

Who Is Accountable for the Algorithm? Agentic Power Dispatch

Governance, liability and financing controls for autonomous electricity-system decisions

Authored by Chennakeshav Adya

Independent Researcher

October 2026

ABSTRACT

Artificial intelligence is moving from forecasting and advisory analytics towards systems that can form plans, select tools, coordinate workflows and recommend or execute operational actions. Electricity dispatch is a demanding test of that transition. Decisions can affect system balance, congestion, reserves, market positions, asset safety, contractual performance and public service continuity. The central governance question is therefore practical: who remains accountable when an agentic system contributes to a dispatch decision? This paper develops a Dispatch Accountability Stack for system operators, utilities, generators, storage owners, aggregators, technology providers, boards, lenders and infrastructure investors. It assigns legal and organisational responsibility to entities and named roles; classifies autonomy by the authority granted to the system; defines operating envelopes, human-oversight requirements, evidence records and incident allocation; and connects these controls to procurement, insurance, project finance and investment approval. The framework treats the algorithm as a governed operational component. Accountability remains with the organisations and people that design, approve, deploy, supervise and use it. Three figures and three tables use hypothetical author-developed assumptions to demonstrate the method. They compare autonomy levels, show how control strength should rise with operational consequence, and illustrate a staged commercial case. The values are analytical examples. They are not forecasts, market quotations, client results, legal conclusions or investment recommendations. The paper concludes that agentic dispatch can become financeable only when authority is bounded, safety constraints remain independent, human intervention is practicable, logs reconstruct each material decision, vendors accept defined obligations, and the asset retains a tested fallback. Boards and lenders should approve a progression from shadow operation to supervised recommendations and then to constrained execution. Broad autonomous authority should require sector-specific evidence, independent assurance and explicit regulatory acceptance.

JEL Classification: G32, K13, L94, M15, O33, Q40

Keywords: agentic AI, power dispatch, electricity system operations, accountability, human oversight, model risk, operational technology, project finance, artificial intelligence governance, critical infrastructure

Introduction

Power systems already use automation. Protection relays act in milliseconds, automatic generation control adjusts output, energy-management systems support control rooms, and market platforms dispatch resources through codified rules. Agentic artificial intelligence introduces a different combination of capabilities. A system may interpret an objective, retrieve data, choose analytical tools, generate a plan, communicate with other systems and recommend or initiate action. That combination can compress decision time and coordinate complex portfolios. It can also obscure which input, model, tool call or approval produced the final instruction.

Electricity operations require accountability before scale. A dispatch action can alter physical flows, market exposure and contractual performance. An error can remain local, or it can interact with congestion, correlated controls and stressed system conditions. The International Energy Agency's work on modernising grids identifies AI-enhanced solutions across planning, operations and maintenance while emphasising governance as applications become more safety critical. The United States Department of Energy is testing large-language-model and agentic capabilities for critical-infrastructure environments. European energy institutions are developing policy for digitalisation and AI while system operators have called for clear liability, auditability and human oversight.

The investment question extends beyond whether a model can produce a useful recommendation. The asset owner must know who approves the objective, who defines the authority boundary, who validates performance, who can override the system, who records evidence, who carries loss and who funds remediation. A lender must know whether the borrower can continue operating when the agent, data service or vendor fails. A regulator must be able to identify the entity that remains responsible for compliance.

This paper answers those questions through a governance and financing architecture designed for real operating decisions.

1 Define agentic dispatch precisely

Agentic dispatch is the use of an AI-enabled system that performs more than a single prediction. It can decompose an operating objective, assemble information, invoke analytical or transactional tools, evaluate alternatives and produce a recommendation or action. The objective might be to maintain a contracted delivery profile, minimise balancing cost, preserve battery state of charge, respond to a system instruction or coordinate flexible demand.

The term should not be applied to every automated control. A deterministic protection relay, an optimisation routine with fixed inputs and constraints, and a machine-learning forecast present different risks. Governance begins with a functional inventory that identifies each component and its authority. The inventory should distinguish data ingestion, forecasting, optimisation, language interfaces, planning, tool selection, action approval, execution and monitoring.

The decisive attribute is authority. A system that drafts an explanation for an operator has limited operational authority. A system that can submit a bid, change a set point or commit a battery has material authority even when its output resembles a recommendation. Management should classify the actual permission path, including API credentials, default approvals and operator response times. Descriptions such as decision support can become misleading when the workflow makes human rejection impracticable.

2 Keep accountability attached to legal entities and named roles

An algorithm cannot hold a board mandate, regulatory licence, operating certificate or contractual obligation. Responsibility must remain assigned to legal entities and named roles. The regulated operator remains responsible for its operating duties. The asset owner remains responsible for the plant and contracts. A vendor remains responsible for the obligations it accepts. Individuals remain responsible for decisions within their delegated authority.

The board should approve the risk appetite and highest permitted autonomy class. An accountable executive should own the use case and economic outcome. The control-room or market-operations leader should own operating procedures and authority. A model owner should maintain the system. An independent validator should challenge performance and limitations. Cybersecurity and operational-technology leaders should control access and resilience. Legal, compliance and insurance teams should allocate obligations before deployment.

