INDIA ENERGY STACK

विद्युत मंत्रालय
MINISTRY OF
POWER

सत्यमेव जयते

INDIA ENERGY STACK

Architecture Document

VERSION 0.4

MARCH 2026

India Energy Stack (IES)

The Ministry of Power is reimagining the digital backbone for India's Power sector by creating a Digital Public Infrastructure (DPI) for the Power sector through the India Energy Stack (IES). The IES is being advanced under a whole-of-ecosystem approach through a phased programme of design, unified architecture blueprint, pilot implementation, and national rollout.

Published in March 2026 by the Ministry of Power

India Energy Stack (IES) Architecture Document

Recommended citation:

Ministry of Power (2026). India Energy Stack (IES) Architecture Document

***Note**: The India Energy Stack documents are currently in draft form, going through various revisions and will be released regularly as the idea evolves for public consultation and review. The contents of these documents should currently be seen purely as a draft.*

Version history

Date	Version	Description
24 Nov 2025	0.1	Initial draft version
17 Dec 2025	0.2	Second draft version
4 Feb 2026	0.3	Third draft version
27 Mar 2026	0.4	Fourth draft version

Table of Contents

Introduction

Electricity systems across the world are changing rapidly. Demand is rising because of electric vehicles, data centres, and smart appliances. At the same time, more renewable energy, such as solar and wind, is entering the grid, creating new challenges as their output varies with the weather. These trends mean the grid must become more intelligent, flexible, and secure. India is no different.

India's power sector is a large ecosystem with thousands of actors - DISCOMs, generators, system operators, OEMs, installers, service providers and millions of consumers and prosumers - all operating through physical and digital systems. Over the years, significant progress has been made - smart metering programs, outage management systems, billing platforms, forecasting tools, and renewable integration initiatives - these have created valuable digital assets and local capabilities.

Yet, most of these operate in isolation across states, utilities and vendors, resulting in fragmented data, inconsistent interfaces and limited interoperability. This slows down decision making, complicates coordination across sectors and creates barriers to integrate new services like EV charging, prosumer programs, demand response, and open access. As renewable energy, rooftop solar, EVs, batteries and prosumers multiply, the grid is shifting from a simple one-way system to a highly dynamic, two-way, high-transaction network. India's priorities - 24x7 reliable supply, 500 GW of renewables by 2030, and rapid electrification of transport and industry- require real-time visibility, trusted data exchange and seamless integration across millions of decentralized energy systems and assets.

The India Energy Stack (IES) offers a comprehensive set of program, policy, architecture, and set of specifications that enables these actors and systems to interact in a trusted, consistent and verifiable manner. Instead of replacing existing systems, IES provides common protocols, schemas (data models), standards, and interaction patterns that connect these actors - allowing the sector to make faster decisions, drive innovation, reduce friction, and move towards a dynamic future.

For full context, refer to the India Energy Stack (IES) Strategy Document.

Purpose of this document

This Architecture Document translates the India Energy Stack (IES) strategy into an architectural blueprint. It defines the core building blocks, protocols, data models, interaction patterns, and other supporting technology elements that will enable interoperable, trusted, and seamless interactions (transactions and data exchanges) across various actors and systems within India's power sector. While the Strategy Document explains why IES is needed and what it must achieve, this document explains how those goals are realised as a Digital Public Infrastructure.

Alignment with Strategy Document

This document should be read as one of the three “AAA” tracks described in the Strategy Document: Architecture, Adoption, and Accelerator.

- **IES Architecture**: Defines the overall architecture of the IES ecosystem, which includes the architecture blueprint, core building blocks, protocols and supporting specifications to enable transactions and data exchanges between various actors within the ecosystem.
- **IES Accelerator**: Implements sandbox environments, tools, reference solutions, and other ecosystem enablers as necessary to accelerate adoption of IES across the ecosystem.
- **IES Adoption Strategy**: Identifies a set of program and policy initiatives to incentivise and encourage adoption by various entities/stakeholders within the IES ecosystem.

Wherever relevant, this document references sections of the Strategy Document (e.g., Vision & Mission, Strategic Objectives & Core Use Cases, Roadmap & Phasing) and should remain consistent with it.

What is IES and what is it not?

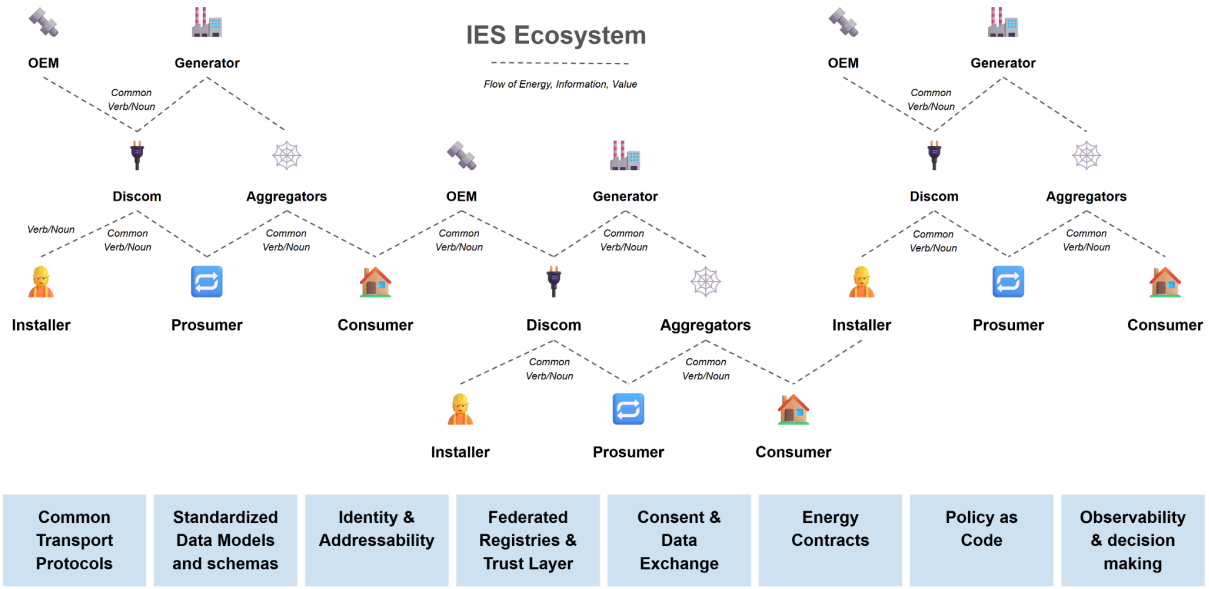

An illustrative set of IES building blocks

Scope of IES

IES is...

- **A set of protocols/specifications** that makes **interoperability** between grid entities/stakeholders **uniform, reliable, verifiable and trustworthy**. Essentially, IES defines how entities (DISCOM-Apps; DISCOM-GENCO; DISCOM-Transco; Genco-Transco; EV charger- Charger booking apps; consumer-producer etc.) interact with one another
- **A set of services** (API definitions and calls) **plus a data model/taxonomy** (how components are identified, organised and relate to each other, i.e., how the specified services are used together).

Therefore, IES is a **foundational digital layer** combining a common interaction framework, standardised **taxonomy/data models**, **assurance** mechanisms for **trust and validation**, and baseline technical standards. This enables the development of scalable solutions for the energy sector.

IES is NOT…

- A centralised database/data lake.
- **A centralised service** that pulls or receives data from various entities.
- A software package.
- **Dependent on a strict hierarchy.** It does not presume a hierarchical structure of the energy sector to function.
- A set of specifications and standards whose sole purpose is to integrate internal systems of any stakeholder/entity/utility.

The following diagram illustrates this:

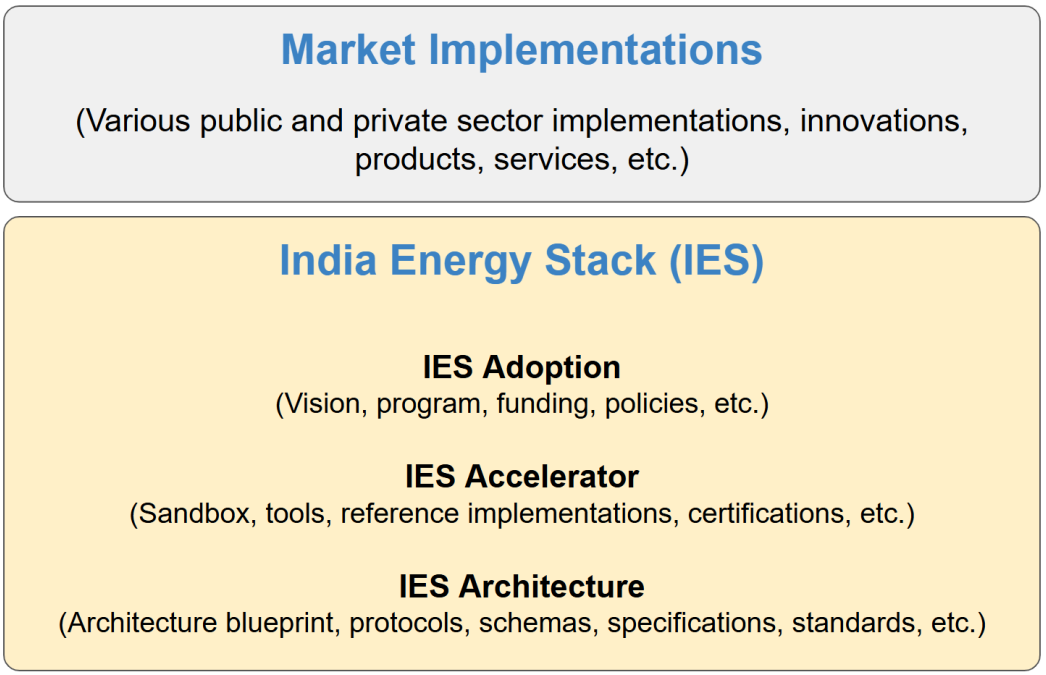

Stakeholders & Users of IES

Building on the strategy document, IES recognises the following primary stakeholder categories and their roles in the architecture:

- **Central Government** – MoP, CEA - Policy direction, regulatory oversight, and stewardship of IES standards and governance processes.
- **System Operators** – POSOCO/GRID-India, RLDCs/SLDCs - Operation of critical system registries and services relevant to scheduling, dispatch, and system reliability; integration of IES into operational planning and control workflows.
- **Utilities & DISCOMs** - Frontline adoption of IES specifications; operation of local/federated registries (e.g., consumer, asset); exposure and consumption of APIs; integration with legacy IT/OT systems.
- **Private Innovators & Market Participants** - Development of applications, analytics, and services that consume and publish IES-compliant data; implementation of provider or consumer platforms that participate in IES interactions.
- **Consumers & Prosumers** - Active participation in energy markets and programs through IES-compliant applications, with control over their data and consent.
- **Think Tanks/Research/Capacity Building Institutes** - upskill DISCOMs, system operators, and market participants on IES standards and tools.

