.3%	[73]–[78], [156], [159], [162], [163], [174]
G2 Multidimensional, multisubset, and query-oriented aggregation	17	11.3%	[51], [56], [58], [79]–[82], [157], [165], [166], [170], [173], [175], [177]–[179], [181]
G3 Noisy and differentially private operational outputs	8	5.3%	[27], [61], [86], [87], [89], [167], [168], [172]
G4 Failure handling and dynamic participation	14	9.3%	[39], [83], [90]–[92], [154], [155], [158], [160], [161], [169], [171], [176], [180]
G5 Integrity, verifiability, and malicious inputs	7	4.7%	[60], [84], [88], [93]–[95], [182]
G6 Demand response and flexibility coordination	2	1.3%	[72], [96]
Communication, authentication, and access			
C1 Privacy-aware reporting and identity protection	5	3.3%	[43], [44], [97]–[99]
C2 Anonymous authentication and key establishment	10	6.7%	[71], [100]–[108]
C3 Device-assisted and PUF-based authentication	3	2.0%	[109]–[111]
C4 Access control and multi-recipient disclosure	3	2.0%	[112], [113], [115]
Analytics, data sharing, and value-added services			
A1 Sanitized, minimized, and synthetic data release	6	4.0%	[7], [66], [117]–[120]
A2 Federated and distributed learning	7	4.7%	[121]–[127]
A3 Encrypted statistics, queries, and private inference	7	4.7%	[49], [53], [57], [67], [128]–[130]
A4 Controlled sharing and recipient-specific disclosure	3	2.0%	[85], [114], [164]
A5 Incentivized value-added-service collection	1	0.7%	[131]
Prosumer, distributed-energy, and e-mobility services			
P1 Renewable metering and prosumer analytics	3	2.0%	[64], [133], [134]
P2 Shared-storage coordination	1	0.7%	[135]
P3 Prosumer trading and governance	2	1.3%	[116], [136]
P4 E-mobility authentication, aggregation, and settlement	5	3.3%	[137]–[141]
Customer-side privacy control			
U1 Local data minimization and transformed reporting	1	0.7%	[5]
U2 Battery-assisted appliance-signature hiding	3	2.0%	[4], [23], [142]
U3 Information- and estimation-theoretic load shaping	5	3.3%	[132], [143]–[146]
U4 Cost-aware and DP-oriented control	7	4.7%	[147]–[153]
Total	150	100.0%	Each study appears exactly once