This allocation avoids a common governance failure: collective participation without identifiable decision ownership. A committee can review evidence, yet one accountable executive should accept the residual risk. Escalation should name the person authorised to suspend the system. Incident protocols should identify the operational commander, technical investigation lead, regulatory-notification owner and financial-loss owner.

3 Classify autonomy by operational authority

Table 1. Autonomy classes for agentic power dispatch

Class	System role	Permitted authority	Human control	Minimum evidence before use
0 Observe	Monitor and explain data	No recommendation or control action	Operator uses existing procedures	Data lineage, access control and output labelling
1 Advise	Generate analysis and options	Recommendation only	Named operator decides and records disposition	Back-test, limitation statement, operator training and decision log
2 Recommend	Rank actions within an approved objective	Prepared instruction requiring affirmative approval	Operator can inspect, amend or reject before deadline	Shadow trial, calibrated confidence, feasible response time and fallback
3 Constrained execute	Execute within a bounded envelope	Pre-authorised actions inside limits	Operator supervises, receives alerts and can stop or reverse	Independent validation, hard constraints, rollback, cyber test and incident drill

Class	System role	Permitted authority	Human control	Minimum evidence before use
4 Automatic protective or reserve action	Perform narrowly defined fast response	Deterministic or tightly bounded response under specified triggers	Human sets mode and limits; intervention may follow action	Sector-approved design, independent safety layer, fail-safe state and continuous monitoring
5 Broad autonomous operation	Plan and execute across assets and markets	Wide discretion over physical and commercial actions	Human supervision may be remote or delayed	Explicit regulatory acceptance, system-level assurance and evidence exceeding current ordinary practice

The classification is an author-developed governance framework. Applicable operating rules, licences and approval requirements vary by jurisdiction and system.

The class should follow the highest authority that the system can exercise, rather than the authority management expects it to use. Credentials, network pathways and integration design determine capability. A recommendation engine connected to an auto-approval script should be classified as execution.

Class 3 can create substantial value while retaining governable boundaries. It may allow a battery to respond within power, energy, degradation, reserve and contractual limits. Hard constraints should sit outside the agent's editable reasoning process. Class 4 applies to narrow protective or reserve functions where speed is essential and deterministic safeguards dominate. Class 5 requires a much stronger evidence base because the system can change objectives, coordinate resources or take actions beyond a compact envelope.

4 Build the Dispatch Accountability Stack

Figure 1. The Dispatch Accountability Stack



The figure is an author-developed control model. Each layer requires asset-specific legal, technical, operational and contractual evidence.

The stack is cumulative. Strong model validation cannot compensate for an undefined objective. A clear objective cannot compensate for weak execution controls. Comprehensive logs cannot recover accountability if nobody owns the operating decision.

The top layer connects the use case to board risk appetite, regulated obligations and system reliability. The second assigns an accountable executive. The third translates policy into executable boundaries. The fourth validates the models, data and tools. The fifth defines human supervision and operating procedures. The sixth enforces independent technical constraints. The seventh creates evidence and allocates incident response and financial consequences.

An investment committee should assess every layer. A red rating at the authority, execution or fallback layer should stop live deployment. A weakness in economic attribution may stop capital approval even when the system is technically safe. This separation prevents safety approval from being treated as proof of value and prevents an attractive business case from bypassing operational controls.

5 Specify the operating objective and prohibited objectives

An agent requires an objective. Objectives such as minimise cost or maximise revenue are incomplete because they can conflict with reliability, contracts, equipment limits and market conduct. The approved objective should express priority: comply with system instructions; preserve safe operation; meet contractual obligations; maintain required reserves; then optimise economic outcome within risk limits.

Prohibited objectives should be equally clear. The system should not learn to withhold capacity, create misleading bids, bypass curtailment, exploit data unavailable to permitted market participants or trade outside delegated limits. It should not change its own control policy, credentials or safety constraints. Any learning or configuration update should enter a controlled change process.

Objective conflicts need deterministic resolution. When a commercial instruction conflicts with a system or safety instruction, the hierarchy should decide the outcome. The agent should identify the conflict, stop the affected action when required and escalate to the named role. The log should preserve the evidence available at the time.

6 Define the authority envelope in machine-enforceable terms

Policy documents are insufficient when execution occurs through software. The authority envelope should be translated into controls that the agent cannot amend. Limits may include assets, markets, instruments, time horizons, megawatts, state of charge, ramp rates, price bounds, credit exposure, reserve commitments, counterparty permissions and cumulative loss.

The envelope should address combinations. A series of individually permitted actions can create an impermissible aggregate position. Portfolio controls should calculate total physical and financial exposure before execution. Correlated agents should share constraint information through an independent control service rather than rely on each agent to infer the portfolio state.