IES Architecture

Architecture Principles

Following are the architecture principles of IES:

1. **Federation and decentralization**: Energy assets and responsibilities are widely distributed, so data and control should remain close to their sources while still being easily discoverable and interoperable across the ecosystem. Architecture should be designed with decentralization at its core to prevent dependency on central authorities, ensure resilience, eliminate single point of failure, avoid monopolization, and facilitate wider participation.
2. **Identity and addressability**: Ensuring every actor and asset has appropriate identifiers and addresses to ensure mutual trustworthiness, transaction traceability, contract adherence, and application of various policies and rules.
3. **Minimal and generalized:** Principle of minimalism is critical to IES as the core remains lean and context invariant across sector and use cases. IES will standardise only what must be common at ecosystem boundaries—interface contracts, core schemas, registry architecture, and trust assurance hooks—while avoiding over-prescribing internal implementations. This is key to ensure IES is stable and generic, allowing many use cases, extensions, and context aware solutions to be built by the ecosystem.
4. **Unified, not uniform**: The infrastructure should promote multiplicity without imposing uniformity or a one-size-fits-all approach. Rather than a single standard or centralized solution, multiple implementations and standards can coexist, as long as they conform to the specifications adopted under IES, allowing institutions to build or procure solutions suited to their context while still participating in a unified ecosystem. This also allows existing technology investments to be upgraded to comply with IES specifications without having to rebuild from scratch.
5. **Inclusive and diversity enabling**: Given the diversity of the Indian ecosystem and consumer base, the architecture must accommodate a diverse range of personas and contexts. This includes solutions tailored for users and institutions with high degree of formalisation and digital readiness as well as those operating in resource-constrained environments.
6. **Trusted and secure**: Security, confidentiality, and privacy preservation are fundamental. Data is treated as a sensitive resource, its capture and transmission are carefully calibrated, its use and sharing are driven by explicit rules and consent, and its capture is based on the principles of “optimal ignorance”. All system interactions ought to be encrypted, signed using cryptographic techniques and privacy-preserving technologies to reinforce trust across the ecosystem.
7. **Open**: The entire digital backbone must avoid being closed-loop, proprietary, or restricted by bilateral or multilateral contracts, and should not be hardwired

to geographic regions. And should instead be an unbundled, loosely coupled, and universally accessible open-loop architecture. IES will leverage and reuse existing standards, components, governance mechanisms, and implementation patterns to accelerate adoption.

8. **Interoperable**: A central goal is a loosely coupled ecosystem enabling seamless movement of information across actors. Open protocols and specifications should enable seamless exchange across actors and allied sectors, while allowing multiple vendors and solutions to compete and interoperate without lock-in or geographic hardwiring. IES will also leverage and interoperate with existing investments in the power sector as well as various DPI building blocks wherever possible rather than reinventing.
9. **Innovation friendly**: The architecture must promote market innovation by encouraging public and private stakeholders to create tailored solutions, applications, products, and services that meet the diverse and evolving needs of their respective contexts. It should also be AI-ready by design—making high-quality, consented data and event signals interoperable and machine-consumable (with standard semantics and provenance), so AI systems can be developed and deployed safely without creating new central dependencies or lock-in.
10. **Inclusive and asynchronous adoption**: IES will support phased, non-disruptive adoption allowing participants to adopt and implement in stages, upgrade independently, and coexist across versions without synchronization across the rest of the ecosystem.

Technology Trends for Architectural Considerations

IES architecture is being designed in a period of sustained technological change. While specific applications, programs, and market models will evolve, several underlying trends are likely to persist over the lifetime of the infrastructure. These trends shape the design of IES to ensure it is designed for the future rather than for the past.

Software-Defined Infrastructure

A fundamental trend across sectors is the increasing softwarisation of physical infrastructure making these physical assets discoverable, observable, and programmable. Sensors, embedded computing, APIs, and remote control systems are becoming integral to how physical assets are monitored and managed. In the power sector, this is visible in smart meters, digitally controlled inverters, storage systems, and software-driven grid operations. IES will be designed to take advantage of this trend.

Increasing Machine Intelligence

Advances in artificial intelligence and machine learning are increasingly influencing human-computer interfaces, productivity, prediction, optimisation, and operations. While AI driven systems do not replace policy or accountability, they shape how decisions are informed and executed. IES will be designed to be an AI and agentic world ready from the get go through its programmable and machine-readable core. Data models, policies, and interactions are machine represented to support automated evaluation, while ensuring that decisions—automated or assisted—leave clear, verifiable audit trails.

Decentralization and Unbundling

Energy systems are no longer dominated by a small number of large producers and operators. They increasingly include millions of distributed participants—prosumers, aggregators, devices, and micro-systems—each generating data and participating in transactions. In addition, large vertical systems are being unbundled into many smaller systems operated by different players. This shift introduces scale, heterogeneity, and interoperability needs that cannot be addressed through bilateral integrations, traditional IT approaches or centralised controls. IES responds by emphasising federated registries, common protocols and specifications, verifiable identifiers, and extensive use of cryptographic techniques to ensure smooth and trusted transactions and coordination across the system.

Rise of High-volume, Low-value Interactions

Another structural change is the move from infrequent, bulk transactions to continuous, granular interactions. India saw this with systems like UPI where low-volume, high-value systems had to be reimagined for a high-volume, low-value environment. Pricing signals, demand response events, flexibility services, and distributed trades generate large volumes of small, time-sensitive interactions. Such environments favour event-driven architectures, deterministic data models, and strong observability. The IES architecture therefore prioritises asynchronous interactions, explicit references, and auditable event flows over tightly coupled transactional workflows.

Mobile-First Participation

For many participants, especially consumers, installers, and field staff, mobile devices are the primary, mostly only, interface to digital systems. Participation in energy programs, access to data, and consent management increasingly occur through lightweight clients operating over variable connectivity. This reinforces the need for thin clients and robust backend trust mechanisms. IES avoids assumptions of heavy enterprise software at the edge and instead supports asynchronous, API-driven participation through mobile and other light-weight interfaces (e.g. chat).

Pervasive Connectivity

Today connectivity is nearly pervasive and will continue to improve in the future. IES must take advantage of this trend and take energy systems towards a connected and coordinated environment. While the new age digital infrastructure is increasingly being built to take advantage of connected systems, many parts of it still work under conditions of heterogeneous, intermittent, and uneven connectivity. Energy systems in particular operate across dense urban networks, remote rural areas, industrial zones, and critical infrastructure environments, each with different latency, bandwidth, and reliability characteristics. Connectivity can range from high-availability fibre networks to mobile, satellite, or occasionally disconnected systems at the edge. IES is therefore designed with the assumption that connectivity cannot be taken for granted. The architecture favours asynchronous, event-driven interactions in general over tightly coupled synchronous workflows unless when necessary, supports store-and-forward patterns, and avoids dependencies on continuous central availability. Protocols, data models, and trust mechanisms are designed to function reliably even under partial connectivity, enabling systems to operate locally and reconcile state when connectivity is restored. This approach ensures that IES remains usable across diverse geographies and operating conditions, and that participation in the energy ecosystem is not constrained by network quality alone.

Privacy, Consent, and Trust as Design Constraints

As public digital infrastructure expands, trust becomes a first-order concern. Concerns around data misuse, opacity, and overreach can undermine adoption even when systems function correctly. IES therefore adopts principles of minimalism, purpose limitation, decentralization, consent-driven access, and auditability. Trust is treated as an architectural outcome rather than as a feature added later.

Evolution of Cryptography

Cryptographic techniques have evolved significantly over the last decade to provide verifiability, mutual trust, encryption, immutability, transactability, and portability. Digital signatures, end-to-end encryption, verifiable credentials, blockchains, tokenization, zero-knowledge proofs, etc are now providing capability to ensure security and trust across digital systems without having to create verification models or human interventions or process driven assurances, all of which result in high cost and becomes non scalable. IES therefore emphasises cryptographic methods to drive agility, trust anchors, verifiability, confidentiality, and auditability for a low-value, high-volume future across a large number of decentralized systems.

Illustrative Flows

Modernising a cyber-physical sector like energy requires applying these architectural principles consistently across the core flows that shape the overall electricity system: **information, services, trust, and money.** In such an environment, digital interactions and physical grid behaviour continuously influence one another. For the system to work at a population scale, these flows must operate on shared rules without centralising data or redesigning internal systems. The aim is to define common interaction boundaries so that every participant can coordinate safely while retaining full control of their local operations.

The **flow of information** like meter data, sensor data from transformers and switchgears, SCADA, and outage management systems. When expressed through common data models, taxonomies, and machine-verifiable formats, this information becomes reliable and composable across utilities, operators, service providers, and devices. The **flow of services** refers to the flow of energy and allied services across the various actors and stakeholders in the grid ecosystem. These include new connection, demand response, ancillary services and various market related services. These actions must follow predictable handshake rules so that service orchestration remains coherent, safe, and extensible across stakeholders. These flows of services are expected to extend beyond energy services to financial services, commerce services like buying and installation, skilling and maintenance services (e.g. installers and service maintenance professionals).

The **flow of trust** – identity, digital addresses, credentials, digital signatures, provenance, and auditability – is what enables a decentralised ecosystem to function without bespoke bilateral agreements. Trust must be embedded directly into each interaction through verifiable credentials, digital proofs, and policy-attached validations, ensuring that every actor and asset can be authenticated, verified, and held accountable. The **flow of money** connects these interactions to economic finality: billing, settlement, incentives, flexibility payments, and financial guarantees. These depend on consistent, tamper-evident information and standardised transaction structures.

These flows are informed and constrained by, and not independent of, the underlying grid critical services, technology and infrastructure towards management and resilience of the physical grid

Together, these flows, recognising the constraints, require a set of **universal building blocks** that allow the system to scale safely from a few actors to millions of assets and applications. By applying the architectural principles across all flows, the ecosystem can support rapid innovation at the edges while preserving the reliability, stability, and coherence of the underlying physical grid.

The following are the list of core building blocks of IES, which are to be seen as a necessary not sufficient, for enabling at-scale coordination between the actors and systems in the energy sector:

1. **Identity and addressability**
2. **Registries and Directories**
3. **Energy Credentials**
4. **Policies and Data**
5. **Protocols and schemas**

Identity and addressability

For IES to behave like a unified fabric rather than a set of bilateral integrations, the ecosystem needs a shared way to know who and what is involved in any interaction. Identity and addressability answer two questions: who is the actor? (party, system, device) and what is the resource or asset? (meter, feeder, DER, contract, data product). IES therefore needs identifiers at multiple levels: legal entities (DISCOMs, generators, regulators, aggregators, OEMs), the platforms and applications they operate, physical and virtual assets in the grid, and logical resources such as Energy Resource Addresses where energy is injected, withdrawn, or measured.

At the party level, IES envisages globally unique, registry-backed identifiers for organisations and (where needed) individual roles. These IDs are anchored in authoritative registries and bound to verifiable credentials encoding licences, roles, and permissions. A DISCOM, for example, may run many systems and APIs, but they all present a common IES Party ID; likewise, an aggregator or OEM is recognised consistently whether interacting with a DISCOM, SLDC, or market operator. Separating “who you are” (party identity) from “what you run” (endpoints, applications) enables intent-based routing, access control, and audit to be applied uniformly.

At the asset and resource level, IES encourages uniform identification of grid assets across the ecosystem. Generators, transformers, lines, substations, meters, batteries, EV chargers, rooftop solar systems, and virtual portfolios should each have a stable, unique identifier that can be referenced across systems and over time, often grouped or exposed via Energy Resource Addresses. Today, a single asset may have multiple internal IDs—one in GIS, one in billing, one in the OEM portal; IES instead allows a common “identifier mechanism” that all parties can use to refer to a system/actor/asset when exchanging data or enforcing contracts.