TABLE XXIII: Compact study-level protocol matrix for the 150 primary studies

Study	Year	Protocol scope			Assurance evidence				Operational evidence			
		App.	Mech.	Out.	Adv.	Priv.	Sec.	Corr.	Life.	Perf.	Eval.	Repr.
Foundational diversification (2010–2013)												
Bohli et al. [27]	2010	AG, BI	DP, AUTH	EX, AP, BL	P	D	P	D	P	P	–	P
Efthymiou et al. [97]	2010	AG, BI, DR, AN, ID	AUTH	BL, ML	P	P	P	P	P	–	–	P
Kalogridis et al. [23]	2010	CS	LS	SH	–	P	P	P	P	P	P	P
Li et al. [73]	2010	AG	HE	EX, QU	P	P	P	D	P	P	–	P
Molina-Markham et al. [45]	2010	AG, BI, DR, ID, CS	HE, AUTH	BL, QU, ML	P	P	–	P	–	P	P	P
Kursawe et al. [74]	2011	AG, BI, AN	SMC	EX	D	P	P	D	P	D	D	P
McLaughlin et al. [4]	2011	AG, CS	LS	SH	P	P	P	P	P	D	P	P
Rial et al. [46]	2011	BI, AN	HE, AUTH	BL	D	D	D	D	P	D	D	P
Ács et al. [86]	2011	AG, CS	SMC, DP	AP	P	D	P	D	D	P	P	P
Lu et al. [78]	2012	AG	HE	EX	P	P	P	D	–	D	P	P
Liang et al. [69]	2013	AG, BI	SMC	EX, BL	P	P	P	D	P	P	–	P
Ruj et al. [112]	2013	AG, ID	HE, FE, AUTH	EX, KY	P	P	D	D	P	P	–	P
Sankar et al. [117]	2013	AG, CS	ENC	AP	P	D	P	P	–	P	–	P
Tan et al. [143]	2013	AN	SMC, LS	SH	P	D	P	D	–	D	P	P
Expansion toward applications and resilience (2014–2018)												
Fan et al. [75]	2014	AG	HE, SMC	EX, KY	P	P	D	D	–	D	P	P
Li et al. [96]	2014	AG, DR	HE	EX, KY	P	P	P	D	P	P	P	P
Yang et al. [147]	2014	CS	LS	SH	–	P	–	D	P	D	P	P
Zhao et al. [151]	2014	CS	DP, LS	SH	P	D	D	P	P	P	P	P
Bao et al. [87]	2015	AG	HE, DP	AP	P	D	P	P	D	P	P	–
Chen et al. [79]	2015	AG, BI, AN	HE, DP	EX, AP, QU	P	P	P	P	–	P	P	P
Chen et al. [39]	2015	AG, BI	HE, SMC	EX	D	P	P	D	D	P	–	P
Chim et al. [71]	2015	AG, BI, ID	HE, AUTH	EX	P	P	P	P	P	P	P	P
Diao et al. [98]	2015	AG, AN, ID	SMC, AUTH, HW	QU	P	P	–	P	P	P	–	–
Eibl et al. [5]	2015	CS	LS	SH	P	P	–	P	–	P	P	P
Shi et al. [88]	2015	AG	SMC, DP	AP, QU	P	D	P	P	D	D	P	P
Yang et al. [148]	2015	CS	LS	SH	–	P	–	D	P	D	P	P
Barbosa et al. [62]	2016	BI, AN, CS	DP	AP, BL, ML	P	D	P	D	–	D	P	D
Busom et al. [76]	2016	AG	HE, SMC	EX	D	D	D	D	–	P	P	P
Gong et al. [70]	2016	AG, BI, DR	AUTH	TR	P	D	P	P	D	P	P	P
Jo et al. [41]	2016	AG, ID	HE, AUTH	EX	P	P	D	D	P	D	P	P
Odelu et al. [101]	2016	ID	AUTH	KY	D	P	P	D	–	P	P	P
Tonyali et al. [61]	2016	AG, BI, AN	SMC	EX, BL	P	–	P	D	P	D	P	P
Tsai et al. [100]	2016	ID	AUTH	KY	P	P	P	P	–	P	–	P
Abdallah et al. [137]	2017	BI, DR, ID, DE, EV	AUTH	AG, BL, TR	P	–	P	P	P	D	P	–
Abdallah et al. [42]	2017	BI, DR, AN, CS	LS	AG, BL, ML	P	P	P	P	P	P	P	P
Boudia et al. [165]	2017	AG	HE	EX, KY	P	P	P	D	–	P	P	P
Giaconi et al. [132]	2017	AN, CS, DE	LS	SH	P	D	D	D	P	D	P	P
Li et al. [56]	2017	AG, BI	HE, SMC	BL, QU	P	P	P	P	P	D	P	P
Ni et al. [47]	2017	AG, BI	HE	EX, BL, KY	P	P	D	P	P	D	P	P
Ni et al. [89]	2017	AG, AN	HE, DP, AUTH	QU	P	D	P	P	P	D	P	P
Shen et al. [80]	2017	AG	HE	EX, QU	D	P	P	P	P	D	P	P
Zhang et al. [152]	2017	BI, CS	DP, LS	SH	P	D	P	P	P	P	P	P
Abbasinezhad-Mood et al. [102]	2018	ID	SMC, AUTH	KY	D	P	P	D	P	D	D	P
Abdallah et al. [163]	2018	AG	HE	EX	P	P	P	P	P	D	P	P
Alsharif et al. [60]	2018	AG, BI	HE, SMC	EX, BL	D	P	P	P	–	D	P	P
Balli et al. [72]	2018	DR	HE, SMC, AUTH	TR	P	P	P	P	–	P	–	P
Farokhi et al. [145]	2018	CS	SMC, LS	SH	P	D	D	D	P	P	P	P
Gope et al. [59]	2018	AG, BI, ID	SMC, AUTH	EX, BL, QU	D	P	P	D	P	D	P	P
Li et al. [144]	2018	CS	LS	SH	P	D	D	D	P	P	–	P
Lyu et al. [172]	2018	AG	SMC, DP	AP	P	D	P	P	D	D	P	P
Sun et al. [142]	2018	AN, CS	LS	SH	P	P	–	D	P	D	P	P

Continued on the next page

TABLE XXIII: Compact study-level protocol matrix (continued)