Time matters. A human approval requirement is meaningful only when the operator has enough time and information to decide. For fast actions, the approved envelope must carry more of the control burden. When communication fails, the system should enter a predefined state that reflects the asset's physical and contractual needs. Safe does not always mean zero output; it may mean holding the last valid set point, following a local droop curve or reverting to a certified controller.

7 Separate the agent from the independent safety layer

The same system should not propose an action, define the limits and certify compliance. Independent safety and market-risk controls should evaluate the proposed instruction. Protection systems, equipment interlocks and certified plant controllers should remain independent where required by engineering practice and regulation.

An execution gateway can check identity, asset availability, technical limits, market permissions, credit limits and freshness of data. It should reject malformed or stale instructions and record the reason. Emergency stop and rollback pathways should operate without reliance on the agent's language model or planning service.

Independence also applies to validation. The development team can test the system, while a separate competent function should validate model performance, failure modes, security and human factors. Material changes to models, prompts, tools, data sources or permissions should trigger reassessment proportional to risk.

8 Make human oversight operationally credible

The European Union AI Act requires deployers of high-risk systems to assign human oversight to competent people with the necessary authority and support. That principle is useful beyond its direct legal scope. A person shown a complex recommendation seconds before execution may have nominal oversight and little practical control.

Credible oversight requires understandable information, adequate response time, authority to reject, a clear alternative and protection from automation bias. The interface should display the proposed action, objective, binding constraints, material inputs, confidence or uncertainty, expected consequence and reason for escalation. Operators should be trained through normal, stressed and failure scenarios.

Workload should be measured. Excessive alerts can make supervision ineffective. Long periods of correct automation can reduce operator readiness. Rotations, simulation and periodic manual operation can maintain competence. Management should test whether an operator can detect a plausible wrong recommendation, intervene within the available time and restore a safe operating mode.

9 Govern data freshness, provenance and access

Dispatch depends on time-sensitive data. Telemetry, topology, outages, weather, prices, nominations and asset constraints may change at different rates. Each input should have a freshness threshold. The agent should reject or downgrade actions when critical inputs exceed that threshold.

Provenance records should identify the source, collection time, transformation, quality flag and version. External data and model services should be covered by service levels, security requirements and continuity arrangements. A vendor outage should not prevent the asset from retrieving its own operating history.

Access should follow least privilege. Read access, recommendation access and execution access should be separated. Credentials should be short-lived where practicable, protected outside prompts and rotated under controlled procedures. Tool calls should be restricted to approved functions and parameters. A language interface should not provide an unrestricted path into operational technology.

10 Validate the complete system, not only the model

Component accuracy is insufficient. Validation should test the integrated path from data to plan, tool selection, constraint check, human decision, execution and evidence. An accurate forecast can still produce a poor dispatch recommendation if the price input is stale or the optimiser misunderstands a contract.

Tests should include normal operations, stressed grid conditions, missing data, conflicting instructions, malicious input, tool failure, network delay and unusual market prices. Historical replay should preserve the information available at the original decision time. Simulation and hardware-in-the-loop testing can examine physical consequences before live control.

Agentic systems can produce variable reasoning paths. Validation should focus on observable behaviour, permissions and outcomes. A representative test suite should be retained and repeated after material changes. The approval record should state known limitations and conditions under which the system must be disabled.

11 Control model and workflow change

An agent may depend on a foundation model, domain model, prompt library, retrieval source, optimiser, tools and orchestration logic. Any of these can change behaviour. The configuration should be versioned as a complete release, with a reproducible record of dependencies.

Vendor model updates require notice and assessment when they affect a live use case. Automatic upgrades can undermine prior validation. Contracts should provide version controls, testing access, incident notice, vulnerability remediation and transition support. The operator should retain the right to freeze, roll back or replace a component.

Drift monitoring should cover data, performance and behaviour. A model can remain statistically accurate while its recommended actions become less valuable after market or contract changes. Operating limits should tighten when uncertainty rises. Material drift should trigger investigation, retraining or revalidation before authority is restored.

12 Design the evidence ledger

Every material recommendation and action should be reconstructable. The evidence ledger should connect the objective, system version, input snapshot, retrieved documents, tool calls, proposed action, constraint results, human disposition, executed instruction, asset response, settlement and incident status.

Logs should be tamper-evident, synchronised to an authoritative clock and retained according to legal, regulatory, market and contractual requirements. Sensitive information should be protected. Access to evidence should be monitored because logs may contain system topology, credentials, personal data or trading information.

Explanation has limits. A fluent narrative generated after an event may not reflect the causal path. Evidence should prioritise actual inputs, permissions, tool calls and executed instructions. A concise operator-facing explanation can support oversight; it should not replace the technical record.