A consistent identifiers and addresses are needed to ensure:

- **Traceability**: unambiguous reference to the same asset across databases and time.
- **Interoperability**: seamless data exchange between DISCOMs, SLDCs, NLDC, regulators, market operators, and consumers.
- **Lifecycle management**: tracking installation, ownership, operation, decommissioning, and recycling against a single identity.
- **Trust and audit**: preventing duplication/spoofing and enabling authentic markets for RECs, carbon credits, and flexibility services.

Consent:

Consent must be **granular, purpose-specific, time-bound, and revocable**. A prosumer sharing meter data with an aggregator for demand response enrollment is not granting blanket access to all their information forever. The consent artefact must specify: *what* data is being shared, *with whom*, *for what purpose*, *for how long*, *under what conditions* it can be used or further shared, and its revocation mechanism.

IES recommends that such consents are captured as a machine-readable, digitally signed artefact. It travels with the data, enabling downstream systems to verify that appropriate permissions exist before processing. If consent expires or is revoked, access stops. IES aligns with MeitY defined electronic consent artifact and adapts this appropriately.

Registries

Registries and verifiable credentials enable IES to function as a “trust bridge” across the sector. They give all actors a shared, reliable way to prove who they are, validate information, and verify actions—without centralising data or control. For example, if a rooftop solar installer holds a verified license, a DISCOM or aggregator can instantly validate it during DER onboarding. This reduces fraud, eliminates manual checks, and speeds up prosumer onboarding

As energy markets decentralize and distributed energy resources become more affordable, millions of small-scale participants - rooftop solar owners, community battery operators - can join the formal energy ecosystem. Unlike large institutions whose claims can be audited at scale, these small actors lack cost-effective verification mechanisms, creating a trust challenge: how to verify that "green" energy is genuinely solar-sourced, or that EV-charging subsidies reach intended recipients. This trust infrastructure relies on two complementary components: registries and credentials.

Registries hold authoritative, machine-readable, cryptographically secured information that serves as the root of trust. Registries can be private and public. Registries are not centralized databases controlled by a single entity. Instead, they

are managed by whichever authority—public or private—is responsible for managing a particular registry. Attestation and registration can happen in a federated manner across multiple nodes.

Private Registries

Private registries are registries that contain private/sensitive data. The data within such registries are typically not public by design and is available to 3rd parties only via authorized and consented access flows.

Core Capabilities:

- **Configurable schemas** with automatic API generation for seamless integration with existing utility systems (CIS, GIS, MDMS)
- **Attestation and verification flows** for validating claims (e.g., a DISCOM verifying an installer's qualifications)
- **Consent and authorization management** for data access and verification
- **Privacy preservation** through digital signatures and encryption for PII protection
- **Scalability** to handle millions of consumers and assets with asynchronous updates

Examples: Consumer registries, DER asset registries, installer/contractor registries, meter registries, etc.

Public Registries (or Directories)

Also known as directories, public registries contain information that should be openly accessible to all ecosystem participants without any need for authorization/consent and require universal discoverability, cryptographic verifiability, and real-time update propagation.

Core Capabilities:

- **Universal API specifications** (lookup, query, search, watch) eliminating need for bilateral integrations
- **Cryptographic proof of authenticity** providing tamper-evident verification
- **Decentralized availability** with no single point of failure
- **Trust chain propagation** where organizations formally recognize each other's registries (e.g., CEA recognizing SERC registries)
- **Real-time revocation propagation** ensuring license suspensions and certification expiries are instantly reflected across all connected systems
- **Agent and AI readiness** through machine-readable formats for automated verification

Examples: Lists of certified solar panel manufacturers, approved charge point operators, public keys of participating entities, revoked business license grid interconnection approvals, accredited certification bodies, tariff schedules, etc.

Energy Credentials

Today, critical information about energy actors and assets is trapped in institutional silos. DISCOMs hold consumer connection data. Certification bodies hold equipment compliance records. Banks hold creditworthiness assessments. Government agencies hold subsidy eligibility information. Regulators hold license status.

When a third party needs to verify any of this information, they face two problematic paths: bilateral API integrations (costly, not scalable with n × m connections, creating walled gardens where data holders control access) or manual verification (paper documents, phone calls, physical visits—slow, expensive, and fraud-prone).

In this institution-centric model: institutions issue, institutions store, institutions decide when you can access your own information. The prosumer, aggregator, or asset owner waits, requests, and depends on institutional gatekeepers. This fundamentally limits the speed, scale, and inclusivity of energy market participation.

The User-Centric Alternative: Verifiable Credentials

Verifiable credentials allow us to flip this model: institutions issue, you hold, anyone can verify without n:m integrations.

A verifiable credential is a machine-readable, cryptographically signed attestation that can be used to make verifiable claims. Any data an institution holds about an actor or asset can be credentialized—turning static database records into portable, tamper-evident proofs that the holder controls and can present to any verifier, anytime, without requiring the issuer's involvement.

In the energy sector, credentials can represent green energy certification, grid interconnection approvals, safety and standards compliance, ownership and property rights, subsidy eligibility (linked to government DBT programs), payment history and creditworthiness, capacity and performance attestations, maintenance and inspection records, transaction history and reputation, renewable energy certificates, carbon intensity verification, and insurance coverage—essentially any attestation that today exists as a paper certificate, database entry, or approval letter.

How It Works

Trusted entities (regulators, utilities, certification bodies, testing laboratories, private institutions) issue credentials by cryptographically signing attestations and delivering them to holders. Holders store credentials (e.g, in digital wallets) and present them to relying/accepting parties during transactions. The relying party checks the

credential's cryptographic integrity against the issuer's public key (available in public registries) and also confirms revocation status—all without contacting the issuer. If a credential needs to be revoked, the issuer updates the revocation registry and the status propagates instantly across the ecosystem.

Verifiable credentials also support selective disclosure—holders can prove specific claims without revealing full credentials. An aggregator proves they hold a valid state license without disclosing complete compliance history. A prosumer proves subsidy eligibility without revealing full consumption data.

Energy credentials operate within a trust framework built on the registry infrastructure described earlier: public key registries where verifiers look up issuer keys to validate signatures, revocation registries for real-time propagation of suspended or expired credentials, and approved issuer lists identifying authorized certification bodies and recognized authorities. This combination—registries as the root of trust, credentials as portable proofs—enables trustworthy data sharing at scale without centralized control, manual audits, or institutional gatekeeping.

Credential Revocation

A credential's lifecycle does not end at issuance. Licences may be suspended, certifications expire, equipment can be decommissioned, connections are transferred, and so on. The ecosystem must trust not only that a credential was validly issued, but that it remains valid at the moment of presentation/use.

Issuers maintain revocation registries — authoritative, machine-readable records of credentials that have been suspended or revoked. These registries that contain this information should be made available as public directories in a privacy protecting manner alongside the issuer's public key. This ensures any verifier can check revocation status as part of the standard verification flow publicly without engaging with any bilateral contract with the issuer.

Three principles govern revocation in IES. Revocation status must be made available without significant time delay (ideally in near realtime) so that revoked/suspended credentials are not used by credential holders. Every relying party (entity that is accepting a credential as part of their workflow) ideally should check revocation status as part of credential verification.

IES supports all widely available schemes for credential revocation and issuers can choose to use any open URL that is made available within the credential to allow relying parties to check status. Conformance requires that the chosen mechanism is publicly documented, machine-queryable, and referenced in the credential's metadata so any verifier can locate and query it.

Digital Energy Contracts

Digital energy contracts are typically the result of a transaction between two or more actors. Digital contracts in IES provide a machine-readable way to formalise agreements, commitments, and obligations between actors in the energy ecosystem. They capture the operational and economic terms governing interactions such as demand response participation, flexibility services, peer-to-peer transactions, open access arrangements, EV charging, and subsidy-linked programmes. By expressing these terms in standardised, computable formats, IES enables consistent interpretation and execution of contracts across systems operated by DISCOMs, aggregators, market platforms, OEMs, and financial institutions.

Unlike traditional paper contracts, digital energy contracts are designed to be machine-readable and tamper-evident by default. This allows contract parameters to be verified and processed by software and AI techniques resulting in faster, cheaper processes as well as easier dispute resolution.

Key elements including contracting parties, referenced assets, identifiers/addresses, validity periods, capacity or quantity commitments, pricing structures, performance conditions, and penalties are represented using structured schemas that systems can evaluate automatically. This allows contracts to be validated at runtime during discovery, scheduling, dispatch, settlement, and audit, ensuring transactions conform to agreed terms while remaining traceable and auditable.

Digital energy contracts are also portable across platforms and jurisdictions. They are not bound to specific vendor systems or marketplaces but reference shared identifiers, registries, and schemas defined under IES. As a result, the same contract can be recognised by utilities, market operators, and regulators without re-encoding or reinterpretation, reducing lock-in and enabling competitive participation.

Digital energy contracts work alongside machine readable policies. While such policies define the rules and constraints that govern what is permissible, digital contracts encode consent, commitments, and accountability for specific participants and transactions. Together, they enable scalable, verifiable, and enforceable energy transactions across a distributed ecosystem.

Machine Readable Policies

Machine-readable representations of policies, rules, and constraints that govern energy transactions and system operations enable both humans and automated agents to understand, interpret, and apply regulatory frameworks, business rules, and operational guidelines consistently across distributed networks.

Unlike traditional policy documents written in natural language, machine readable policies express rules in structured formats (JSON, YAML, XML) that systems can

parse and validate against, ensuring automated compliance checking throughout the transaction lifecycle.

Importantly, a machine readable policy object is not the entire policy itself—it abstracts only those elements that machines can interpret and execute within appropriate flows. The underlying policy may contain context, rationale, exceptions, and interpretive guidance that remain in natural language. This captures the enforceable rules that systems can evaluate deterministically.

The authorship of a machine readable policy object can be separate from the policy signatory. A regulator issues a policy, but the machine-readable encoding may be authored by a technical body, industry association, or any entity with the expertise to faithfully translate regulatory intent into computable form. This separation allows domain experts to handle codification while regulators retain authority over the policy itself. **These machine readable policies are published—onto public registries—for anyone to consume, and are versioned independently from the underlying policy.** For instance, a tariff regulation may remain unchanged while its machine representation undergoes revisions to fix encoding errors, improve precision, or add support for new data formats. Clear versioning ensures that all parties know exactly which encoding they are validating against, and that updates propagate transparently across the ecosystem.

Policies can take many forms and serve diverse use cases: demand response program rules defining eligibility criteria, event triggers, and compensation formulas that participating entities and systems can read to determine qualification and response actions; dynamic pricing policies encoding time-of-use tariffs, surge pricing triggers, and seasonal adjustments that marketplaces automatically apply; geographic and resource-specific constraints like "*only allow P2P trading within 5km for residential prosumers with valid green certificates*" or "*EV charging reservations must be between 1-100 kWh*" that discovery and selection APIs validate automatically etc. By expressing these policies as structured data IES enables consistent interpretation, transparent verification, dynamic updates without code changes, policy composition from reusable building blocks, and hierarchical policy application from national regulations down to platform-specific rules - creating guardrails within which decentralized energy transactions can occur compliantly, efficiently, and at scale.

Data Exchange

Systems, entities and actors in the power sector generate data continuously — regulatory filings, tariff policies, consumer records, meter readings, DER registrations, flexibility transactions. Today, most of this data moves through bilateral integrations, bespoke templates, or not at all. The result is fragmented information, weak auditability, and friction that compounds at every sector boundary.

IES takes a different position: data exchange is not a collection of point-to-point problems to be solved use case by use case. It is a single structural problem — structured data moves between authorised parties, requires discovery, must be validated and attested, and produces verifiable receipts. The only things that change are who can see it, how it moves, and what trust proof is attached.

