

The Water Network's Working Model: Turning Digital Twins into Capital Decisions

A governance and investment framework for hydraulic evidence, uncertainty and network capital allocation

Authored by Chennakeshav Adya

Independent Researcher

September 2026

Abstract

Background. Water utilities are combining network telemetry, geographic information, hydraulic models, work history and customer data to create digital representations of their systems. These models can improve operating awareness and reveal interactions that are difficult to see from isolated asset records. A visually sophisticated twin does not by itself establish that a renewal, reinforcement or resilience investment is justified. Objective. This paper develops a capital-decision framework that connects a digital twin's technical credibility to the economic choice it is asked to support. Method. It synthesises current digital-water practice, asset-management standards, hydraulic-model guidance, digital-twin credibility research and infrastructure investment principles. It then defines a decision hierarchy, validation envelope, evidence register, option-comparison method, uncertainty adjustment, benefits ledger, governance model and staged implementation roadmap. Findings. A twin should be accepted for a defined use only when its data lineage, calibration, validation, uncertainty, operating integration and change control are proportionate to the financial consequence of that use. Renewal, capacity, leakage, energy and resilience decisions require different evidence thresholds. Model outputs should expose alternatives and confidence ranges; authorised people should approve capital. Implications. Utilities can convert a twin into a working capital-allocation system by linking every recommendation to a physical mechanism, validated evidence, economic counterfactual, accountable owner and post-investment learning loop. All unattributed numerical exhibits are hypothetical decision tools; they are not forecasts, valuations, offers or descriptions of a named utility or transaction.

JEL Classification: G31, L95, O32, Q25, Q55

Keywords: water network, digital twin, hydraulic model, asset management, capital allocation, resilience, leakage, uncertainty, investment appraisal, water utility

This paper is an educational and structural analysis prepared for research purposes. It is not investment advice, an offer, or a solicitation. It contains no client information.

1. THE CAPITAL QUESTION

A water-network digital twin earns its place in an investment process when it improves a consequential decision. The relevant question is specific: should the utility renew a main, reinforce a pressure zone, add storage, reconfigure pumping, create redundancy or wait for better evidence? The

model should compare credible alternatives, quantify consequences and show where uncertainty could reverse the preferred choice.

The investment committee should define the decision, boundary, service objective, counterfactual, time horizon, capital constraint and approval authority before requesting an answer from the twin. Otherwise the programme risks producing an impressive representation without a disciplined use case.

2. A DECISION INSTRUMENT

A twin is a decision instrument composed of physical assets, observations, models, interfaces, governance and people. Current United Kingdom guidance defines a digital twin through a fit-for-purpose digital representation and two-way communication appropriate to the required decision timeframe [1]. This definition places purpose and timeliness ahead of visual fidelity.

For capital allocation, the representation should explain how an intervention changes hydraulic behaviour, service risk, operating cost and lifecycle value. It should also retain the evidence that connects those changes to the recommended expenditure.

3. REPRESENTATION IS NOT REALITY

Every network model simplifies. Pipe condition is partially observed. Demands vary. Valves may be misrecorded. Sensors drift. Customer behaviour changes. Construction records can be incomplete. A model can reproduce selected observations while remaining wrong about the mechanism that matters for a proposed investment.

Decision governance should therefore prohibit language that treats a simulated result as an observed fact. Each output should state its input period, calibration conditions, validation evidence, confidence range, excluded mechanisms and known limitations.

4. THE VALIDATION ENVELOPE

The validation envelope is the combination of geography, operating state, demand range, asset configuration, time resolution and decision type within which the model has demonstrated acceptable performance. A model calibrated for average-day planning may not be credible for a transient-event decision. A model validated in one pressure zone may not support a citywide leakage estimate.

Capital approval should remain inside the envelope or include additional field testing. The envelope belongs in the investment paper, not solely in technical documentation.

5. DECISION-SPECIFIC CREDIBILITY

Credibility is use dependent. A pump-scheduling decision can require accurate short-interval flow, pressure and energy data. A 25-year trunk-main reinforcement decision needs robust demand scenarios, network topology, asset life, climate exposure and option staging. A leak-localisation use case needs sensitivity to night-flow and pressure signatures.

NIST's digital-twin work emphasises verification, validation and uncertainty quantification across the lifecycle [2]. The practical implication is a separate credibility test for each decision rather than a single certificate for the whole platform.

6. THE PHYSICAL-TO-CAPITAL ARCHITECTURE

The working architecture begins with the network and its operating states. Sensors, inspections, work orders, customer contacts and laboratory observations create evidence. Data services reconcile identity and time. Hydraulic, statistical and optimisation models convert evidence into scenarios. Economic models translate scenarios into cash flow, risk and service outcomes. Governance determines which outputs may inform, recommend or automate an action.