13 Allocate incident responsibility before deployment

Table 2. Illustrative incident-accountability matrix

Incident	Immediate accountable lead	Required first action	Evidence owner	Contract or financing question
Unsafe dispatch recommendation	Control-room leader	Reject, isolate system and preserve safe operation	Operations assurance	Was the recommendation inside the warranted use and tested envelope?
Unauthorised executed action	Operational-technology security lead	Block credentials, restore approved mode and assess exposure	Cybersecurity	Did access control, vendor software or internal procedure fail?
Stale or corrupted input	Data-service owner	Quarantine feed and activate approved fallback	Data governance	Which party owed data quality, availability and notice?
Model or workflow drift	Model owner	Reduce authority or suspend release	Independent validation	Was monitoring adequate and was change control followed?
Operator approval error	Accountable operations executive	Stabilise position, review interface and competence	Compliance	Did the workflow provide practical oversight and sufficient time?
Market or contractual breach	Market-compliance lead	Stop affected actions and notify under applicable rules	Legal and compliance	Who controlled the objective, limit and execution credential?
Vendor outage	Business-continuity owner	Move to tested alternative and invoke service process	Procurement	Are service credits, indemnities and transition rights adequate?

This matrix is an author-developed starting point. Final allocation depends on applicable law, licences, contracts, insurance and the facts of an incident.

The matrix prevents delay during an incident. It identifies who commands the immediate response while preserving later factual and legal investigation. The first priority is safe and compliant operation. Financial allocation follows evidence and the governing agreements.

Contracts should avoid ambiguous phrases that transfer all AI risk to a party unable to control it. Obligations should follow control: the asset owner controls operating authority; the vendor controls specified software and services; the data provider controls contracted data quality; the integrator controls agreed implementation work. Shared causes require cooperation, evidence access and defined dispute procedures.

14 Address product, professional and operational liability

The European Union's Product Liability Directive expressly includes software within the concept of a product and addresses related services, updates and the capacity of some systems to change after placement on the market. The legal effect depends on jurisdiction, implementation and facts. Organisations should obtain qualified advice for each deployment.

Operational liability may also arise through electricity regulation, grid codes, market rules, contract, negligence, employment duties, cybersecurity obligations and professional services. A vendor limitation of liability does not remove the operator's regulated duties. An operator indemnity does not excuse a vendor from statutory obligations.

The contract architecture should define intended use, excluded use, performance measures, data responsibility, security, change control, audit, incident cooperation, intellectual-property rights, insurance, indemnities, liability caps, consequential loss, termination and exit support. Caps should reflect plausible loss and insurance availability. Some obligations may require separate caps or exclusions from a general cap.

15 Treat cybersecurity as an operational and financing condition

Agentic systems expand the attack surface through data feeds, models, plugins, APIs and tool credentials. Prompt injection, malicious retrieval content, poisoned data, compromised software supply chains and stolen execution credentials can become operational risks. Critical-infrastructure guidance from CISA, NIST, NERC and national cybersecurity agencies supports secure development, supply-chain controls, segmentation and incident response.

The architecture should isolate enterprise AI services from operational technology. Communication paths should be allowlisted and monitored. The agent should receive only the minimum tools and permissions. High-consequence commands should require an independent gateway, authenticated source and context validation.

Lenders should review cyber controls because a successful attack can affect availability, revenue, remediation cost and insurance. The financing model should include realistic downtime and restoration assumptions. Insurance coverage, exclusions and incident-notification requirements should be reconciled with vendor and financing agreements.

16 Measure operational consequence before model sophistication


Governance effort should follow consequence. A system that summarises maintenance notes can tolerate different uncertainty from a system that changes a dispatch set point. The organisation should score safety, system reliability, contractual exposure, market exposure, customer impact, reversibility, speed and geographic reach.

Reversibility is important. A recommendation can be reviewed; an executed market order may be offset at a cost; a battery action may be reversed within physical limits; a protection action may occur before human intervention. The score should reflect the time available to detect and correct an error.

Correlated scale can make a modest unit-level risk material. Many agents using similar models and signals may act together. Portfolio testing should examine common data outages, shared model errors, crowded trading responses and simultaneous fallback. System operators may need visibility over aggregated autonomous behaviour.

17 Use a technology-readiness curve for authority

Figure 2. Illustrative readiness requirements by autonomy class



Scores are hypothetical author-developed values on a 0 to 100 scale. They illustrate how evidence and controls should strengthen as authority rises.

The curve shows a threshold concept rather than a universal score. Recommendation systems need credible validation and oversight. Constrained execution requires independent limits, cyber assurance, fallback and incident drills. Narrow automatic action requires sector-grade engineering evidence. Broad autonomy requires evidence close to the highest available standard across every dimension.

Progress should pause when one dimension lags. Strong technical performance cannot compensate for an unavailable fallback. A comprehensive policy cannot compensate for untested execution. The gate should assess the weakest critical control.

18 Build an economic case that prices control cost


Agentic dispatch can create value through faster response, portfolio coordination, lower imbalance cost, improved storage use, reduced curtailment, better reserve participation and lower routine workload. The case should deduct the full cost of data, integration, validation, control engineering, cybersecurity, licences, training, monitoring, insurance and fallback.

Management should separate realised cash, avoided cost, capacity value and option value. A modelled improvement is not realised cash. A time-saving estimate should identify whether headcount, overtime, response capacity or quality changes. Reliability benefits should use an approved valuation method and avoid double counting with insurance or contractual protections.