Study	Year	Protocol scope			Assurance evidence				Operational evidence			
		App.	Mech.	Out.	Adv.	Priv.	Sec.	Corr.	Life.	Perf.	Eval.	Repr.
Architectural and functional integration (2019–2022)												
Alsharif et al. [113]	2019	AG, ID	FE, AUTH	EX	P	P	P	D	P	D	P	P
Alsharif et al. [51]	2019	AG, BI, AN	SMC, FE	EX, BL, QU	D	P	P	D	–	D	P	P
Chen et al. [166]	2019	AG, AN	HE	EX, QU	–	–	P	D	–	D	P	D
Gai et al. [116]	2019	ID, DE	AUTH, BC	KY, TR	P	P	P	P	–	D	P	P
Gope et al. [109]	2019	ID	SMC, AUTH, HW	KY	P	D	D	P	D	D	P	P
Hassan et al. [64]	2019	AG, BI, AN, DE	SMC, DP	EX, BL	P	P	–	P	–	P	P	P
Hossain et al. [150]	2019	BI, CS	AUTH, LS	SH	P	P	P	P	P	P	P	P
Liu et al. [156]	2019	AG, CS	HE, SMC	EX	P	P	P	D	–	D	P	P
Natgunanathan et al. [149]	2019	CS	LS	SH	P	D	P	D	P	P	P	P
Shateri et al. [118]	2019	AN, CS	ML	SH	P	D	P	P	–	D	P	P
Xue et al. [50]	2019	AG, BI, DR, AN	HE	EX, BL	P	P	P	P	P	P	P	P
Yao et al. [129]	2019	AG, AN	HE	EX, QU	P	P	P	P	–	P	P	P
Zhang et al. [103]	2019	ID	SMC, AUTH	KY	P	P	P	P	P	D	P	P
Zhang et al. [38]	2019	AG, BI, AN	HE	EX	D	P	P	D	–	P	P	P
Zhu et al. [44]	2019	AG, BI, ID	HE, AUTH	EX, BL	–	P	–	P	–	P	–	–
Fioretto et al. [167]	2020	OT	DP	QU	P	D	P	P	P	D	D	P
Gai et al. [168]	2020	AG, BI	DP	AP, QU	P	D	P	P	P	D	P	P
Gope [110]	2020	ID	AUTH, HW	KY	P	P	P	P	P	P	P	–
Ibrahim et al. [48]	2020	AG, BI	FE	EX, BL	P	P	P	D	–	D	P	P
Khan et al. [104]	2020	ID	AUTH	KY	P	P	–	P	–	P	–	–
Liu et al. [57]	2020	AG, BI, AN	HE, SMC	EX, QU	P	D	P	D	–	D	P	P
Ou et al. [120]	2020	AN, CS	DP	ML, SH	P	D	D	P	–	P	P	P
Saleem et al. [174]	2020	AG	HE, SMC	EX	P	P	P	P	P	P	P	P
Shateri et al. [119]	2020	AN, CS	ML	SH	P	D	P	P	–	P	P	P
Son et al. [136]	2020	BI, DR, DE	FE, AUTH, BC	BL, TR	P	D	D	P	–	D	D	P
Xue et al. [90]	2020	AG	HE, SMC	EX	D	P	D	P	D	D	P	P
Yu et al. [106]	2020	DR, ID	SMC, AUTH	KY	P	P	P	–	P	P	–	–
Chen et al. [160]	2021	AG	HE, SMC	EX	P	P	P	D	D	D	P	P
Chin et al. [146]	2021	CS	LS	SH	P	D	P	D	P	D	P	P
Ibrahim et al. [49]	2021	AG, BI, AN	FE	EX, BL, ML	P	P	P	P	–	D	P	P
Ibrahim et al. [99]	2021	AG, CS	HE, LS	AG	P	P	P	D	–	P	P	P
Kamil et al. [128]	2021	AG, AN	HE, AUTH	EX, QU	P	P	P	–	P	D	–	P
Khan et al. [169]	2021	AG, BI, AN	HE	AG	P	–	P	–	D	D	P	P
Lee et al. [135]	2021	DR, AN, DE	FL	ML	–	–	–	P	P	D	P	P
Liu et al. [170]	2021	AG, AN	SMC	EX, QU, KY	D	P	P	P	P	P	P	P
Liu et al. [153]	2021	AG, CS	DP, LS	SH	P	D	P	P	P	P	P	P
Merad-Boudia et al. [173]	2021	AG, AN	HE, SMC	EX, QU	P	P	P	D	P	D	P	P
Mohammadali et al. [81]	2021	AG	HE, SMC	EX	D	P	P	P	P	D	P	P
Qian et al. [162]	2021	AG	HE, SMC	EX	P	P	P	D	P	P	P	P
Sadhukhan et al. [105]	2021	ID	AUTH	KY	P	P	P	–	P	P	P	P
Singh et al. [175]	2021	AG, AN	HE, FL, BC	AG, ML	–	P	–	–	–	P	P	–
Vallent et al. [139]	2021	ID, EV	AUTH	KY	P	P	P	P	P	P	–	P
Wang et al. [83]	2021	AG	HE, SMC	EX, QU	P	P	P	P	P	P	–	–
Zhao et al. [52]	2021	AG, BI, AN	HE	AG, BL, QU	P	P	D	D	–	D	P	P
Zuo et al. [157]	2021	AG	HE	EX, QU	P	P	–	P	–	P	P	–
Badr et al. [124]	2022	AN	SMC, FE, FL	AG, ML	P	P	P	P	–	D	P	P
Fernández et al. [123]	2022	AN	SMC	ML	–	P	P	P	–	D	P	P
Gough et al. [65]	2022	AG, BI	DP	AP, BL, TR	P	D	P	P	–	D	P	P
Lin et al. [121]	2022	AN	FL	ML	–	–	–	D	P	D	P	P
Tran et al. [66]	2022	AG, BI, AN	SMC	EX, BL, ML	P	P	D	D	P	D	D	P
Wen et al. [125]	2022	AN	HE, SMC, DP, FL	EX, ML	D	D	P	P	–	P	P	P
Wu et al. [91]	2022	AG	HE, SMC	EX	P	P	D	D	P	D	D	P