This chain makes failure visible. A weak asset identifier can undermine a sound hydraulic equation; a sound simulation can still be paired with a weak valuation.




Figure 1. From physical network evidence to an authorised capital decision

Generic decision architecture; actual governance, asset boundaries and approval authorities depend on the utility and jurisdiction.

7. ASSET IDENTITY

Pipes, valves, pumps, reservoirs, meters, customers and work orders need stable identifiers. Duplicate or inconsistent identities can attach an incident to the wrong asset, fragment failure history and distort risk rankings. Spatial proximity alone is insufficient where parallel mains or reconstructed alignments exist.

The data model should preserve source identifiers, reconciliation rules, confidence and effective dates. Asset merges and splits require controlled lineage because historical evidence can otherwise be silently reassigned.

8. NETWORK TOPOLOGY

Hydraulic behaviour depends on connectivity and status. A closed valve recorded as open can alter simulated paths, pressures and available fire flow. Topology assurance should combine records, field confirmation, operational knowledge and anomaly analysis.

The programme should maintain a topology exception register. Exceptions with a material effect on the proposed investment should be resolved, tested as scenarios or reflected in the uncertainty allowance. A visually complete map should not conceal uncertain connectivity.

9. DEMAND EVIDENCE

Demand allocation affects flows, pressure, residence time and reinforcement need. Billing data, district-metered-area balances, smart-meter profiles and land-use forecasts can each contribute, subject to purpose and quality. Historical consumption may not represent future peak patterns or new industrial loads.

Capital scenarios should separate committed connections, evidenced development, policy-led growth and speculative demand. The model should show which demand category drives the investment date.

10. SENSOR FITNESS

Sensors should be assessed for location, range, accuracy, frequency, availability, clock alignment and maintenance. A large sensor count does not guarantee observability. Correlated gaps or devices placed away from decision-sensitive locations can leave important parameters unidentified.

The utility should design monitoring around the decisions it expects to make. Temporary loggers and targeted field tests can create more value than indiscriminate permanent instrumentation.

11. DATA TIMELINESS

Decision timeframes differ. Operational control may need minute-level data; renewal planning may rely on monthly or annual patterns. Two-way communication should be fast enough for the use case and slow enough to preserve quality controls appropriate to the consequence.

The architecture should define latency, reconciliation and late-data rules. A provisional data stream can support situational awareness while a reconciled dataset supports investment evidence.

12. DATA QUALITY DIMENSIONS

Completeness, accuracy, consistency, timeliness, uniqueness and lineage should be measured separately. A single quality score can hide a critical defect. The proposed investment may depend on a small number of pressure sensors whose accuracy matters more than the completeness of unrelated attributes.

Quality reporting should be weighted by decision sensitivity. The investment pack should identify the inputs that dominate the result and the evidence supporting them.

13. HYDRAULIC CALIBRATION

Calibration adjusts uncertain parameters so simulated behaviour aligns with observations under defined conditions. EPA guidance treats calibration and validation as essential and links the required standard to the intended use [3]. Calibration should include flows, pressures, tank levels, pump operations and boundary conditions relevant to the decision.

Parameter changes should remain physically plausible. An excellent statistical fit created by unrealistic roughness or demand patterns can weaken causal credibility.

14. INDEPENDENT VALIDATION

Validation tests the configured model against observations not used to calibrate it. The test should include different demand periods, operational states and relevant events. Performance criteria should be set before reviewing the validation result.

Where independent data are scarce, the utility can reserve monitoring periods, zones or events. Repeated tuning against every observation may improve fit while removing a meaningful test of predictive performance.

15. RESIDUAL DIAGNOSIS

A residual is the difference between simulated and observed behaviour. Its pattern can reveal missing demand, bad topology, sensor bias, operational misunderstanding or model-structure error. Average error alone can cancel material positive and negative deviations.

The review should examine residuals by location, time, operating state and magnitude. Persistent structure is a diagnostic signal and should influence confidence in the proposed capital mechanism.

16. VERIFICATION OF COMPUTATION

Verification asks whether the equations, code, interfaces and numerical processes perform as intended. It includes unit handling, solver settings, boundary transfer, version compatibility and reproducibility. Validation asks whether the representation is adequate for the real-world use.

Both are required. A correctly implemented model can be an inadequate representation; a plausible result can emerge from a computational defect. Release records should retain software version, configuration and test evidence.

17. UNCERTAINTY TAXONOMY

Uncertainty can arise from measurement, parameters, future demand, asset condition, model form, climate, operations and human response. These sources should be named because they behave differently. More sensors may reduce measurement uncertainty while leaving demand growth unresolved.

The capital case should distinguish reducible uncertainty from structural uncertainty. It should state what additional evidence would change the decision and what uncertainty must be carried commercially.