This leads to a unified data exchange architecture organised around three categories.

Public Data Exchange

Public data is discoverable by anyone, verifiable by anyone, and hosted at source. Custodians publish metadata records to cryptographically verifiable, decentralised directories (DeDi). Any participant — a DISCOM, an aggregator, a developer, an AI system — can discover what exists and fetch directly from the custodian's endpoint without bilateral integration or prior relationship. There is no central repository; discovery is global, data stays local.

What lives here: published tariff policy packs, approved participant directories, public regulatory disclosures derived from accepted filings, DER registry entries, revocation lists, certified equipment catalogs, and any data a custodian designates as open. Trust is established through content hashes, issuer signatures, and version chains — any party can verify integrity without contacting the issuer.

A representative flow: a DISCOM publishes an attested tariff policy pack to the public channel. A market platform discovers it via directory lookup, fetches it directly from the DISCOM's endpoint, verifies the issuer signature against the public key registry, and confirms the version chain. The same pack is independently consumed by an aggregator evaluating demand response eligibility and by a regulator auditing applied tariff versions — all without any of them coordinating with each other or with a central system.

Private Data Exchange (System-to-System)

Private data is access-controlled, licensed, optionally priced, and transacted through a standard asynchronous protocol between systems. This is the channel for institutional, system-to-system data exchange — a regulator requesting a DISCOM filing, a researcher requesting aggregated data, an aggregator requesting operational telemetry, or an energy sector startup accessing data for providing value added layers. The same interaction grammar applies regardless of what is being exchanged.

IES defines a packet envelope with correlation IDs, an immediate acknowledgement and async payload delivery via paired callbacks. Data does not travel through the protocol packet. The response carries a payload digest (for integrity verification) and

a payload url (a time-limited, dynamically generated, authorisation-enforced URL from which the requester fetches the data directly from the custodian's endpoint or in future requester accesses a secure data enclave of the provider).

This architecture follows a decentralised data management model: every data provider retains full autonomy and agency over their data. The provider decides how they share, with whom they share, and under what conditions they share — all managed at their own endpoint. There is no central database where data is collected, stored, or routed through. The protocol defines the handshake; the data stays where authority lies.

Access is governed by credentials of the requesting party. A regulator presents its role credential to access DISCOM filings. A researcher presents a purpose-scoped credential to access anonymised consumption data under defined terms. An aggregator presents an accreditation credential before receiving dispatch telemetry. If a credential is invalid, expired, or outside the scope of the request, the interaction fails at the protocol boundary before any data moves.

A representative flow: a State Electricity Regulatory Commission sends a standardised data request to a DISCOM's IES-compliant endpoint, presenting its regulator credential. The DISCOM's system issues an immediate ACK with a correlation ID. Within the agreed response window, the DISCOM's system assembles the filing — validated against the IES canonical schema, signed, and timestamped — and delivers it to the SERC's callback endpoint. The SERC's system verifies the signature, stores the receipt binding both parties, the content hash, schema version, and outcome status. The filing is now auditable, traceable, and disputable by reference to that receipt — without either party having surrendered control of their system of record.

For data shared with third parties (researchers, policymakers, market analysts), the same principles apply: data remains at source, the contributor retains control, access terms are codified per accessor class, and every transaction produces a receipt.

Consumer Data Exchange (via Credentials)

Data that needs to be delivered to the consumer — consumption history, meter data, billing records, consumer profile, connection details — follows a fundamentally different path. Consumer-facing data is not exchanged system-to-system through the private data exchange channel. Instead, it is issued directly to the consumer as a set of verifiable credentials.

DISCOMs and other authoritative parties issue these datasets as energy credentials (W3C Verifiable Credentials) — cryptographically signed, machine-readable attestations delivered directly to the consumer through their portal, app, DigiLocker push, or equivalent interface — at regular intervals (e.g., aligned to the billing cycle)

or on consumer demand. The consumer holds these credentials in their wallet and can present them selectively to any relying party — an aggregator, a lender, a rooftop solar installer — without requiring the DISCOM's involvement at the point of presentation.

This is the credential-issuance-first model: the default path for consumer data is issuance to the consumer, not API-based gated access via third parties. An API endpoint for the same data is optional; DISCOMs that choose to offer consumer data via data exchange APIs (as explained in previous section) to support a third-party app ecosystem must obtain consent of the consumer before releasing data (this is directly between data provider and consumer). A credential reference alone is insufficient to authorise release: credentials are bearer artefacts and do not prove the presenting party is the data subject. Acceptable off-channel mechanisms include OTP, app-based challenge, or equivalent. If the API path is implemented, the request must also include a trusted way to validate consent if required in audit, a form of consent artefact specifying what data, with whom, for what purpose, for how long, and its revocation mechanism.

For the full specification of how verifiable credentials work — issuance, presentation, verification, selective disclosure, and revocation — refer to the Energy Credentials section of this document.

The Common Layer

All three categories share a common structure — every data object, whether a tariff pack, a regulatory filing, a DER registry entry, or a consumer credential, uses well structured protocols and data schemas. All carry JSON-LD semantics so data objects are self-describing and interpretable across systems without shared databases.

This approach ensures IES based systems remain federated and connected via IES specifications rather than through centralized databases or platforms. IES defines the envelope, the semantics, the interaction patterns, and the conformance requirements. Not the storage, not the systems, and not the decisions.

Security and resilience

India’s energy system is a unique convergence of operational technology (OT)—generation assets, substations, SCADA, AMI, DER controllers—and modern information technology (IT) layers that enable coordination, analytics, transactions, and market operations. The India Energy Stack (IES) must therefore uphold a security posture that protects both real-world physical assets and the digital coordination fabric that connects millions of actors.

Security and resilience by design

IES follows a “secure-by-default, resilient-by-design” approach, where every building block—identity, registries, credentialing, APIs, event streams, data flows, and contracts—is protected through layered security, cryptography, and governance. The architecture assumes:

- **Zero Trust** across utilities, vendors, DER providers, and market participants
- **Least privilege** for all actors and data flows
- **Defense in depth** across stack layers (device → edge → network → cloud → application)

Core security capabilities envisioned in IES

The initial security envelope includes:

1. **Strong, verifiable digital identities** for assets, actors, devices, and organizations
2. **Cryptographically verifiable credentials** for onboarding, licensing, compliance, device certifications, and audit trails
3. **Secure API gateway patterns**, including signed requests/responses, mutual TLS, and nonce-based replay protection
4. **End-to-end data confidentiality and integrity**, especially for telemetry, control signals, demand response events, and settlement data
5. **Event-level provenance**, ensuring all data used for settlements, forecasting, market operations, and grid balancing is tamper-evident
6. **credential-based authorization chain**, linked to registries

In addition, various deployment level security (network security, systems hardening, access controls, certifications, etc), use of Security Operations Centers (SOC), deployment of detailed incident response protocols, processes, audits, etc are applicable for actors implementing IES.

Leveraging Cryptography

IES leverages modern cryptographic techniques to ensure data integrity, authenticity, confidentiality, and privacy across the ecosystem. Advances in cryptography have made it possible to embed strong security guarantees at low cost, and without a central intermediary or manual verifications enabling even resource-constrained devices and high-volume transaction flows to participate securely.

Key cryptographic leveraged by IES include:

1. **Encryption:** IES ecosystem will employ encryption to protect sensitive data both in transit and at rest. Transport-layer protocols like TLS secure API calls, telemetry streams, and information exchanges, while storage-layer encryption

will protect meter readings, billing records, and contract terms from unauthorized access.

2. **Secured APIs and Transport**: All IES APIs will be secured using standard cryptographic protocols including TLS for transport encryption, mutual TLS (mTLS) for bidirectional authentication, and signed payloads for message integrity and non-repudiation.
3. **Digital Signatures** Digital signatures anchor trust across the ecosystem. Critical payloads - schedules, meter readings, settlement data, consent artefacts—can be cryptographically signed by the originating party, enabling downstream actors to verify authenticity without relying on intermediaries.
4. **Verifiable Credentials**: IES will use verifiable credentials (machine readable and cryptographically signed) to encode attestations about actors and assets. These credentials will be cryptographically signed by issuers and can be verified by any relying party without contacting the issuer, enabling privacy-preserving and scalable trust validation.
5. **Zero-Knowledge Proofs and Hashing for Privacy Protection**: Where sensitive data must be proven without disclosure, IES will support zero-knowledge proof (ZKP) techniques - for example, proving subsidy eligibility without revealing full consumption history. Hashing ensures data integrity for audit trails while being privacy preserving.
6. **Tokenization**: Ability to represent various energy assets and contracts ina tokenized fashion allows them to be self-describing, verifiable, and transactable. This dramatically decreases the transaction costs within the ecosystem and creates future innovation possibilities.

Post-Quantum Readiness

IES is designed to be agile and evolvable supporting co-existence of multiple signing, encryption, and other cryptographic techniques. This ensures the IES ecosystem can adopt newer Post Quantum Cryptography (PQC) algorithms and techniques as they evolve in an asynchronous fashion (meaning not everyone having to adopt at the same time).

Observability

IES treats observability as a fundamental capability baked into the protocols, via a set of telemetry specifications and an architecture design. Observability allows the right actors/systems to watch the right amount of data at the right time to take necessary actions in a proactive manner in a federated manner retaining autonomy and control at the right levels.

Observability is neither about a mega centralized platform accumulating all data nor about some fancy dashboard or tool. Every meaningful interaction in IES produces the right evidence, in the right shape, with the right privacy guarantees, so that (a) each participant can operate and debug their own systems (“self”), and (b) the

programme can understand adoption, friction, and impact (“programme”), without centralising sensitive data or creating surveillance. Observability in IES explicitly balances privacy, confidentiality, and autonomy with the legitimate need for the ecosystem (and the programme) to understand what is happening.

Principles of observability in a federated protocol

1) Protocol evidence via standardized telemetry
If the protocol does not define what must be observable, the ecosystem resorts to bespoke logging, unverifiable spreadsheets, and inconsistent reporting. IES therefore specifies observability artefacts at the interaction boundary.

2) Anonymous aggregates by default
Programme-level insights prefer anonymous/aggregated signals over raw records. Default stance: minimum information needed to learn what we need to learn.

3) Layered observability: business, systemic, technology, infrastructure
Observability is articulated across layers because the questions, consumers, and privacy sensitivities differ.

- **Business:** what happened in the domain? (e.g., trades, volume, settlements)
- **Technology:** is implementation behaving correctly? (e.g., API counts, latency, schema/signature/policy failures)
- **Infrastructure:** is compute/network healthy? (e.g., uptime, p95/p99 latency, error spikes, backlogs, dependency reachability)
- **Systemic:** is the system-of-systems working? (e.g., where cross-org transactions fail, brittle integrations, chronic friction pockets)

4) Observe the boundary, protect the inside
IES does not require deep introspection into internal systems. It observes inter-system interactions: integrity, timing, failure reasons, policy application, and verifiable outcomes.

5) Privacy- and purpose-bound visibility
Not everyone sees everything. Observability access is governed by role, purpose, and policy, with strong identity and auditability.

Without observability: the protocol exists, but operations devolve into bespoke logs, manual RCA calls, inconsistent reporting, and unverifiable adoption claims.
With observability: protocol interactions produce evidence by design (trace + reason + receipt), nodes expose standard boundary metrics, policy decisions are observable, programme metrics come from anonymous aggregates and attestations, and conformance enforces it end-to-end.