Continued on the next page

TABLE XXIII: Compact study-level protocol matrix (continued)

Study	Year	Protocol scope			Assurance evidence				Operational evidence			
		App.	Mech.	Out.	Adv.	Priv.	Sec.	Corr.	Life.	Perf.	Eval.	Repr.
Zeng et al. [178]	2022	AG, BI	HE, SMC	EX, QU	P	P	–	P	P	P	–	–
Zheng et al. [63]	2022	AG, BI, CS	DP	AP, BL	P	D	P	P	P	P	D	P
Recent emphasis on richer guarantees (2023–2026)												
Zhang et al. [93]	2023	AG	HE	EX	P	P	D	D	P	D	P	P
Aurangzeb et al. [164]	2024	AN	BC, HW	ML	–	–	–	–	–	P	P	–
Baksh et al. [108]	2024	ID	AUTH	KY	P	–	P	P	P	P	P	–
Chang et al. [161]	2024	AG, AN	SMC	EX, QU	P	P	P	D	P	D	P	P
Chen et al. [58]	2024	AG, BI	HE, SMC	EX, QU	P	P	P	P	P	D	P	P
Fan et al. [155]	2024	AG	SMC, HW	EX, KY	P	P	P	P	P	D	D	P
Li et al. [84]	2024	AG	HE	EX, QU	P	P	P	P	P	D	P	P
Liang et al. [138]	2024	AG, ID, EV	SMC, AUTH, HW	EX, KY	P	P	P	P	–	P	P	P
Liu et al. [171]	2024	AG	SMC, AUTH	EX, KY	P	P	P	–	P	D	P	P
Manzoor et al. [122]	2024	AG, AN	FL	ML	–	–	–	P	P	P	P	P
Marandi et al. [53]	2024	AG, BI, AN	HE	BL, QU	D	P	P	P	P	P	P	P
Pang et al. [94]	2024	AG	SMC	EX, QU	P	D	P	P	P	P	P	P
Prateek et al. [141]	2024	ID, EV	AUTH, HW	KY	P	P	P	P	P	P	P	P
Rostampour et al. [111]	2024	ID	AUTH, HW	KY	P	P	P	P	P	P	D	P
Wu et al. [176]	2024	AG	HE, SMC	EX	D	P	D	P	D	D	P	P
Wu et al. [130]	2024	AN	HE	ML	P	P	P	P	–	D	P	P
Yu et al. [140]	2024	DR, ID, DE, EV	AUTH, HW	KY	P	P	P	P	P	P	P	P
Zhang et al. [179]	2024	AG	HE, SMC	EX	P	P	P	–	P	P	P	P
Zhang et al. [180]	2024	AG, AN	HE, SMC	EX, QU	P	P	P	–	P	P	D	P
Zhang et al. [181]	2024	AG	HE	EX	P	P	P	P	P	P	D	P
Zhao [159]	2024	AG	HE, AUTH, BC	EX, KY	P	–	–	–	–	P	–	–
Zhao et al. [67]	2024	AG, BI, AN	SMC, DP	AG, BL, ML	P	P	P	D	–	D	D	P
Habib et al. [107]	2025	ID	AUTH	KY	P	–	P	P	–	P	–	–
Li et al. [85]	2025	AG	HE, SMC	EX	D	D	P	P	P	D	D	P
Liu et al. [54]	2025	AG, BI	SMC	EX, BL	P	P	P	–	–	P	–	–
Liu et al. [114]	2025	AG	HE, FE	EX, KY	P	P	P	P	P	D	P	P
Liu et al. [158]	2025	AG, AN	SMC	EX	D	P	P	–	P	D	P	P
Sharma et al. [134]	2025	AN, DE	FL	ML	–	–	–	P	–	P	P	P
Yu et al. [95]	2025	AG	HE, SMC	EX	P	P	P	P	–	P	P	P
Zaman et al. [7]	2025	AN	DP	AP, ML	P	P	P	P	P	P	P	–
Zareidar et al. [68]	2025	BI	SMC, DP, LS	BL	D	P	P	D	P	D	D	P
Zareidar et al. [77]	2025	AG, AN	HE, SMC	EX	D	P	P	D	P	D	D	P
Zhang et al. [43]	2025	AG, BI, AN, ID	AUTH, HW	BL	–	–	P	D	–	P	P	P
Zhang et al. [182]	2025	AG	HE, BC	AG, QU	P	P	–	P	P	P	P	P
Zhang et al. [126]	2025	AG, AN	HE, FL, HW	AP, ML	P	P	P	D	P	D	P	P
He et al. [127]	2026	AN	SMC, DP	AP, ML	P	P	P	P	P	D	P	P
Li et al. [82]	2026	AG, AN	HE, SMC	EX, QU, KY	P	D	P	P	P	D	P	P
Liu et al. [92]	2026	AG, AN	HE	EX	P	P	P	P	P	P	P	P
Naeem et al. [40]	2026	AG, BI	HE, BC	EX, BL	P	P	D	D	–	D	P	P
Verma et al. [133]	2026	AG, DE	HE, SMC, AUTH	EX	P	P	–	P	P	P	P	P
Wang et al. [154]	2026	AG	HE, SMC	EX, KY	P	P	P	–	P	D	P	P
Wang et al. [115]	2026	ID	SMC, AUTH	KY	P	–	P	P	P	D	P	P
Xiao et al. [177]	2026	AG, AN	HE	EX	P	P	P	D	P	P	P	P
Zareidar et al. [131]	2026	DR, AN	DP, AUTH, LS	TR	D	D	D	D	P	D	P	P
Zeng et al. [55]	2026	AG, BI, AN	SMC, FE, AUTH	EX, QU	D	P	D	P	–	P	P	P