18. SENSITIVITY IS NOT PROBABILITY

Sensitivity analysis changes assumptions and observes the result. Probabilistic analysis assigns distributions and estimates outcome ranges. Scenario analysis constructs internally consistent futures. These tools answer different questions.

The investment paper should avoid converting a sensitivity range into an implied probability without evidence. Where probability is weak, structured scenarios and break-even tests can provide a more transparent decision basis.

19. DECISION THRESHOLDS

A threshold identifies when an option becomes acceptable or preferred. Examples include minimum pressure, maximum interruption exposure, required resilience, whole-life cost advantage or risk-adjusted net present value. Thresholds should be agreed before optimisation to avoid choosing criteria that favour a model-produced answer.

The twin should report proximity to the threshold and the uncertainty around it. A narrow apparent margin may justify a pilot, monitoring or staged option.

20. EVIDENCE THRESHOLDS BY USE

Different uses deserve different evidence standards. A low-cost operating trial can tolerate more uncertainty than an irreversible trunk-main project. An emergency action can accept provisional evidence under explicit incident authority. A regulatory filing may require reproducibility and independent assurance.

The following matrix illustrates a proportional approach. The thresholds are governance examples and should be adapted to local obligations.

Table 1. Illustrative evidence thresholds by capital and operating use

Decision use	Primary evidence	Required model test	Economic test	Appropriate authority
Pump schedule	Current telemetry and equipment limits	Short-interval validation	Energy and service trade-off	Operations authority
Leak investigation	Flow, pressure and field confirmation	Local event sensitivity	Avoided loss and repair cost	Network manager
Renewal bundle	Condition, failures and hydraulic consequence	Zone validation and downside cases	Lifecycle cost and risk	Asset investment committee
Capacity reinforcement	Committed and scenario demand	Peak and contingency validation	Timing, option and utilisation	Capital committee
Resilience scheme	Hazard, dependency and recovery evidence	Stress and failure scenarios	Expected loss and service value	Board or delegated committee

21. THE DECISION-RIGHT LADDER

The platform can inform, diagnose, recommend, schedule or control. Each level requires stronger assurance. An informative dashboard may display observed conditions. A diagnostic model may identify likely mechanisms. A recommendation engine may rank interventions. Automated control directly changes the physical network.

Capital approval should remain with named authorised people. The twin should preserve the alternatives, evidence and rationale considered, including cases where management rejects the ranking.

22. RENEWAL DECISIONS

Renewal should combine failure likelihood, consequence, hydraulic criticality, service history, condition evidence, access cost and intervention opportunity. Age alone is a weak proxy. The twin can simulate service consequences and reveal dependencies that are absent from isolated asset scores.

The economic comparison should include repair, planned replacement, coordinated works, monitoring and run-to-failure where lawful and safe. Bundling can change mobilisation cost and outage risk.

23. CAPACITY DECISIONS

Capacity investment is exposed to demand timing, connection certainty, operating configuration and alternative capacity. The twin can test whether reinforcement is required, where the bottleneck lies and whether reconfiguration, pressure management, storage or demand response can defer construction.

The case should state the demand event that triggers expenditure. Committed load and speculative pipeline should receive different weights. Staging can preserve the option to expand when evidence strengthens.

24. LEAKAGE DECISIONS

Leakage reduction requires a physical and economic baseline. The model can prioritise zones, pressure interventions and investigation routes, while field evidence confirms actual leaks. Apparent losses and data defects should be separated from physical loss.

Value depends on marginal water cost, scarcity, energy, treatment capacity, service risk and regulatory incentives. Retail tariff is generally an unsuitable universal proxy for avoided cost.

25. PRESSURE MANAGEMENT

Pressure affects leakage, bursts, customer service and energy. A proposed change can shift risk across elevations and operating states. The twin should test minimum and maximum pressure, transients, fire-flow obligations, controls and failure modes.

Economic value should include valve or controller capex, maintenance, avoided repairs, water and energy effects, and any customer consequence. Field trials can validate the assumed response before wide deployment.

26. PUMPING AND ENERGY

Pump scheduling can reduce energy cost while preserving storage and pressure constraints. Tariffs, carbon intensity, equipment efficiency and maintenance exposure may vary over time. A lower energy schedule can create vulnerability if it removes operational headroom.

The model should value energy, service resilience and asset wear together. The control recommendation should include safe fallback modes when data, communications or optimisation services fail.

27. STORAGE DECISIONS

Storage provides balancing, resilience, pressure support and operational flexibility. Additional volume can also create water-quality, land, embodied-carbon and maintenance consequences. The twin can test how storage interacts with pumping, demand and contingency conditions.

The appraisal should compare physical storage with network reinforcement, operating changes, interconnection and demand measures. Value can arise from deferred investment as well as direct service improvement.