Self observability

Self observability is the ability for any IES participant (utility, market platform, registry operator, technology provider) to answer:

- Are my IES interfaces up, performant, and compliant?
- Are failures due to me, the counterparty, the network, a registry, or a policy/version mismatch?
- Are we applying the correct policy and schema versions?
- Can I prove what happened, when, and why—without ambiguity?

Protocol-level observability artefacts (part of conformance)

A) Trace context (cross-party correlation)
Every IES interaction carries trace context so multi-hop journeys can be correlated during incidents—without exposing payload content.

B) Outcome + reason taxonomy (failure clarity)
Every API response includes outcome status (accepted/rejected/partial/pending) and a standard reason taxonomy (policy vs auth vs schema vs dependency vs internal error), plus (where safe) a short machine-readable debug hint. This enables faster MTTR across organisations.

C) Verifiable receipts for key interactions
For operationally/financially meaningful interactions, responses produce a receipt object storable by both parties, including: request/response hashes, timestamps, identities/roles, schema + policy pack versions, outcome + reason code, and cryptographic signature. This makes audits/disputes evidence-driven.

D) Minimal metrics endpoint (participant-owned, policy-governed)
Each participant exposes an authenticated metrics interface for boundary metrics: API availability/latency (avg, p95, p99), request counts by operation, failure counts by reason, registry sync health, certificate/key rotation health. These are boundary metrics, not internal business logs.

Programme observability

Programme observability answers: adoption depth, where adoption is easy/hard and why, implementer challenges (schema/policy/registry/ops), and overall impact (efficiency, reliability, market outcomes, innovation). Because IES is federated, programme observability is built from anonymous aggregates and verifiable evidence, not central ingestion of raw operational data.

A) Adoption observability (who adopted what)
A programme-level adoption evidence model participants can publish (or allow querying), including: supported IES specs/profiles + versions;

sandbox/pilot/production; transaction classes enabled; operational maturity indicators (uptime/latency/incident bands); conformance/certification status.

B) Friction observability (where and why it's hard)
Non-sensitive signals such as: top failure reasons in sandbox/certification; version mismatch frequencies; time-to-first-success for new integrators; recurring policy interpretation conflicts; dependency bottlenecks (registry reachability, credential issuance delays). This guides spec simplification, reference implementations, and targeted support.

C) Impact observability (what changed because of IES)
Aggregable metrics: business outcomes (e.g., trade counts/volumes where safe as aggregates), operational outcomes (reconciliation cycles, dispute times), ecosystem outcomes (active implementers, interoperable apps). Each use case defines a small, measurable, comparable, privacy-preserving impact set.

Privacy, confidentiality, and governance

Observability must not undermine trust. Therefore: programme reporting defaults to anonymous aggregates; access is purpose-bound (incident response, compliance audit) with explicit authorisation; payload exposure is minimised (hashes, counters, distributions, reason codes over raw payloads); and auditors are audited (access to sensitive artefacts is logged and reviewable).

Reference Use Cases built on IES Architecture

IES is a general-purpose, federated protocol layer for the power sector—not a collection of one-off integrations or "pilot-by-pilot" solutions. The core remains stable and context-invariant while enabling many extensions to be built by the ecosystem (unified, not uniform; generalised and extensible; interoperable and trusted by design). Architecture becomes real only when exercised against concrete, high-value sector problems.

We therefore started with a small set of anchor use cases that (a) cut across stakeholders, (b) expose recurring friction—fragmented identifiers, inconsistent data formats, manual workflows, weak auditability—and (c) demonstrate reuse of the same IES building blocks: identity/addressability, registries, verifiable credentials, machine readable policies, receipts/provenance, and standard interaction patterns.

Note: These use cases are merely a set of initial ones and not a comprehensive final set. The set of use cases supported by IES will continuously grow and evolve.

Below sections provides a brief view of each current use case: why it matters, how IES applies, and what changes after IES—without going into pilot scope or implementation detail.

1. Inter-DISCOM P2P Energy Trading

Cross-state, cross-DISCOM P2P trading is the hardest “system-of-systems” case: buyers and sellers sit under different utilities, while contracting, delivery proof, wheeling/billing, and dispute handling span multiple organisations with no central controller to fall back on. Today, identities and connection–meter linkages don’t carry cleanly across boundaries, evidence comes in mismatched formats, and verification/settlement becomes slow and brittle.

IES makes P2P repeatable by standardising discovery plus contracting interactions and defining shared evidence and receipt artefacts (e.g., signed meter-derived actuals) that every participant can reference consistently while each platform, exchange, and DISCOM stays autonomous. The result is interoperable platforms, consistent DISCOM visibility into scheduled vs actual trades, and evidence-driven reconciliation and disputes through shared protocol receipts and versioned records.

2. Regulatory Data Exchange (Standardised, Verifiable Filings)

Regulatory submissions are still document-heavy and utility-specific, which creates constant compliance friction and expensive reconciliation. In practice, templates are bespoke, validations are ad-hoc, provenance is weak (“what was submitted when, and what was acknowledged”), and comparisons across utilities are painful.

IES turns filings into API-native exchanges by enabling canonical schemas, standard APIs, validation reports, version chains, and tamper-evident receipts—without centralising data. Utilities and regulators keep their systems of record; IES standardises the handshake and the evidence. Outcomes become machine-readable, traceable, comparable-by-default, and far easier to audit and analyse.

3. Energy Policies

Policies arriving as PDF documents create a predictable complexity: interpretation might not be uniform, updates propagate slowly and inconsistently, and disputes flare because it’s hard to prove a billing outcome actually applied the correct effective policy. Utilities and vendors often re-implement the same logic repeatedly, each time with subtle differences.

IES supports and offers methodologies to move towards machine-readable, digitally signed, versioned policy packs with provenance, and deterministic evaluation. Systems can resolve the applicable policy version and produce explainable outputs

that are bound to the policy source and version, making policy portability and verification practical at scale.

4. DER Visibility

DER scale is throttled by fragmented registries and weak linkage between the consumer, premise, connection, meter, and DER asset. Today that shows up as DISCOM-specific portals, inconsistent identifiers, repeated verification, spreadsheet exchanges, and poor comparability for regulatory or market-adjacent reporting.

IES enables federated registries and canonical DER schemas, with strong linkage via credentials and interoperable exchanges that include validation, acknowledgements, and receipts—without creating a central database. DER data becomes consistently discoverable and auditable, reducing reconciliation friction, improving planning, and enabling aggregators/programmes to scale participation across utilities.

5. Consumer Side Flexibility

Flexibility is technically feasible but operationally hard to scale because enrolment, consent, dispatch, measurement & verification, and settlement are fragmented—and therefore frequently disputed. Discoverability is weak, consent is often one-off and non-portable, dispatch/telemetry semantics vary, baselines get contested, and settlements drag with limited auditability.

IES provides building blocks to construct the end-to-end journey: catalogues for programmes, identity-to-connection linkage, reusable consent, capability descriptors, event objects, policy objects, and receipts that tie together enrol → dispatch → verify → settle. That makes flexibility repeatable across DISCOMs and providers with faster onboarding, verifiable outcomes, and itemised settlement—while staying fully federated.

6. Digital Consumer Lifecycle Management

Many downstream services need trusted consumer information plus consent, but relying on centralised storage creates fragility and privacy risk. IES instead enables energy credentials issued by authoritative parties (e.g., DISCOMs) and held in consumer-controlled wallets for secure, selective sharing—so trust travels with the consumer, not with duplicated databases.

Before this, consumer data is fragmented and repeatedly re-verified; consent is app-specific with weak portability and revocation; and audit trails for “who shared what, when, and under what authority” are thin. IES standardises credential schemas and flows for issuance, presentation, verification, renewal/revocation, and policy-governed consent artefacts with verifiable receipts—while CIS/GIS/MDMS

remain systems of record—leading to faster onboarding, fewer repeated checks, and stronger privacy and auditability.

7. EV Charging

EV charging is inherently multi-operator, so interoperability is the difference between a seamless experience and a patchwork of proprietary apps and roaming integrations. Today, discovery and access are fragmented, reservation and session semantics differ by network, and integration costs stay high.

IES provides common semantics for discovery (geo/route, connector type, availability), reservation (time-bound orders), and charging sessions (tracked fulfilments), so multiple providers can participate using a shared interaction grammar. Drivers can discover, reserve, charge, and pay across providers via interoperable apps, while charge point operators retain operational autonomy and the ecosystem scales without bespoke one-off integrations.

Across all the above use cases, the intent is to reuse the same core primitives (identity, registries/credentials, machine readable policies, standard APIs and interaction patterns) rather than building bespoke protocol or patterns per use case. These anchors were selected to prove that IES functions as a common handshake specification across federated systems, even as operational and commercial contexts vary.

Technology governance & compliance

- <TBD>

Adoption Strategy

<TBD>

Developer Experience, Sandbox & Certification

For IES to scale nationally across utilities, aggregators, OEMs, startups, algorithmic market agents, and DER manufacturers, the ecosystem must provide a frictionless, transparent, and predictable developer experience. Early choices here will determine adoption speed and long-term innovation.

Principles guiding developer experience

- Open and well-documented APIs that follow consistent standards and versioning patterns.
- Clear onboarding flows for organizations, devices, and software agents.
- Accessible test environments that mimic real utility/market constraints.
- Rapid iteration cycles, supported through synthetic datasets, mock services, and reference implementations.
- Reference code and tools, for accelerating the developer efforts.

Certification as Trust Infrastructure

IES certification is not merely a compliance mechanism—it is a **foundational architectural component** that operationalizes the "trusted and secure" principle across the federated ecosystem.

In traditional centralized systems, trust is enforced through physical control and bilateral agreements. In a federated DPI where millions of distributed entities interact without central oversight, **certification becomes the architectural mechanism** that enables:

> **Verifiable interoperability** – Validating that implementations of IES specifications (registries, credentials, APIs, consent mechanisms) function correctly across heterogeneous systems before production deployment.
>
> **Graduated assurance** – Proportional validation matching risk profiles: automated conformance for low-risk applications, scenario testing for operational platforms, and compliance audits for safety-critical infrastructure.
>
> **Trust at scale** – Replacing manual verification with programmatic validation, enabling rapid ecosystem growth without compromising reliability or security.

Certification validates implementations across IES's four architectural layers (Data, Identity, Exchange, Consent), ensuring the ecosystem functions as a coherent whole despite its federated nature. This positions certification as **trust infrastructure** rather than regulatory overhead.

For detailed certification profiles, tiered models, governance, and phased rollout, refer to Section 4.x in the IES Strategy Document.

References

1. Consent artefact published by Meity - https://dla.gov.in/sites/default/files/pdf/MeitY-Consent-Tech-Framework%20v1.1.pdf
2. IndEA framework - https://egovstandards.gov.in/sites/default/files/2021-10/IndEA%20Framework%20V%201.0.pdf
3. User-centric credentialing and Personal data sharing - https://drive.google.com/file/d/1kOEvGqjsKxoYvK3KglB77-DmmiZIykhA/view
4. Decentralized Data Marketplace (DDM) - https://beckn.io/wp-content/uploads/2026/02/Decentralised-Data-Marketplace-DDM-1.pdf
5.