A. Protocol-Matrix Symbols and Codes

The symbols and codes used in Table XXIII are defined below. Evidence symbols assess the reported support for each dimension, not overall study quality.

Code	Definition	Code	Definition
Application codes		Evidence symbols	
AG	Aggregation/operations.	D	Detailed or adequate support.
BI	Billing, tariff, or settlement.	P	Partial, informal, or restricted support.
DR	Demand response or incentives.	–	Absent or seriously incomplete support.
AN	Analytics, forecasting, or detection.	Output codes	
ID	Identity, authentication, or access control.	EX	Exact aggregate.
CS	Customer-side privacy or data minimization.	AP	Approximate or noisy aggregate.
DE	Distributed energy, prosumer, trading, or shared storage.	AG	Aggregate whose exactness is not established.
EV	Electric mobility/V2G.	BL	Bill, tariff, or settlement.
OT	Other.	QU	Query or statistic.
Mechanism codes		ML	Model, forecast, classification, or control result.
HE	Homomorphic encryption.	KY	Authentication, access, or session-key result.
SMC	Masking, secret sharing, MPC, or distributed decryption.	TR	Trade, auction, allocation, or incentive result.
DP	Differential privacy or perturbation.	SH	Shaped, coarsened, or sanitized trace.
FE	Functional, attribute-based, or inner-product encryption.	OT	Other.
AUTH	Anonymous-authentication, credential, or pseudonym mechanism.	Evidence dimensions	
FL	Federated learning.	Adv.	Adversary and collusion model.
ML	Learned sanitization.	Priv.	Privacy definition or analysis.
BC	Blockchain or smart contract.	Sec.	Security analysis.
HW	PUF, TEE, TPM, or quantum-assisted hardware.	Corr.	Functional correctness.
LS	Load shaping, coarsening, or data minimization.	Life.	Lifecycle and robustness.
ENC	Other encryption-based construction.	Perf.	Performance reporting.
		Eval.	Implementation or empirical evaluation.
		Repr.	Reproducibility support.