28. WATER-QUALITY INTERACTION

Hydraulic changes can alter residence time, direction, mixing and disinfectant conditions. Capital choices should therefore connect hydraulic and water-quality analysis where relevant. A solution to a pressure or capacity problem should not create an unmanaged quality risk.

The investment process should identify regulatory, sampling and operational requirements. Specialist review remains necessary where the twin's representation does not cover the relevant physical or chemical mechanism.

29. RESILIENCE DECISIONS

Resilience analysis should test hazards, component failures, common causes, dependencies, recovery and customer consequences. A single-component outage study may overlook power, communications, access, treatment or supply-chain dependencies.

The twin can expose cascading effects and compare redundancy, isolation, mobile response, spares, alternative supply and faster recovery. The decision should value both avoided expected loss and severe plausible outcomes.

30. CLIMATE AND SOURCE UNCERTAINTY

Climate can affect source availability, demand, flooding, ground movement, water quality and energy conditions. Long-lived assets should be tested under multiple coherent futures. A single central forecast can disguise the option value of adaptable designs.

The capital case should show which climate assumptions matter, how frequently they will be reviewed and whether a staged pathway performs acceptably across scenarios.

31. INTERDEPENDENCY MAPPING

Water networks depend on electricity, communications, chemicals, transport, workforce and upstream assets. Other sectors depend on water. Digital representations can connect these dependencies if ownership, interfaces and information rights are clear.

Capital prioritisation should identify common failure points and investments that reduce multiple risks. Sensitive infrastructure information should be governed proportionately and shared only with authorised parties.

32. ECONOMIC COUNTERFACTUAL

Every benefit needs a counterfactual: what happens without the proposed action? The counterfactual can include reactive maintenance, rising failures, temporary supply, service penalties, energy cost, planned works and later replacement. It should reflect realistic management behaviour rather than passive deterioration by default.

The twin can improve the counterfactual by simulating network consequences. Finance still needs to translate those consequences into attributable cash and service value.

33. BENEFITS LEDGER

The ledger should identify each benefit, beneficiary, mechanism, evidence, timing, owner and financial treatment. Categories can include avoided repair, reduced leakage, energy saving, capacity deferral, service improvement, resilience, safety and regulatory performance.

Benefits should not be added when they overlap. Avoided water production, deferred capacity and leakage incentive can reflect related physical volume and require reconciliation.

34. CASH VERSUS PUBLIC VALUE

Some outcomes create utility cash; others create customer, environmental or social value. Both can support an investment decision, subject to transparent classification. Debt capacity should rely on enforceable and available cash rather than the entire public-value estimate.

Funding structure can combine utility capital, public support, performance payments or concessional instruments where beneficiaries and cash sources differ. The twin provides evidence; it does not turn non-cash value into cash.

35. WHOLE-LIFE COST

Whole-life cost includes design, data preparation, sensors, licences, integration, cloud or compute, cyber controls, model maintenance, field validation, training, asset intervention, operation, renewal and exit. A low initial platform price can produce high integration or lock-in cost.

The appraisal should separate recurring platform cost from intervention cost and retain a funded plan for model currency throughout the investment horizon.

36. OPTION VALUE

Digital evidence can support options to defer, stage, expand, switch or abandon. A flexible intervention may outperform a larger fixed project when demand or condition is uncertain. Option value requires defined decision points and evidence that will be available at each point.

Management should avoid treating flexibility as free. Monitoring, design interfaces, procurement rights and reserved land can carry cost.

37. PORTFOLIO OPTIMISATION

Capital constraints require comparison across projects. A portfolio model can combine risk reduction, service outcomes, statutory need, financial return, deliverability and strategic dependency. Scores should retain their underlying measures and avoid false commensurability.

The twin can estimate network interactions between projects. The optimiser should show dominated options, binding constraints and sensitivity to weights rather than provide an unexplained ranking.

38. HYPOTHETICAL OPTION SET

Consider a pressure zone with rising bursts, uncertain demand growth and limited redundancy. Management compares targeted renewal, pressure optimisation, a reinforcement main and a staged hybrid. The figures below are hypothetical assumptions used solely to demonstrate the method.

Table 2. Hypothetical option comparison for a pressure zone

Option	Initial capital	Evidence confidence	Principal benefit	Material downside	Staging feature
Targeted renewal	18	Medium	Reduces failure exposure	Bottleneck remains	Add batches after inspections
Pressure optimisation	7	Medium-high	Reduces leakage and bursts	Service margin can narrow	Pilot one district first
Reinforcement main	52	Medium-low	Adds capacity and redundancy	Demand can arrive late	Limited after construction
Staged hybrid	27	Medium	Near-term risk reduction and option	Requires active governance	Trigger reinforcement later

Monetary units are deliberately unspecified. The table is not a valuation or project recommendation.