The baseline should be reconstructable. Shadow operation can compare recommendations with actual decisions while controlling for information that became available later. A controlled live test can measure outcomes within a small authority envelope. The commercial case should report adverse outcomes and missed opportunities as well as successes.

19 Illustrate the staged value and control case

Figure 3. Hypothetical annual net-value progression for an agentic dispatch programme



Values are author-developed illustrations in USD millions. They do not represent a client, asset, market forecast or current quotation.

The hypothetical shadow stage has negative net value because assurance, integration and monitoring cost exceed direct benefit. That result can be rational when the stage creates evidence and avoids premature authority. The supervised stage produces limited positive value. Constrained execution increases value after the organisation proves controls. Portfolio scale creates the largest gross value and the largest risk provision.

Investment approval should use stage-specific capital. The next tranche should depend on evidence: decision quality, operator intervention, execution reliability, incident performance, cash attribution and control completion. A programme that fails a gate should reduce authority, redesign the workflow or stop.

20 Connect governance to project finance

Project-finance lenders assess construction, technology, operations, revenue, counterparty, cybersecurity and legal risks. Agentic dispatch can affect several simultaneously. It may improve operating margin while introducing dependency on data, software, cloud infrastructure or a specialist vendor.

The financing case should identify whether agentic value is necessary to service debt. A prudent base case can treat early benefits as upside until evidence matures. If the benefit supports debt sizing, lenders may require

performance history, independent review, conservative haircuts, reserve accounts, service continuity, step-in rights and replacement arrangements.

Finance documents should address material changes in operating systems, cyber incidents, loss of key licences, vendor termination and compliance breaches. Direct agreements may be relevant where a critical vendor cannot be readily replaced. The project should retain data, configuration and transition support sufficient for continuity.

21 Allocate risk in procurement

Procurement should start with the use case and control architecture, rather than a generic AI platform request. The specification should state operating objective, autonomy class, environments, interfaces, performance measures, authority limits, safety controls, logging, cybersecurity, testing and exit requirements.

Acceptance should occur in stages. Documentation and security review precede sandbox tests. Historical replay precedes shadow operation. Shadow evidence precedes supervised use. Constrained execution follows only after approval. Each stage should have measurable entry and exit criteria.

Commercial models can combine fixed fees, service levels and carefully designed outcome components. Outcome fees require a frozen baseline and protection against risk shifting. A vendor should not receive payment for gross trading gains while the asset owner bears losses, degradation and compliance exposure. Gainshare should be based on verified net value within agreed limits.

22 Preserve market integrity

An agent that participates in electricity markets must comply with bidding, scheduling, disclosure, recordkeeping and market-conduct rules. The organisation should prevent coordination based on confidential information, manipulative patterns and actions inconsistent with physical capability.

Market surveillance should incorporate agent identity, strategy version and tool calls. Limits should cover price, volume, frequency, cancellation and concentration. New strategies should be tested under stressed liquidity and unusual prices. Compliance should be able to suspend a strategy independently of the commercial owner.

When multiple agents interact, emergent behaviour can create risk without an explicit instruction to manipulate. Governance should examine collective outcomes. System operators and regulators may require new reporting or testing as autonomous participation grows.

23 Address regional regulatory differences

The same architecture may operate under different legal regimes. The EU AI Act establishes risk-based obligations and human-oversight requirements for specified high-risk systems. Electricity-system applications require use-case-specific classification. The Product Liability Directive, cybersecurity rules and national energy regulation can add obligations.

North American operators must consider FERC-approved reliability standards, NERC standards, regional market rules and state requirements. NERC standards address operational planning, real-time reliability, monitoring, situational awareness and supply-chain cybersecurity. A deployment should map each function to applicable standards and responsible entities.

GCC, UK, Asian and other markets have distinct licensing, grid-code, market and data requirements. Cross-border vendors should not assume that one approval package satisfies every jurisdiction. The control framework can remain consistent while the legal and technical evidence changes by location.

24 Establish board reporting

Board reporting should focus on authority, consequence and evidence. It should state which use cases are live, their autonomy class, assets and markets, accountable executives, last validation, material changes, overrides, incidents, near misses, cyber events, fallback tests, realised value and open conditions.

Aggregate accuracy is insufficient. Reporting should include performance during stressed conditions and the intervals with greatest consequence. It should disclose when human approval becomes routine or unusually rapid, because that may indicate automation bias. It should show whether operators used the fallback and whether it worked.

The board should approve escalation thresholds. A serious incident, loss of evidence, control bypass or unapproved authority expansion should trigger immediate review. Management should certify that the live configuration matches the approved configuration.