Annexures

Annexure 1: IES Specifications Links

1. India Energy Stack (IES) github links
    a. Specifications - https://github.com/India-Energy-Stack/ies-specs
    b. Documentation - https://github.com/India-Energy-Stack/ies-docs
    c. Reference implementation for EV charging - https://github.com/India-Energy-Stack/ies-docs/tree/main/implementation-guides/EV-Charging
    d. Draft architecture for P2P trading (inter-discom) - https://github.com/India-Energy-Stack/ies-docs/tree/main/implementation-guides/P2P%20Trading
    e. Energy data exchange - https://github.com/India-Energy-Stack/ies-docs/tree/main/implementation-guides/data_exchange
    f.

Annexure 2: Draft Verb-Noun Mapping

Links to visualizations:

- Draft ecosystem map
- Discom centric map

Generic Verbs

discover (catalog/service)

publish (capability/offer)

search (capability/slot)

quote (tariff/offer)

select (offer/slot)

confirm (order/enrollment)

cancel / reschedule (order/slot)

notify (status/update)

pay / settle (invoice/fee)

apply (open-access/net-metering)

enroll (program/market)

Grid-Specific Verbs

schedule (dispatch/energy block)

revise (schedule update)

dispatch (setpoint/ramp)

curtail (generation/load)

shed / restore (load/feeder)

reconfigure (topology/switch)

override (controller/AGC)

acknowledge (interlock/command)

trip / block (device/breaker)

validate (forecast/network constraint)

reconcile (imbalance/deviation)

stream (telemetry/PMU/SCADA)

aggregate (feeder/area/control zone)

alarm / clear (limit breach)

clear (market/auction)

allocate (capacity/transmission right)

settle (energy/imbalance)

Assurance / Compliance Verbs

enroll (party/role)

authorize (role/permission)

revoke / rotate (credential/key)

attest (device/firmware/meter seal)

log / notarize (action/receipt)

conform (profile/standard)

audit (operational logs)

certify (party/license/device)

dispute / resolve (settlement/penalty)

Annexure 3: Bibliography of Standards and Specifications

This appendix lists major technical standards, specifications, and regulatory instruments relevant to the India Energy Stack (IES) architecture across the power sector value chain—generation, transmission, distribution, markets, behind-the-meter resources, EVs, and consumer data sharing.

The focus is engineering and implementation oriented: for each standard family we capture what it does, where it is used in real systems, and which stakeholder groups care about it. This list is not legally exhaustive, but is intended to be broad enough that most IES design decisions can be grounded in one or more of these references.

Reading guide and categorisation

Entries are grouped into nine thematic categories:

1. Grid information models & utility automation
2. Operational communications, SCADA/EMS and synchrophasors
3. DER, flexibility and market integration
4. Metering, settlement and consumer energy data
5. Electric mobility and vehicle-to-grid (V2G)
6. Cybersecurity, privacy and operational security
7. Identity, consent and DPI building blocks
8. Core IT / web interoperability standards
9. Indian grid codes, technical and metering regulations

Cross-cutting note (Beckn/DEG/UEI): In addition to sector protocols (IEC/IEEE/ISO), IES can adopt open transaction-network patterns (discovery → ordering → fulfilment → status) from Beckn-based energy specifications such as DEG/UEI, especially for multi-platform interoperability (markets, DER flexibility, EV charging, consumer services).

Common pattern used per entry

- Standard / spec (organisation, code, title)
- Scope — what it standardises
- Typical usage in power systems — where it is deployed today
- Relevant stakeholders — who primarily cares about it
- IES relevance — why it matters for IES design

This bibliography is intended to be used as a design checklist when defining IES APIs, data models, registries, and implementation profiles.

---

1) Grid information models & utility automation

1.1 IEC Common Information Model (CIM) — IEC 61970 / 61968 / 62325

- Standards
    - IEC 61970 — Energy management system application program interface (EMS-API) — CIM for transmission
    - IEC 61968 — Application integration at electric utilities – System interfaces for distribution management — CIM for distribution
    - IEC 62325 — Framework for energy market communications — CIM extensions for market transactions

- Scope
    - Object-oriented information model of power system assets, topology, measurements and market constructs (UML; serialised via XML/RDF/JSON, etc.)

- Typical usage
    - Transmission EMS/SCADA/planning/asset integration (61970)
    - DISCOM integration across DMS/OMS/GIS/MDM/customer systems (61968)
    - Wholesale market communications: clearing/scheduling/settlement (62325)

- Relevant stakeholders
    - CTU/STUs, RLDCs/SLDCs/NLDC, transmission licensees, DISCOMs, system operators, market operators, OEMs, planning/analytics vendors

- IES relevance
    - Canonical reference for network/asset semantics; strong candidate for IES canonical models for topology, connectivity, assets, and state

1.2 IEC 61850 — Power utility automation

- Standards
    - IEC 61850 series — Communication networks and systems for power utility automation (substations, DER, hydro, wind, etc.)

- Scope
    - Standardised data models (Logical Nodes), services, and communication profiles (MMS over TCP/IP, GOOSE, Sampled Values) for protection/control/monitoring IEDs

- Typical usage
    - Substation automation systems (SAS), protection relays, bay controllers, station/process bus; plant control; increasingly DER/microgrid controllers

- Relevant stakeholders
    - Transmission/distribution utilities, GENCOs, IED OEMs, EPCs, SCADA vendors, testing/certification bodies

- IES relevance
    - Primary reference for field-level semantics and event/messaging patterns when mapping substation/plant data into IES telemetry/event APIs

1.3 Related modelling / semantics references (commonly encountered)

- IEC 62056 / DLMS-COSEM (also central to Category 4)
    - Organisation / code — IEC 62056 series; BIS companions (IS 15959 parts; IS 16444 parts)
    - Scope — COSEM object models, OBIS codes, communication profiles for metering data exchange
    - Usage — Smart meters, AMI head-end, MDM/MDMS, handheld field units
    - Stakeholders — DISCOMs, AMI vendors, meter OEMs, billing/MDM providers, regulators

- OBIS (Object Identification System)
    - Defined within IEC 62056-61 and national companions (e.g., IS 15959) as coding scheme for metering quantities; critical to meter-centric APIs

- Green Button / NAESB ESPI (customer data sharing reference)
    - XML/JSON schemas and web service profiles for publishing/sharing interval usage data with customers/third parties (primarily North America)
    - Useful reference pattern for customer-centric data access APIs under consent

2) Operational communications, SCADA/EMS and synchrophasors

2.1 Telecontrol & SCADA protocols

- IEC 60870-5-101 / 104 — Telecontrol protocols
    - Scope — Serial (-101) and TCP/IP (-104) telecontrol between control centres and RTUs/IEDs

    - Usage — SCADA links for transmission/distribution substations (common in Europe/Asia)
    - Stakeholders — Transmission/distribution utilities, SLDC/RLDC, RTU/IED OEMs, SCADA vendors

- IEC 60870-6 / TASE.2 / ICCP — Inter-Control Centre Communications Protocol
    - Scope — Real-time data exchange between control centres (SCADA/EMS ↔ SCADA/EMS)
    - Usage — RLDC↔SLDC, RLDC↔NLDC, utility↔regional coordination exchanges
    - Stakeholders — System operators, CTU/STUs, RLDCs, NLDC, large generators tied into EMS networks

- DNP3 / IEEE 1815 — Distributed Network Protocol
    - Scope — Object-oriented telecontrol/automation with robust time-tagging and event reporting
    - Usage — Distribution automation (reclosers/capacitor banks), some transmission RTUs (notably North America)
    - Stakeholders — DISCOMs, DA/AMI vendors, OEMs

2.2 Synchrophasors and wide-area monitoring

- IEEE C37.118.1 / C37.118.2
    - Scope — PMU measurement requirements (-.1) and synchrophasor data transfer (-.2)
    - Usage — WAMS, oscillation analysis, dynamic security assessment, PMU-based protection
    - Stakeholders — RLDCs/NLDC, transmission licensees, system operators, PMU/PDC vendors

- IEC/IEEE 60255-118-1
    - Joint IEC/IEEE alignment for synchrophasor measurement requirements; used in modern relay/PMU specifications

3) DER, flexibility and market integration

3.1 IEEE 1547 series — DER interconnection

- IEEE 1547-2018 — Interconnection and interoperability requirements for DER with EPS interfaces
- Scope
    - Ride-through, reactive power, voltage regulation, interoperability/information exchange; testing in IEEE 1547.1

- Typical usage
    - Utility interconnection rules for rooftop solar, storage, small wind, etc.
- Relevant stakeholders
    - DISCOMs, DER aggregators, inverter/storage OEMs, installers, regulators
- IES relevance
    - Device-level profiles for DER registration, capabilities, telemetry via IES to aggregators/markets

3.2 IEEE 2030.x and Smart Energy Profile (SEP 2.0)

- IEEE 2030.5 — Smart Energy Profile application protocol (SEP 2.0)
    - Scope — IP-based application layer for utility↔DER/meter/EV/customer energy systems
    - Usage — Utility-to-DER comms for PV/storage/EV charging/demand response
    - Stakeholders — DISCOMs, aggregators, DER gateway vendors, inverter/EVSE OEMs

- Related IEEE 2030 family
    - IEEE 2030 (interoperability reference)
    - IEEE 2030.7 / 2030.8 (microgrid controller guidance/testing)

3.3 OpenADR — automated demand response

- OpenADR 2.0 / 2.0b / 3.0 (OpenADR Alliance)
    - Scope — machine-readable price/event DR messages (VTN↔VEN)
    - Usage — DR, peak shaving, flexibility services for C&I and aggregators; some residential programmes
    - Stakeholders — System operators, DISCOMs, aggregators, C&I consumers, BMS vendors
    - IES relevance — reference for event-based flexibility APIs/payload semantics

3.4 Market messaging & scheduling

- IEC 62325 (CIM for markets)
    - Supports scheduling, nominations, capacity allocation, settlement, market process messaging

- Regional / national market patterns
    - ENTSO-E (Europe), NAESB (North America) and similar frameworks as design references (even if not adopted verbatim)

4) Metering, settlement and consumer energy data

4.1 IEC 62056 DLMS-COSEM family

- Scope
    - COSEM data model, OBIS codes, and transport profiles (HDLC/TCP-IP/PLC/RF/cellular, etc.)
- Typical usage
    - AMI smart meters, head-end, MDMS, on-prem displays/gateways; basis of many national meter standards
- Stakeholders
    - DISCOMs, AMI vendors, meter OEMs, billing vendors, regulators, conformance labs
- IES relevance
    - Canonical source for metering quantities/code lists; maps into IES data models and customer data access APIs

4.2 BIS IS 15959 series — Indian companion specs to IEC 62056

- IS 15959:2011 and Parts 1/2/3 — Data Exchange for Electricity Meter Reading, Tariff and Load Control – Companion Specification
- Scope
    - India-specific companion specs: DLMS/COSEM objects, OBIS usage, comms profiles for Indian deployments
- Typical usage
    - De-facto / mandatory basis for meter data exchange in Indian DISCOM AMI deployments
- Stakeholders
    - DISCOMs, AMI vendors, meter OEMs, BIS, CEA, testing labs
- IES relevance
    - IES meter-centric APIs and consented data sharing should align to IS 15959/IEC 62056 to minimise duplication

4.3 BIS IS 16444 series — AC static smart meters (India)

- IS 16444 (Part 1):2015 — Direct connected watthour smart meter (Class 1/2)
- IS 16444 (Part 2):2017 — Transformer-operated smart meters (Class 0.2S/0.5S/1.0S)
- Scope
    - Technical/performance requirements: accuracy, load profiles, comms modules, tamper detection
- Typical usage
    - Procurement baseline under national programmes (e.g., RDSS) and BIS/QCO compliance
- Stakeholders