39. CAPITAL TIMING

Timing can dominate value. Early investment can reduce risk while creating underutilised capital. Delay can preserve cash while exposing service and construction inflation. The model should identify trigger dates and consequences under each demand, condition and delivery scenario.

An approval can authorise design or enabling work before full construction. This separates readiness from irreversible commitment.

40. DISCOUNTED CASH FLOW

Discounted cash flow should include attributable cash effects, capex, operating cost, tax where relevant, residual value and working capital. Public-service outcomes can be presented alongside rather than forced into weak cash proxies.

The discount rate should match the risk and cash-flow basis. Model uncertainty is better represented through scenarios, ranges and decision staging than through an unexplained premium alone.

41. RISK-ADJUSTED VALUE

Expected value can combine consequence and probability where both have evidence. Severe low-frequency outcomes may also require threshold or resilience treatment beyond a simple average. Regulatory and statutory obligations can constrain the option set.

The investment paper should show the decision under central, downside and severe plausible conditions. It should identify which assumptions can reverse the ranking.

42. CONFIDENCE ADJUSTMENT

Confidence should affect action, evidence spending and staging. It should not become an arbitrary percentage haircut. A low-confidence, high-consequence decision may justify field investigation; a low-cost reversible action may proceed with monitoring.




Figure 2. Evidence confidence, decision consequence and the validation envelope
Illustrative governance map; thresholds should be approved for the specific use and utility.

43. VALUE OF INFORMATION

The value of information is the expected improvement in decision quality from an additional test, minus its cost and delay. It is greatest when uncertainty can change the choice and the decision has material consequences.

The twin can help target field tests by identifying sensitive locations and parameters. Testing should stop when additional evidence is unlikely to change the decision enough to justify the cost.

44. SCENARIO DESIGN

Scenarios should combine coherent assumptions for demand, climate, tariffs, energy, asset deterioration, delivery and regulation. Mixing optimistic inputs from unrelated futures can create an implausible case. Scenarios should remain few enough to understand and broad enough to reveal strategy failure.

The investment committee should see the option ranking under each scenario and the signposts that indicate which future is emerging.

45. STRESS TESTING

Stress tests examine conditions beyond the normal forecast. Examples include sensor outage, communications loss, pump failure, drought restrictions, major bursts, power interruption, cyber incident or delayed growth. The objective is to understand failure and recovery, not to claim a precise likelihood without evidence.

Results should inform contingency, redundancy, insurance, reserves and operational procedures.

46. BENEFITS REALISATION

Approval should create a benefits register with baseline, owner, measure, timing and verification method. Post-investment data should test whether the intervention delivered the physical and financial mechanism expected.

Realisation evidence should update deterioration assumptions, response curves and future option comparisons. A twin becomes a learning system when completed projects improve later decisions.

47. MODEL DRIFT

Network configuration, demands, controls, assets and data systems change. A once-validated model can drift outside its evidence base. Drift monitoring should compare predicted and observed behaviour and detect changes in inputs, residuals and decision-sensitive parameters.

Material drift should trigger investigation, recalibration, restricted use or withdrawal of a decision approval. The model register should record status and review date.

48. CHANGE CONTROL

Changes to topology, parameters, algorithms, interfaces, economic assumptions and thresholds require classification and approval. The control should distinguish routine data refresh from a change that can alter a capital recommendation.

Each released model should be reproducible from versioned inputs and configuration. Emergency corrections should be documented and reviewed after the event.

49. MODEL INVENTORY

The utility may have several hydraulic, demand, deterioration, optimisation and financial models. The inventory should name owner, purpose, version, dependencies, validation envelope, limitations, users and criticality.

Duplicate models can create conflicting answers. Consolidation is useful where purposes overlap; specialised models can remain where their differences are deliberate and governed.

50. ACCOUNTABILITY

Asset management owns the investment decision; operations owns safe operating knowledge; engineering owns technical design; data and technology own platform controls; finance owns economic integrity; risk and assurance challenge material assumptions. Clear accountability prevents the platform supplier or model team from becoming the de facto capital approver.

The committee minutes should retain the model's role, human judgement and accepted residual risks.

51. INDEPENDENT CHALLENGE

High-consequence decisions may need independent model review, engineering challenge, economic assurance or regulatory scrutiny. Independence should be proportionate and free from incentives tied solely to project approval.

The reviewer should examine purpose, data lineage, calibration, validation, uncertainty, alternatives and economic translation. Re-running the same model is only one part of challenge.

52. DATA RIGHTS

Contracts should establish ownership, access, portability, retention, permitted use and exit. The utility should be able to retrieve the data, configuration and decision record needed to operate, assure and migrate the capability.