25 Give lenders and investors a focused diligence pack

Table 3. Lender and investor diligence for agentic dispatch

Diligence area	Core evidence	Adverse finding	Financing response
Authority	Autonomy class, permissions, limits and approval records	Execution capability exceeds approved class	Suspend live authority; require independent remediation
Technical assurance	Replay, simulation, stress tests, failure modes and validation	Tests omit stressed or correlated conditions	Exclude benefit from base case; fund additional assurance
Human oversight	Competence, interface, response time, overrides and drills	Operator cannot intervene within the decision window	Reduce autonomy; redesign workflow and staffing
Cybersecurity	Architecture, segmentation, credentials, supply-chain controls and incident response	Agent has broad operational-technology access	Make remediation a condition precedent
Commercial value	Frozen baseline, realised cash, full cost and downside	Benefit relies on modelled gross value	Apply haircut or treat as upside only
Vendor dependency	Service levels, version control, data portability, escrow and exit	No viable fallback or replacement route	Require reserve, direct agreement or alternative supplier
Liability and insurance	Warranties, indemnities, caps, exclusions and coverage	Loss falls into an uninsured contractual gap	Reallocate risk, increase reserve or decline reliance
Continuity	Manual mode, certified controller, recovery time and test evidence	Fallback is undocumented or untested	Prevent debt reliance and require successful drill

The table is a commercial diligence framework and does not replace legal, technical, insurance or regulatory advice.

The pack should be proportionate to authority. A class 1 advisory tool may need a concise assessment. A class 3 execution system requires independent technical, cyber, legal and commercial review. A system material to debt service requires ongoing reporting and change controls.

Investors should ask whether governance enables value rather than only adds cost. Clear authority and evidence allow the asset to expand successful use cases. Weak governance can force a valuable system to remain in permanent pilot status because management, lenders and regulators cannot rely on its outcomes.

26 Design insurance around the actual loss pathways

Relevant coverage may include property damage, machinery breakdown, business interruption, cyber, technology errors and omissions, professional indemnity and directors' and officers' liability. Coverage varies. Policies can exclude infrastructure events, operational technology, war, systemic cyber incidents, fines, contractual liability or known vulnerabilities.

The insured should map scenarios to policies, deductibles, sublimits, waiting periods and exclusions. Vendor insurance certificates are insufficient without reviewing the scope and claims basis. Contractual indemnities should align with available insurance while preserving recourse for material obligations.

Incident-notification clauses should be coordinated. A delayed notice to the insurer, regulator, system operator, lender or vendor can impair rights. The response plan should assign each notification and preserve privilege where applicable.

27 Create a four-stage commercialisation roadmap

Stage one is controlled shadow operation. The system receives real inputs and generates recommendations without execution rights. The organisation tests data, behaviour, operator comprehension, evidence and hypothetical economics.

Stage two is supervised recommendation. Named operators can approve, modify or reject actions. The programme measures intervention, timing, errors and realised results. The system remains inside a narrow asset and market scope.

Stage three is constrained execution. The agent receives limited credentials through an independent gateway. Hard limits, alerts, stop controls, fallback and incident response are proven. Authority expands only after independent validation and governance approval.

Stage four is scaled portfolio operation. Scale should preserve asset-level constraints and portfolio risk controls. Common-mode failures, correlated actions and vendor concentration receive additional testing. The board approves the scale threshold and system-wide exposure.

28 Apply explicit stop conditions

Stop conditions should be defined before live use. They may include loss of critical data, time synchronisation failure, unexpected tool access, breached limits, anomalous action frequency, validation failure, unexplained performance drift, cyber alert, evidence-ledger failure or inability to reach a supervisor.

The response should be automatic where delay creates unacceptable risk. Other conditions can require rapid human judgement. The stop should move the asset to a known operating state and preserve evidence. Restart should require documented approval after root-cause analysis and remediation.

Near misses matter. A rejected unsafe recommendation or prevented credential misuse demonstrates control performance and reveals system weakness. Reporting should reward detection rather than suppress inconvenient evidence.

29 Use accountability as a value-creation asset

An accountable architecture can accelerate commercial adoption. Operators understand their authority. Vendors understand acceptance criteria. Lenders can diligence continuity. Insurers can assess loss pathways. Regulators can identify responsible entities. Management can attribute realised value.

The evidence ledger also improves optimisation. It shows where recommendations fail: data, model, contract, operator, market liquidity, equipment or execution. Capital can then target the binding constraint. The organisation can expand authority when evidence supports it and reduce authority when conditions change.

This approach supports transaction diligence. A buyer can assess whether an AI-enabled operating claim is transferable, whether licences and data rights survive change of control, whether the vendor can terminate, and whether historical value can be reproduced. The governance system becomes part of operational quality and valuation.

30 Make the accountability answer visible

The answer to who is accountable should be available before an incident. The regulated entity remains responsible for regulated operation. The board sets risk appetite. The accountable executive owns the use case. Operations owns the live decision process. The model owner maintains the system. Independent assurance challenges it. Vendors and data providers carry defined contractual and statutory obligations. Cybersecurity protects the execution path. Finance measures value and resilience.

These roles can vary by organisation, yet responsibility cannot disappear into the algorithm. Every material action should connect to an approved objective, authority envelope, responsible entity, named role and evidence record.