    - DISCOMs, meter OEMs, test/cert labs, regulators, programme agencies
- IES relevance
    - Sets baseline capabilities/data items that IES-based AMI and consumer data APIs should assume

4.4 CEA metering regulations (India)

- CEA (Installation and Operation of Meters) Regulations, 2006 (+ amendments)
    - Metering requirements: installation, accuracy, location, operation; responsibilities for reading/testing/replacement
- Typical usage
    - Basis for metering practice across T&D and retail; defines metering hierarchies/responsibilities
- Stakeholders
    - CEA, CERC/SERCs, transmission licensees, DISCOMs, large consumers, meter OEMs
- IES relevance
    - IES energy accounting, loss computation, and reconciliation must respect metering points/responsibilities/accuracy classes

4.5 Customer-centric energy data sharing (reference patterns)

- Green Button / NAESB ESPI
    - Retail customer data download/share patterns (interval data); useful reference though IES will likely rely on DEPA-style consent roles
- Utility / jurisdiction APIs
    - UK/EU smart meter hubs and other bespoke interfaces (design references for usability/security expectations)

5) Electric mobility and vehicle-to-grid (V2G)

5.1 EV conductive charging systems — IEC 61851 + connectors

- IEC 61851 series — Electric vehicle conductive charging system
    - Scope — charging modes, safety/control functions, EV↔EVSE communication requirements
- IEC 62196 series — plugs, socket-outlets, connectors and inlets (Type 1/Type 2, etc.)
- Typical usage
    - Core physical and safety interoperability references for EVSE/EV OEMs/CPOs/utilities
- Stakeholders

    - EVSE OEMs, CPOs, DISCOMs, OMCs, city authorities, regulators, auto OEMs

5.2 Vehicle-to-grid communication — ISO 15118

- ISO 15118 series — Road vehicles – Vehicle to grid communication interface
    - Scope — EV↔EVSE charging control, tariffing, plug-and-charge; later parts for V2G
    - Usage — intelligent AC/DC charging; V2G pilots where EV acts as controllable load/storage
    - Stakeholders — EV OEMs, EVSE/CPOs, aggregators, utilities, roaming platforms, payment providers
    - IES relevance — informs EV resource registration/control semantics inside IES flexibility models

5.3 EV charging network protocols — OCPP / OCPI / related

- OCPP (Open Charge Alliance)
    - Scope — charge point ↔ CSMS: config, transactions, diagnostics, smart charging
    - Usage — de-facto global standard (1.6, 2.0.1 widely deployed)
    - Stakeholders — CPOs, EVSE OEMs, platform vendors

- OCPI
    - Scope — roaming/inter-operator protocol: locations/tariffs, authorisation, sessions, settlement
    - Usage — roaming networks and hubs; national roaming platforms
    - Stakeholders — EMSPs, CPOs, roaming hubs, regulators
- Other roaming protocols
    - OCHP, eMIP, proprietary (reference patterns for “network of networks” interoperability)
- IES relevance
    - Mature message flows/resource models that can be mapped/bridged into IES EV APIs

6) Cybersecurity, privacy and operational security

6.1 IEC 62351 — security for power system communications

- IEC 62351 series — Data and communications security for IEC TC57 protocols (60870-5, 60870-6/ICCP, 61850, 61970, 61968)
- Scope
    - Authentication, encryption, key management, RBAC, security event logging for power protocols
- Typical usage

    - Securing SCADA/EMS/substation automation and CIM integrations
- Stakeholders
    - Utilities, system operators, OEMs, cybersecurity teams, regulators
- IES relevance
    - IES security guidelines should be compatible with 62351 protocol-specific profiles where TC57 protocols are used

6.2 IEC 62443 — industrial automation & control systems security

- IEC 62443 series
    - Scope — defence-in-depth framework; security levels; secure product development lifecycle
    - Usage — OT cybersecurity programmes; often used by regulators/insurers as baseline
    - Stakeholders — GENCOs, DISCOMs, TRANSCOs, OEMs, integrators, cyber service providers
    - IES relevance — guides security architecture and lifecycle for IES-connected OT systems

6.3 NISTIR 7628 — smart grid cybersecurity guidelines

- NISTIR 7628 Rev.1
    - Scope — control catalogue + privacy considerations tailored to smart grid (AMI/DR/DER)
    - Usage — reference for national frameworks and utility risk assessments
    - IES relevance — helpful baseline for IES security controls and assurance levels

6.4 NERC CIP (North America) — compliance pattern reference

- NERC CIP standards (CIP-002…CIP-014, etc.)
    - Scope — mandatory reliability standards for cyber/physical security of bulk electric system cyber systems
    - Usage — North American BES entities; used elsewhere as benchmark
    - IES relevance — operational compliance patterns relevant as reference for high-assurance operations

6.5 ISO/IEC 27000 family — information security management

- ISO/IEC 27001:2022
    - Scope — ISMS requirements: policies, processes, continual improvement
    - Usage — common baseline across utilities/operators/service providers

    - Stakeholders — all IES participants processing customer or operational data
    - IES relevance — recommended ISMS baseline for shared IES infrastructure (registries, hubs, platforms)

6.6 Indian cyber and grid security guidance

- CEA Guidelines for Cyber Security in Power Sector (2021 + updates)
    - Governance, network security, incident response, monitoring requirements
- CERC Indian Electricity Grid Code 2023 — cybersecurity provisions
    - Cyber-secure operation responsibilities
- IES relevance
    - IES security/operational processes must align with CEA/CERC guidance, especially for “critical” entities

7) Identity, consent and DPI building blocks

7.1 OpenID Connect / OAuth 2.0 (and related)

- OpenID Connect Core 1.0
    - Identity layer on OAuth 2.0; REST/JSON profile claims
- OAuth 2.0
    - IETF standard for delegated authorisation (token-based)
- Usage in energy
    - Utility portals, EV platforms, data-sharing services
- IES relevance
    - Natural choice for IES identity/authorisation/token models for human and machine actors; supports federation and multi-tenant deployments

7.2 W3C Verifiable Credentials (VC) and DIDs

- W3C Verifiable Credentials Data Model 2.0
    - Cryptographically verifiable credentials; issuer–holder–verifier model; privacy-respecting disclosure
- Usage
    - Digital wallets, permits, certifications, compliance attestations
- IES relevance
    - Candidate for portable, verifiable licences/registrations (e.g., aggregator accreditation, installer certificates, equipment compliance)

7.3 India Stack data layer — DEPA and Account Aggregators

- DEPA

    - Consent artefacts + consent managers enabling user-controlled data sharing
- Account Aggregator (AA)
    - RBI-regulated NBFC-AA ecosystem for financial data using standard APIs with revocable consent
- IES relevance
    - Strong template for energy data consent and standardised roles/APIs (adapted for energy telemetry and consumer data)

7.4 Beckn protocol and open transaction networks

- Beckn Protocol
    - Standard APIs/message formats for decentralised discovery/ordering/fulfilment/status
- Usage
    - ONDC and other open networks
- IES relevance
    - Pattern for open, federated service marketplaces (energy retail, flexibility, EV charging); potential interoperability with Beckn

8) Core IT / web interoperability standards

These are generic but foundational for IES-compliant implementations:

- IETF / W3C protocols
    - HTTP/HTTPS, TLS, TCP/IP, DNS, URI/URL, REST, WebSockets
- Data formats & serialisation
    - JSON (ECMA-404 / RFC 8259), XML (W3C), JSON Schema, YAML, Protocol Buffers, Avro
- Messaging / IoT protocols
    - MQTT, AMQP, CoAP (often used for device-to-cloud/event streaming; may carry mapped IEC/DLMS semantics)
- Security tokens / signatures
    - JWT, JWS/JWE, X.509, PKI profiles, TLS certificates

In IES these manifest as API design guidelines, event streaming patterns, and security best practices.

9) Indian grid codes, technical and connectivity regulations

9.1 CEA technical standards for connectivity

- CEA (Technical Standards for Connectivity to the Grid) Regulations, 2007 (+ amendments)
    - Connectivity conditions for generating stations/transmission lines; voltage/frequency/protection/FRT/comms facilities

- CEA connectivity below 33 kV Regulations, 2013
    - Distribution-level connectivity requirements
- IES relevance
    - IES registries/models for generators/DER/grid compliance should capture parameters referenced (short-circuit levels, protection schemes, comms requirements)

9.2 Indian Electricity Grid Code (IEGC) — CERC

- IEGC Regulations, 2023 (CERC)
    - Reliability, operations, scheduling/despatch, renewables integration, ancillary services, cybersecurity duties
- Typical usage
    - Binding on inter-state connected entities; backbone of operational discipline and data exchange expectations
- IES relevance
    - Many obligation flows (telemetry, scheduling, outage coordination, performance reporting) become digital flows supported by IES APIs

9.3 Other CEA regulations relevant to IES

- CEA Safety Requirements Regulations, 2011
    - Safety/operation requirements for plants/lines; informs safety attributes in asset registries
- Draft Communication Standards for Power System Operations (CEA)
    - Direction on comms standards for system operators (likely IEC TC57); important to track for future IES alignment

Summary

- Canonical data models
    - Prefer CIM (IEC 61970/61968/62325) and DLMS/COSEM + IS 15959/IS 16444 as primary semantic sources; use IEC 61850 for field/substation semantics
- Interoperability with existing utility systems
    - Explicitly map IES resource models to operational protocols: IEC 60870-5-101/104, IEC 60870-6/TASE.2, IEEE 1815/DNP3, IEC 61850, IEEE C37.118
- Flexibility, DER and EV services
    - Reference IEEE 1547, IEEE 2030.5, OpenADR, ISO 15118, IEC 61851, OCPP, OCPI for message flows and controllable attributes
- Security, identity and consent
    - Align with IEC 62351, IEC 62443, ISO/IEC 27001, NISTIR 7628, Indian CEA cyber guidelines, OpenID Connect/OAuth 2.0, W3C VC, and DEPA/AA

- Consented consumer data sharing
    - Use Green Button/ESPI as a usability reference, but base the actual approach on DEPA-style consent artefacts and AA-like roles adapted to energy

Annexure 4: Examples (Indicative)

SLDC ↔ GENCO / IPP / Aggregator

- schedule (energy block [15-min / 96 blocks]) →
- revise (schedule revision [R0…Rn, gate closure]) →
- dispatch (active power setpoint [MW], ramp [MW/min]) →
- dispatch (reactive power setpoint [MVAr], voltage target [kV]) →
- curtail (generation output [%/MW], reason code) →
- start / stop (unit command [on/off], min-up/min-down) →
- acknowledge (command receipt [id, ts, signature]) ←
- report (actual generation [MW/MWh], deviation [Δ]) ←
- reconcile (actual vs schedule [interval list]) ↔
- notify (contingency / outage / derate [MW, duration]) ↔
- certify (unit accreditation [AGC capable, primary reserve]) →
- attest (telemetry quality [availability %, latency]) ←

SLDC ↔ TRANSCO (STU/CTU)

- reconfigure (switching plan [element ids, steps, interlocks]) →
- approve / reject (clearance for work [permit id]) ←/→
- allocate (transmission capacity [ATC/NTC]) →
- compute (loss factors, TTC/ATC updates) →
- stream (line flows, bus voltages, frequency) ↔
- alarm / clear (limit breach [thermal/voltage/stability]) ↔
- authorize (role-bound access to SCADA/EMS views) →
- log (SOE / switching logs, time-synced) ↔