Restrictions on training, derivative models or cross-client use should be explicit. Sensitive infrastructure and customer data require controlled handling throughout the supplier chain.

53. INTEROPERABILITY

Interoperability supports competition, reuse and long-lived asset decisions. Open formats and documented interfaces can reduce dependence on a single vendor, subject to security and performance needs. Semantic alignment matters as much as technical connectivity.

ISO 55013 links asset-data management to strategic, operational, financial and compliance functions [4]. The twin should serve that broader governance environment.

54. VENDOR LOCK-IN

Lock-in can arise from proprietary models, undocumented transformation, non-portable history, specialist skills or commercial terms. Exit cost should be assessed before procurement. A low platform fee can be offset by expensive migration or lost evidence.

The contract should provide export, documentation, transition support, escrow or continuity measures where proportionate, and a test of recoverability before the service becomes critical.

55. CYBERSECURITY

A connected twin can enlarge the attack surface and expose operational information. Controls should cover identity, segmentation, logging, supplier access, secure development, vulnerability management, backup and incident response. Automated control requires stronger separation and safe fallback.

Cyber scenarios should enter both operational assurance and capital economics. A benefit case that depends on continuous connectivity should price the controls and resilience needed to sustain it.

56. PRIVACY AND ETHICS

Customer-level consumption can reveal occupancy and behaviour. Data minimisation, purpose limitation, access control, retention and lawful processing should be designed into the system. Workforce analytics can also create legitimate concerns.

The investment case should include compliance cost and the operational consequences of restricted data use. Governance should prevent technical convenience from expanding the approved purpose.

57. PROCUREMENT STRATEGY

Procurement can separate data foundation, modelling, integration, assurance and managed service, or combine them with clear interfaces. Evaluation should test a representative decision rather than rely on demonstrations using curated data.

Bidders should disclose assumptions, limitations, recurring cost, dependency, model maintenance and evidence portability. Commercial incentives should reward accepted outcomes and sustainable capability.

58. BUILD, BUY OR PARTNER

The appropriate boundary depends on utility capability, market maturity, criticality, differentiation and lifecycle economics. Core decision knowledge and accountability should remain with the utility even when technology is supplied externally.

A partner can accelerate delivery and transfer skills. The operating model should define who diagnoses, validates, approves, updates and responds when the model disagrees with field judgement.

59. MINIMUM VIABLE TWIN

A minimum viable twin should solve one valuable decision with bounded geography and measurable evidence. It needs enough integration and governance to test the complete decision chain. A broad visual platform without a closed decision loop is a weak pilot.

The first use case should be material, observable, repeatable and safe to stage. It should create reusable data and controls for later applications.

60. PILOT DESIGN

The pilot should define baseline, decision, alternatives, validation data, operational owner, economic measure, acceptance thresholds and stop criteria. Parallel operation can compare the recommendation with current practice before authority expands.

The pilot should include adverse conditions and data failures. Success should require a better documented decision, not merely a functioning interface.

61. TECHNICAL ACCEPTANCE

Technical acceptance covers data ingestion, identity, topology, computation, calibration, validation, performance, security, resilience, documentation and reproducibility. Tests should reflect operational scale and poor-quality conditions.

Open defects need severity, owner and permitted-use restriction. Acceptance for one use should not be represented as acceptance for all future uses.

62. DECISION ACCEPTANCE

Decision acceptance tests whether the twin changed a real choice with traceable value. Management should review whether alternatives were credible, uncertainty visible, economics consistent and field response practical.

A technically accurate model can fail decision acceptance if its output arrives too late, cannot be explained or does not connect to accountable workflows.

63. STAGE GATES

The programme should pass gates for purpose, data readiness, model credibility, operational integration, economic value, controlled deployment and benefits realisation. Each gate needs evidence, authority and permitted next activity.

Capital for the platform and capital for network interventions should be distinguishable. A platform gate should not automatically approve the projects it recommends.

64. FUNDING THE CAPABILITY

The digital capability can be funded as utility operating expenditure, capital programme support, innovation funding, regulatory allowance or a shared service, subject to local accounting and regulatory rules. The funding basis should match the benefits and lifecycle.

Ofwat's innovation programme and PR24 investment context illustrate the scale of current experimentation and infrastructure need in England and Wales [5]-[6]. Local approval remains specific to the regulatory regime.

65. PORTFOLIO BUSINESS CASE

The platform business case should aggregate only compatible benefits and avoid counting a network intervention benefit twice. It should include the cost of data remediation, integration, change, assurance and ongoing model stewardship.

The committee should see a no-platform alternative, targeted-tool alternatives and the cost of maintaining current practice. A platform can be justified through multiple use cases once each is evidenced.

66. RISK ALLOCATION

Risk should sit with the party able to manage it, subject to affordability and system accountability. The following matrix is illustrative.