Conclusion

Agentic power dispatch can improve coordination and speed across increasingly complex electricity systems. Its commercial promise depends on a credible answer to accountability. That answer is an operating architecture, rather than a disclaimer.

The Dispatch Accountability Stack connects board mandate, use-case ownership, bounded authority, model and data assurance, human oversight, independent guardrails, evidence and incident allocation. Autonomy should rise in stages as assurance, resilience and realised value mature. The complete control cost should be included in the investment case. Financing should remain serviceable under agent or vendor failure.

The organisations most prepared to use agentic dispatch will know exactly which decisions the system may take, who can stop it, how each action is reconstructed, who bears each loss and how operation continues when technology fails. That discipline creates the foundation for safe scale, regulatory confidence and financeable value.

Frequently Asked Questions

Who is legally accountable when an AI agent contributes to a dispatch decision?

Accountability remains with the relevant legal entities and people under applicable regulation, licences, contracts and delegated authority. The algorithm is a governed component. Use-case-specific legal advice is required because responsibility depends on jurisdiction and facts.

What makes a dispatch system agentic?

An agentic system performs a sequence of tasks such as interpreting an objective, retrieving data, choosing tools, forming a plan and recommending or executing an action. The governance classification should follow the authority it can exercise.

Can a human approval button provide adequate oversight?

Adequate oversight requires competence, understandable information, sufficient response time, authority to reject, an alternative action and protection from automation bias. A nominal approval step may be ineffective when the operator cannot meaningfully assess the recommendation.

When should an agent be allowed to execute dispatch actions?

Execution should follow successful replay, simulation, shadow operation, independent validation, cyber testing, operator training, fallback drills and approval of a machine-enforceable authority envelope. Initial authority should remain narrow.

What records should be retained?

The evidence ledger should connect the objective, system version, inputs, retrieved material, tool calls, recommendation, constraint results, human disposition, executed instruction, asset response and settlement. Retention and access should follow applicable requirements.

How should lenders treat projected value from agentic dispatch?

Early modelled value can be treated as upside. Inclusion in a debt base case requires repeatable evidence, conservative assumptions, continuity controls and a financing model that remains resilient when the benefit is reduced or removed.

What should happen when the agent or data service fails?

The asset should enter a predefined, tested operating mode using certified local control, an alternative service or competent manual operation. The response should preserve safety, compliance, contractual performance and evidence.

What is the first decision a board should make?

The board should approve the risk appetite, prohibited uses and highest permitted autonomy class. Management can then design objectives, controls, evidence and stage gates within that mandate.