SLDC / REMC ↔ RE Generator / OEM

- forecast (day-ahead / intraday generation profile) ←
- validate (forecast vs minimum technical constraints) →
- commit (availability declaration [Pmax, Pmin, ramp]) ←
- curtail (RE spill instruction [MW/%, cause]) →
- override (inverter settings: Volt/VAR, Freq-Watt) →
- report (plant availability, outage ticket) ←
- attest (firmware hash, certified settings, date) ←
- audit (change log for setpoints/config) ↔

Market Operator (PX/DSM) ↔ Participants (GENCO/Trader/DISCOM)

- publish (auction notice, product spec) →
- quote (bid/offer [price, quantity, curve]) ←
- clear (market result [MCP/MQV, awards]) →
- allocate (congestion revenue rights / corridor rights) →
- settle (energy/ancillary invoices; DSM deviations) →

- dispute / resolve (settlement item, evidence bundle) ↔
- certify (participant accreditation, bank guarantees) →
- suspend (participant for non-compliance) →

DISCOM ↔ Consumer / Prosumer / Open-Access

- discover (tariff / program catalog, DR offers) →
- enroll (DR program / ToD plan / net-metering) ←
- shed / restore (contracted load [kW/feeder/group]) →
- request (DER dispatch profile for prosumer) →
- stream (AMI meter interval data, outage events) ←
- reconcile (billing determinants: kWh, kVAh, ToD) ↔
- certify (KYC, consumer category, sanctioned load) →
- authorize (data-sharing consent scopes & duration) ↔

DISCOM ↔ Aggregator / ESCO

- publish (portfolio on-boarding requirements) →
- select / confirm (aggregator nomination) ↔
- dispatch (DR event signal: level, window, baseline method) →
- verify (delivered reduction vs baseline) ↔
- settle (performance payment / penalty) ↔
- audit (measurement & verification trail) ↔

DISCOM ↔ OEM / DER Operator

- dispatch (active/reactive setpoints to fleet) →
- override (safety mode / emergency trip) →
- reconfigure (protection/grouping, feeder mapping) →
- attest (device identity, cert chain, firmware) ←
- rotate / revoke (device keys, access tokens) →
- log (config changes, tamper flags, evidence) ↔

AMI Headend ↔ MDMS ↔ Billing System

- stream (meter reads [15/30-min], event codes) →
- aggregate (transformer/feeder boundary energy) →
- estimate (VEE rules: validation-edit-estimation) ↔
- compute (losses, technical/non-technical KPIs) →
- generate (billing determinants, ToD splits) →
- notarize (interval datasets, hash receipts) →
- audit (VEE rule application logs) ↔

Metering Entity ↔ Market / Settlement

- attest (signed interval data, meter id, CT/PT ratio) →
- timestamp (UTC/PTP sync status, offset) →

- aggregate (boundary / state / control-area totals) →
- notarize (immutable receipt, Merkle proof) →

TSO (POSOCO) ↔ SLDCs

- coordinate (inter-state schedules, HVDC setpoints) ↔
- issue (system advisories, reserve requirements) →
- activate (ancillary services: RR/FRR/MFR) →
- monitor (frequency, ACE, tie-line flows) ←
- certify (SLDC conformance, reporting) →

STU ↔ CTU

- exchange (planned outages, corridor derates) ↔
- compute (state ↔ national TTC/ATC) ↔
- allocate (LTA/MTOA/STOA rights) →
- log (coordination minutes, approvals) ↔

Generator ↔ OEM / Plant Controller

- tune (governor/AVR parameters within bands) →
- enable / disable (AGC participation) →
- report (AGC response, droop compliance) ←
- attest (protection settings, firmware lock) ←

Renewable Prosumer / Rooftop ↔ DISCOM

- apply (net-metering / gross-metering connection) ←
- certify (installer, equipment, commissioning report) →
- limit (export cap %, anti-islanding status) →
- notify (reverse power flow event) ←

EV Charging Operator ↔ DISCOM / Market

- discover (time-of-use prices, incentives) ←
- request (capacity increase / temporary demand) →
- throttle (charging profile, demand cap) →
- report (site load, session telemetry) ←
- attest (metered kWh per session, signed) →

Open-Access Consumer ↔ SLDC / STU / PX

- apply (open-access, banking, wheeling) ←
- revise (schedule blocks before gate close) →
- settle (implied losses, deviation charges) ↔
- furnish (bank guarantee, contract copies) →

Cyber / Trust Fabric (cross-cutting)

- enroll (party/device/role with regulator CA) →
- authorize (scope: read/stream/command; resource ids) →
- rotate (keys/certs; validity windows) →
- revoke / suspend (credentials, reason code) →
- log / notarize (action receipts, WORM store) ↔
- conform (profile versions; test reports; badges) →

Annexure 5: Risks & Mitigation

(This may later move into the Adoption Strategy section. Several mitigations are built into the IES Architecture; others will require policy, governance, or implementation measures.)

Building national digital infrastructure involves a mix of strategic, technical, and governance risks. These do not undermine the value of IES but highlight the conditions needed for successful adoption and sustained interoperability.

1. **Strategic Risks**
    a. Adoption: If DISCOMs, system operators, OEMs, and market participants do not adopt IES standards, the ecosystem may remain fragmented. Adoption depends on incentives, tooling, and regulatory clarity.
    b. Incentives: Different actors have different priorities. Without aligned incentives, participants may prefer bilateral or proprietary integrations instead of shared infrastructure.
    c. Alignment with regulation: Lack of explicit regulatory backing for registries, credentials, or digital contracts can create uncertainty and slow implementation.
    d. Capacity: Lack of capacity and capability within a large section of key ecosystem players may result in improper or slow adoption.
2. **Technical Risks**
    a. Legacy Readiness: Many utilities and grid operators run legacy systems that may struggle to implement modern APIs, schemas, or profiles.
    b. Inconsistent Data Models: If schemas are interpreted differently by actors, interoperability breaks despite “compliance.”
    c. Weak Trust Mechanisms: If identity, credentials, or registries are not robust or well-governed, the ecosystem cannot rely on the authenticity of assets, actors, or contracts.
    d. Performance Constraints: High-frequency telemetry, events, and scheduling flows must work at scale and low latency; otherwise operational systems may bypass IES.
    e. Technology lock-in: IT vendors or OEMs may push proprietary alternatives without adequate interoperability and portability options thus locking the ecosystem into closed solutions.
3. **Governance Risks**
    a. Unclear Roles: Ambiguity around who operates registries, who issues credentials, or who maintains schemas can stall adoption and fragment implementations.

b. Lack of Certification: Without a strong certification framework, “IES-compliant” systems may behave differently in practice, undermining interoperability.
c. Inadequate Tooling: If sandboxes, SDKs, examples, and reference implementations are weak, adoption slows and vendors build point-to-point integrations.
d. Data Privacy & Consent: Lack of clarity about how data is used, shared, or governed can undermine trust and slow adoption, even when the technical system works correctly.
e. Version Fragmentation: Without strict versioning and change-control processes, different states or vendors may drift to incompatible profiles.

IES’s emphasis on open architecture and standards, accelerator strategy, and multi-stakeholder governance is intended to mitigate many of these risks while leaving room for policy and program design to address the rest.

Annexure 6: More about machine readable policies

Machine readable policies give the mechanism through which IES transforms regulatory, operational, and program rules into machine-readable logic that can be evaluated automatically across the energy ecosystem. Instead of relying on traditional policy documents written in natural language—which require human interpretation, implementation in multiple IT systems, and often lead to inconsistent enforcement—machine represented policies express rules in structured formats such as JSON, YAML, or domain-specific rule languages. These representations can be parsed, verified, and executed by digital systems, enabling uniform interpretation of policies across states, DISCOMs, market platforms, aggregators, OEM systems, and consumer applications.

In a highly decentralised energy environment with millions of participants, DERs, and digital services, manual or discretionary enforcement is neither scalable nor consistent. Representing policies in machine readable formats in public directories provides a shared, authoritative source of truth for rules governing eligibility, tariffs, constraints, limits, and program conditions. Whether an interaction involves onboarding a rooftop solar system, validating Open Access eligibility, determining if a P2P transaction complies with feeder constraints, checking dynamic pricing conditions, or assessing compensation under a Demand Response event, the same set of rules can be evaluated automatically by any compliant system. This reduces ambiguity, eliminates inconsistent interpretations, and ensures fairness and transparency for all ecosystem participants.

Policies in IES may cover a wide range of use cases:

- **Demand Response and flexibility programs:** machine-readable definitions of event triggers, baselines, compensation formulas, and participant eligibility that any aggregator or DISCOM system can evaluate uniformly.
- **Tariffs and dynamic pricing:** structured representations of time-of-use tariffs, congestion charges, seasonal adjustments, or surge pricing triggers that marketplaces and billing platforms can apply automatically.
- **Open Access and P2P trading constraints:** rules such as "*allow only within 5 km for residential prosumers*" "*respect feeder capacity limits*" or "*disallow export above sanctioned load*" etc. enforced consistently at transaction time.
- **EV charging and mobility:** constraints like "*reservations must be between 1–100 kWh*" or "*fast-charging pricing applies only during peak periods*" which discovery and selection systems validate before offering services to users.

By expressing policies programmatically, IES enables:

- **Consistent interpretation:** Everyone—DISCOMs, aggregators, market operators, OEMs, apps—evaluates policies in exactly the same way.

- **Transparent verification:** Any actor can check which rules applied to which transaction and why a decision was allowed or denied.
- **Low-discretion governance:** Reduces reliance on case-by-case approvals and manual interpretation, lowering friction, delays, and the potential for bias.
- **Dynamic updates without system rewrites:** Regulators can update rules once, and compliant systems automatically apply them.
- **Composable and hierarchical rules:** National regulations, state variations, DISCOM-level constraints, and program-specific rules can all be layered and combined safely.

IES allows policies to be encoded and enforced reliably across a distributed network of participants without centralising execution or decision-making.

Illustrative Example for Rooftop Solar

1. Example Policy - "*Residential consumers may install rooftop solar up to their sanctioned load, subject to feeder capacity limits and installer accreditation.*"
2. Policy
    a. Check prosumer's sanctioned load from the Participant Registry
    b. Check DER capacity does not exceed sanctioned load
    c. Validate feeder load from the ERA Registry
    d. Confirm installer credential from the Credential Registry
3. Rule Pseudocode (Illustrative)

```
rule RooftopSolarEligibility:
      require der.capacity_kw <= prosumer.sanctioned_load_kw
      require era.feeder.current_loading_percent
            < 0.8 * era.feeder.max_loading_percent
      require installer.credentials.exists(
            type   = "ACCREDITED_INSTALLER",
            status = "VALID")
```

4. Outcome - Any DISCOM, aggregator, or rooftop solar portal can instantly and uniformly determine whether the proposed installation complies.

Human Oversight, Exceptions, and Contestability

This approach does not attempt to hard-code all policy decisions. It defines the default, machine-verifiable path for enforcing clear rules, while recognising that real-world energy systems involve ambiguity and exceptions. Cases requiring human interpretation are routed to authorised stakeholders, with all decisions—automated or manual—recorded in a transparent, auditable, and contestable manner. Over time, recurring exception patterns can inform refinement of policies, allowing frequently used interpretations to be progressively formalised where appropriate.

INDIA ENERGY STACK

सत्यमेव जयते

विद्युत मंत्रालय
MINISTRY OF
POWER

IES ARCHITECTURE DOCUMENT

VERSION 0.4 | MARCH 2026