Table 3. Illustrative allocation of digital-twin decision risks

Risk	Natural owner	Contract or governance response	Capital-decision consequence
Asset-record error	Utility with field support	Data remediation and confidence flags	Restrict affected recommendations
Software defect	Supplier and utility assurance	Testing, warranty and incident process	Reproduce and reapprove material outputs
Demand forecast	Utility leadership	Scenarios and staged triggers	Preserve option to defer or expand
Sensor failure	Operations and supplier	Redundancy, maintenance and fallback	Reduce confidence or suspend use
Model drift	Model owner	Monitoring and scheduled validation	Withdraw use outside tolerance
Cyber incident	Shared	Segmentation, response and continuity	Safe mode and evidence review
Benefit shortfall	Benefit owner	Baseline, verification and action plan	Update portfolio and future assumptions

67. GOVERNANCE CADENCE

Operational review can monitor data and residuals frequently. Model owners can review performance monthly. Capital committees can review material recommendations at stage gates. Board oversight should focus on aggregate exposure, major decisions and control effectiveness.

Cadence should follow decision velocity. A static annual review is inadequate for a rapidly changing operational model; continuous meetings without clear thresholds can create noise.

68. MANAGEMENT INFORMATION

Useful reporting shows model status, validation envelope, material exceptions, decisions supported, realised benefits, drift, incidents, changes and upcoming approvals. It should separate observed measures from simulated outcomes and management assumptions.

Dashboards should link to evidence rather than replace it. Senior users need concise decision information and access to deeper technical records when challenged.

69. REGULATORY EVIDENCE

Regulators may examine need, efficiency, deliverability, customer outcomes and risk. A twin can strengthen evidence when its method, data and uncertainty are transparent. Proprietary opacity can weaken confidence even when the visual output is persuasive.

The utility should preserve reproducible cases, versioned assumptions and explanations suitable for independent review. Regulatory acceptance should not be assumed from technical adoption.

70. ORGANISATIONAL CAPABILITY

Digital water requires hydraulic, asset, operational, data, cyber, finance and commercial capability. Training should focus on decision roles, limitations and escalation as well as software use. Experienced operators hold contextual knowledge that should be captured and challenged constructively.

Succession and supplier transition plans matter because the twin is a continuing operating capability rather than a one-off model build.

71. ADOPTION

Adoption grows when the system addresses a real pain point, preserves accountability and makes evidence easier to test. Users should see how recommendations are formed and how to challenge them. Early wins should be documented with limitations.

Management should track use in actual decisions, not logins alone. A widely viewed dashboard can remain commercially irrelevant if approval papers ignore it.

72. FAILURE MODES

Common failure modes include unclear purpose, weak asset identity, calibration without validation, hidden manual transformations, overconfident precision, missing economic counterfactual, automated recommendations without authority, vendor dependence and unfunded model maintenance.

The control response is explicit scope, evidence lineage, independent tests, versioning, economic reconciliation, decision rights, exit planning and lifecycle budget.

73. RED FLAGS

Red flags include a claim that the twin is universally accurate; absence of validation data; unexplained parameter changes; reliance on one central forecast; benefits valued at gross tariff without cash analysis; no field confirmation; results that cannot be reproduced; and project approval before the alternative set is defined.

Any red flag material to the decision should pause, restrict or condition approval.

74. FIRST 180 DAYS

The first month should select the decision and name owners. The next two months should reconcile critical data, define the counterfactual and design validation. Months four and five should run field tests, option economics and user trials. Month six should complete independent challenge and a controlled capital decision.

The output is an approved or rejected investment with an auditable evidence chain, plus a reusable capability baseline.




Figure 3. A 180-day route from decision definition to controlled capital use
Generic implementation sequence; timing depends on utility readiness, procurement and decision consequence.

75. INVESTMENT-COMMITTEE PACK

The pack should include decision statement, network boundary, service objective, alternatives, counterfactual, evidence register, topology status, calibration, independent validation, uncertainty, scenario results, whole-life economics, benefits ledger, risk allocation, delivery plan, cyber review, data rights, assurance and stage gates.

All management estimates should be identified. The pack should state where the twin changed the analysis and where professional judgement remains decisive.

76. DECISION RULE

Approve an intervention when the physical mechanism is credible, the model is validated for the intended use, uncertainty is visible, alternatives are fairly compared, downside performance is acceptable, delivery is feasible and accountability is clear. Condition approval where targeted evidence or staged triggers can protect value.

Do not use a digital twin to conceal a weak counterfactual or speculative demand case.

77. SYSTEM-LEVEL LEARNING

Completed interventions create evidence about failure rates, demand response, leakage, energy, delivery cost and customer impact. That evidence should update the twin and the capital framework. The value compounds when learning transfers across zones and programmes with appropriate context.