References

1. National Institute of Standards and Technology, *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*, 2023. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>
2. National Institute of Standards and Technology, *AI RMF Playbook*. https://airc.nist.gov/AI_RMF_Knowledge_Base/Playbook
3. National Institute of Standards and Technology, *NIST AI Resource Center*. <https://airc.nist.gov/>
4. National Institute of Standards and Technology, *Cybersecurity Framework 2.0*, 2024. <https://www.nist.gov/cyberframework>
5. U.S. Department of Energy, *AI for Energy*. <https://www.energy.gov/cet/articles/ai-energy>
6. U.S. Department of Energy Office of Cybersecurity, Energy Security, and Emergency Response, *Stormbreaker: AI Testbed for Critical Infrastructure*, 2026. <https://www.energy.gov/ceser/stormbreaker>
7. U.S. Department of Energy, *Grid Modernization Initiative and Grid Test Bed Inventory*. <https://www.energy.gov/gdo/grid-modernization-initiative>
8. International Energy Agency, *Modernising Grids in the Age of Electricity*, 2026. <https://www.iea.org/reports/modernising-grids-in-the-age-of-electricity>
9. International Energy Agency, *Energy and AI*, 2025. <https://www.iea.org/reports/energy-and-ai>
10. International Renewable Energy Agency, *Digitalisation and AI for power system transformation: Perspectives for the G7*, 2025. <https://www.irena.org/Publications/2025/Oct/Digitalisation-and-AI-for-power-system-transformation-Perspectives-for-the-G7>
11. European Commission, *Digitalisation of the energy system*. https://energy.ec.europa.eu/topics/eus-energy-system/digitalisation-energy-system_en
12. European Commission, *AI Act*. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>
13. European Union, *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence*. <https://eur-lex.europa.eu/eli/reg/2024/1689/oj>
14. European Union, *Directive (EU) 2024/2853 on liability for defective products*. <https://eur-lex.europa.eu/eli/dir/2024/2853/oj>
15. European Network of Transmission System Operators for Electricity, *Vision 2050*. <https://www.entsoe.eu/publications/system-development-reports/vision-2050/>
16. European Network of Transmission System Operators for Electricity, *Response to the European Commission consultation on AI in energy*, 2025. <https://www.entsoe.eu/>
17. Office of Gas and Electricity Markets, *Ethical AI use in the energy sector*, updated 2026. <https://www.ofgem.gov.uk/publications/ethical-ai-use-energy-sector>
18. Office of Gas and Electricity Markets, *AI assurance in the energy sector*, 2026. <https://www.ofgem.gov.uk/>
19. International Organization for Standardization, *ISO/IEC 42001: Artificial intelligence management systems*. <https://www.iso.org/standard/81230.html>
20. International Organization for Standardization, *ISO 31000 Risk management*. <https://www.iso.org/iso-31000-risk-management.html>
21. North American Electric Reliability Corporation, *Reliability Standards*. <https://www.nerc.com/pa/Stand/Pages/Default.aspx>
22. North American Electric Reliability Corporation, *TOP-001 Transmission Operations*. <https://www.nerc.com/pa/Stand/Reliability%20Standards/TOP-001-6.pdf>
23. North American Electric Reliability Corporation, *TOP-010 Real-time Reliability Monitoring and Analysis Capabilities*. [https://www.nerc.com/pa/Stand/Reliability%20Standards/TOP-010-1\(i\).pdf](https://www.nerc.com/pa/Stand/Reliability%20Standards/TOP-010-1(i).pdf)
24. North American Electric Reliability Corporation, *CIP-013 Supply Chain Risk Management*. <https://www.nerc.com/pa/Stand/Reliability%20Standards/CIP-013-2.pdf>
25. North American Electric Reliability Corporation, *Reliability Guideline: Situational Awareness for the System Operator*. <https://www.nerc.com/>
26. Federal Energy Regulatory Commission, *Technical Conference on increasing market and planning efficiency through improved software*. <https://www.ferc.gov/>
27. California Independent System Operator, *Automated Dispatch System Guide*. <https://www.caiso.com/>
28. California Independent System Operator, *Exceptional Dispatch*. <https://www.caiso.com/market-operations/products-services/exceptional-dispatch>
29. Cybersecurity and Infrastructure Security Agency, *Secure by Design*. <https://www.cisa.gov/securebydesign>
30. Cybersecurity and Infrastructure Security Agency, *Industrial Control Systems*. <https://www.cisa.gov/topics/industrial-control-systems>
31. Cybersecurity and Infrastructure Security Agency, National Cyber Security Centre and international partners, *Guidelines for Secure AI System Development*, 2023. <https://www.cisa.gov/resources-tools/resources/guidelines-secure-ai-system-development>
32. National Cyber Security Centre, *Guidelines for secure AI system development*. <https://www.ncsc.gov.uk/collection/guidelines-secure-ai-system-development>
33. World Bank, *Project Bankability: Risk Allocation and Impact on PPA Prices*, 2023. <https://documents1.worldbank.org/curated/en/099120623171525006/pdf/P1742021cf52b60e6196b81854984124388.pdf>
34. World Bank, *Making Greenfield Projects Bankable*, 2021. <https://blogs.worldbank.org/en/ppps/making-greenfield-projects-bankable>
35. International Finance Corporation, *Utility-Scale Solar Photovoltaic Power Plants: A Project Developer's Guide*. <https://www.ifc.org/en/insights-reports/2015/utility-scale-solar-photovoltaic-power-plants>
36. Organisation for Economic Co-operation and Development, *OECD AI Principles*. <https://oecd.ai/en/ai-principles>
37. Institute of Electrical and Electronics Engineers, *IEEE 7000-2021 Model Process for Addressing Ethical Concerns during System Design*. <https://standards.ieee.org/ieee/7000/6781/>
38. Basel Committee on Banking Supervision, *Principles for Operational Resilience*, 2021. <https://www.bis.org/bcbs/publ/d516.htm>

About the Author

Authored by Chennakeshav Adya

Independent Researcher

Chennakeshav (CK) is a corporate finance and investment banking executive with 25+ years of global experience in deal origination, structuring and execution across M&A, growth capital and corporate strategy. He has led value-creation mandates for founders, corporates and funds — bridging the boardroom view to hands-on execution and close.

His career spans Morgan Stanley, HSBC, Lloyds Banking Group, EWEC, ADQ portfolio companies and Emirates Growth Fund, across TMT, real estate, fintech, deeptech, cleantech, infrastructure and energy. He has partnered with C-suite leaders, private equity and venture funds, sovereign wealth funds and family offices to finance complex fund raises and scale-up ventures, and has led M&A due diligence, post-merger integration and business-transformation initiatives to create value.

At Matchpoint Partners he is Managing Partner, leading the firm's corporate finance, M&A and capital-raising practice. He holds an MBA from London Business School, an engineering degree from VTU and a Master of Laws (LLM, in progress) from UCL London.

An active start-up mentor, CK mentors at Techstars, DIFC FinTech Hive, Startup Grind, Founder Institute and IN5, serves as Entrepreneur Mentor in Residence (EMiR) at London Business School, and judges the Entrepreneurship World Cup.

<https://www.linkedin.com/in/ckadya/>

<https://www.matchpoint-partners.com/team/ck-adya.html>

This paper is part of a continuing series on the structure of private and alternative markets. The views expressed are the author's own. The paper is for information only, describes market structure in general terms, and does not constitute investment, legal, tax or regulatory advice or a recommendation in respect of any security, vehicle or counterparty.