Learning requires stable definitions and willingness to record disappointing outcomes. Benefits governance should reward accuracy and correction rather than optimistic persistence.

78. STRATEGIC VALUE

A credible twin can improve network understanding, investment sequencing, regulatory evidence and cross-functional dialogue. It can reveal options that isolated asset plans miss and support faster response when conditions change.

Strategic value depends on trust earned through use-specific evidence. The platform should remain subordinate to the utility's public-service obligations and governance.

79. BOARD QUESTIONS

The board should ask which decisions the twin is authorised to support, where it has been validated, what uncertainty could reverse current priorities, who owns the model, how value is measured, how cyber and data risks are controlled, and how the utility can operate or migrate if a supplier fails.

These questions connect digital investment with fiduciary, service and resilience accountability.

80. CONCLUSION

A water-network digital twin becomes a working model when it is tied to a real decision, a defined validation envelope, transparent uncertainty and an accountable approval process. Its purpose is to improve the quality and timing of capital choices across renewal, capacity, leakage, energy and resilience.

The strongest programmes connect physical evidence, model credibility, economic alternatives, stage gates and post-investment learning. They use simulation to expose choices and consequences while preserving human responsibility for public infrastructure capital.

REFERENCES

[1] UK Government. Digital twin: official definition. 29 October 2025. <https://www.gov.uk/government/publications/digital-twin-definition/digital-twin-official>

- [2] National Institute of Standards and Technology. Credibility Consideration for Digital Twins in Manufacturing; Digital Twins programme resources. <https://www.nist.gov/publications/credibility-consideration-digital-twins-manufacturing> and <https://www.nist.gov/digital-twins>
- [3] United States Environmental Protection Agency. Water Distribution System Analysis: Field Studies, Modeling and Management; hydraulic-model calibration guidance. <https://nepis.epa.gov/Exe/ZyPURL.cgi?Dockey=2000D2C2.TXT>
- [4] ISO Technical Committee 251. ISO 55013:2024, Asset management: Guidance on the management of data assets. <https://committee.iso.org/sites/tc251/home/news/content-left-area/news-and-updates/iso-launches-new-standards-in-th.html>
- [5] Ofwat. Water Innovation Fund Annual Report 2025. <https://www.ofwat.gov.uk/wp-content/uploads/2025/07/Ofwat-Innovation-Fund-Annual-report-2025-v5.pdf>
- [6] Ofwat. Ofwat approves £104 billion upgrade to accelerate delivery of cleaner rivers and seas and secure long-term drinking-water supplies. 19 December 2024. <https://www.ofwat.gov.uk/ofwat-approves-104bn-upgrade-to-accelerate-delivery-of-cleaner-rivers-and-seas-and-secure-long-term-drinking-water-supplies-for-customers/>
- [7] Ciaponi, C. et al. A digital water strategy based on the digital water service concept to support asset management in a real system. *Journal of Hydroinformatics*, 2023, 25(5), 2004-2022. <https://iwaponline.com/jh/article/25/5/2004/96957/A-digital-water-strategy-based-on-the-digital>
- [8] Japan International Cooperation Agency and Boston Consulting Group. Digital Transformation in the Water Sector. 2025. https://www.jica.go.jp/about/dx/publication/2025/_icsFiles/afieldfile/2026/01/05/20250416-jica-bcg-water-dx-report.pdf
- [9] California Energy Commission. Advancing Demand Response in the Water Sector. 2024. <https://www.energy.ca.gov/publications/2024/advancing-demand-response-water-sector>
- [10] American Water Works Association. Digital Twins: The New Frontier for Water Utilities. *Journal AWWA*. <https://awwa.onlinelibrary.wiley.com/doi/pdf/10.1002/awwa.1979>
- [11] National Institute of Standards and Technology. Security and Trust Considerations for Digital Twin Technology. 2025. <https://www.nist.gov/news-events/news/2025/02/security-and-trust-considerations-digital-twin-technology-nist-releases-ir>
- [12] United States Environmental Protection Agency. Water Audits and Water Loss Control for Public Water Systems. <https://nepis.epa.gov/Exe/ZyPURL.cgi?Dockey=P1009VCZ.TXT>
- [13] United States Environmental Protection Agency. Formal Bayesian Approach for Real-Time Calibration of Water Distribution System Models. https://cfpub.epa.gov/si/si_public_record_report.cfm?Lab=NRMRL&datebeginpublishedpresented=01%2F01%2F2010&dirEntryId=350607&showcriteria=0&sortby=revisionDate&subject=Water+Research
- [14] ISO Technical Committee 251. ISO 55001:2024, Asset management system requirements. <https://committee.iso.org/sites/tc251/home/projects/published/iso-55001.html>

ABOUT THE AUTHOR

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.

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

Profile: <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.