



RESEARCH COMPANION

Editorial date: 25 September 2026

A Model of Distributed Authority During Disruption: The Limits of Safe Local Decision-Making in Digital State Services

Delegated Authority Under Network Partition: An Exhaustive Toy-State Study and a Research Protocol for Safe Offline Public Services

MERIDIAN FUTURES LAB | M02 Scientific Companion | Working Paper 1.0 | Editorial date: 25 September 2026 | Evidence cutoff: 25 September 2026

Accurate classification: A critical conceptual paper containing a fully executed, bounded computational state-space examination (Conceptual Working Paper with an Executed Bounded Exhaustive State-Space Exploration). **It is not presented as an original Saudi operational study, a complete security proof, or a peer-reviewed paper.**

Proposed program period: 2026–2035.

Arabic Abstract — English Translation

This paper examines a question distinct from service availability during disruption: what **limit of local authority** can a digital system exercise when it cannot verify shared records in real time, while preserving rights invariants and the safety of beneficiaries? It distinguishes

viewing a signed copy, receiving a provisional request, completing a scarce or unique entitlement, and obtaining a final enforceable decision. We connect this distinction to Consistency as Logical Monotonicity (CALM), Coordination Avoidance, the limits of Conflict-Free Replicated Data Types (CRDTs), and the offline-payment design challenges explored in Project Polaris.

We use three bodies of evidence: the Iberian power outage of 28 April 2025 and the accounts by the Banco de España and European Central Bank (ECB) of disrupted payment channels, the advantages of cash, and the limits of its distribution; the continuation of some British Library functions while others remained unavailable after the October 2023 attack; and published Saudi texts on business continuity management in banking and digital government. None of this evidence permits an estimate of the probability that a Saudi service will fail or reveals its actual local authority. It does, however, define questions that can be tested.

We executed a complete computational exploration of a **bounded synthetic model** comprising two disconnected sites, three units of shared authorization, five possible levels of local demand at each site (zero through four), and five decision policies: **25 demand states and 125 policy evaluations**. Under blind local commitment, the shared limit was violated in **13 of the 25 states**. An allocation split into two units and one preserved the limit in every state under the assumption of honest local enforcement, but across all states it **forwent ten opportunities for safe completion** relative to a connected benchmark. This is an **exhaustive count of constructed abstract states**, not a statistical probability or a result from a real facility. A separate example shows that a valid cryptographic signature alone does not prevent a copy of the same voucher from being used at two sites without coordination or corresponding hardware guarantees.

The paper proposes an Offline Authority Profile (OAP) for specifying each service's authority during disruption, and an authorized testing protocol that balances safety of rights, actual completion, backlog, social equity, and the cost of review. Its contribution is to organize the questions, conditions, and tests while identifying where stopping or making a provisional promise is more honest and safer than an irreversible final decision.

Keywords: Bounded Local Authority; Network Partition; Application Invariants; Coordination Avoidance; Escrow Allocation; Double Spending; Signature Validity; Reconciliation.

English abstract

The research question is **which public-service decisions can be safely authorized locally during a network partition**, rather than simply which interfaces remain operational. A distinction is made among viewing signed and time-bounded information, receiving provisional intent, consuming preallocated scarce authorization, and making final decisions that require shared, fresh knowledge. The synthesis draws on CALM, invariant-confluent

coordination avoidance, conflict-free replication, and the Bank for International Settlements' Project Polaris research on offline payment design. Documented case evidence covers the April 2025 Iberian blackout and payment disruption, the British Library's post-2023 partial-service restoration, and Saudi rules on institutional business continuity.

An **executed, fully synthetic, bounded exhaustive analysis** explores 25 demand states at two disconnected sites under five decision policies. The global scarce authorization budget equals three invented units. Blindly letting each site independently commit against the old global balance violates the shared quota in 13 of 25 cases. Honest, disjoint escrow allocation preserves the quota in all enumerated cases but yields 10 fewer safely finalized units in aggregate than an always-connected benchmark. These are **counts of constructed abstract states, not estimated risks, probabilities, Saudi results, or guarantees against adversarial compromise**. A separate replay counterexample illustrates why a valid issuer signature alone does not ensure unique offline redemption. We propose an Offline Authority Profile and an authorized measurement protocol with explicit outcomes, safety invariants, revocation constraints, appeal handling, and distributional assessment.

Keywords: offline public services; delegated authority; safety invariants; network partition; application semantics; bounded escrow; digital identity; payment resilience; Saudi Arabia.

1. Unit of Analysis: The Act That Changes a Right

Many discussions of business continuity become lists of infrastructure: a second server, an alternative region, backup connectivity, a reduced-function mode. These are important measures, but they do not directly answer whether an act's effect remains **valid** when the shared information it requires is absent. A local terminal may continue receiving requests, but receiving a request is not the same as finally disposing of a legal or financial right. We therefore choose a different unit of analysis: the **Authorized State-Changing Action**, or a read-only action that must disclose how current its evidence is.

The primary question is: **under what legal, informational, cryptographic, and operational conditions can a branch, site, device, or service representative complete a specified act during a disruption of unknown duration from a shared authority, without exceeding delegated authority or creating a conflict that cannot be repaired?**

We define five subsidiary questions:

1. Which decisions remain valid when additional information arrives later, and which depend on knowing that no right or restriction exists at a distant site?
2. What is the difference between cryptographic signature validity and authorization to use a resource or delegation only once?

3. How can a scarce right be divided among independent sites without exposing global-register invariants to violation, and what is the cost of unused allocations?
4. When does recording provisional intent create social value without turning a receipt into a false final promise?
5. How should local capability be tested against service law, real conditions of human access, and unequal distributions of devices, power, and time?

The paper's purpose is conceptual and design-oriented, with a completed illustrative computational component. This version does not test national causal relationships, infer rates of failure or exposure, or validate the security of a real payment system. Functional, legal, security, and social evaluations remain distinct. Preserving a numerical limit in a toy state machine does not make that machine fit for institutional adoption.

Hypotheses and Counterindications

H1 — Conditional prior delegation: Some acts have invariants that can be preserved locally through bounded, disjoint delegation from an authorized body, supported by evidence, signatures, audit, and frank disclosure to the user about the request's status.

H2 — Limits of knowledge during disruption: Other acts require a current register or a unique resource. If every site can use the same stale copy, excess or conflicting rights may be created.

H3 — The cost of limits: Prior distribution can guarantee exclusivity and reduce harm, but it may freeze allocations at sites that do not need them during the crisis, preventing other sites from performing otherwise valid work.

H4 — Value of a provisional promise: Documented provisional intake may preserve the submission time and the applicant's right to review, but it is not final execution and is not suitable for every kind of right.

H5 — A plausible alternative: For most sensitive functions, the better design may be a safe stop, limited protected connectivity, or a stronger shift to in-person service rather than complex local delegation. These alternatives must remain in the comparison.

The practical-value thesis can be rejected if authorized tests show that a specified local delegation does not increase valid completion relative to a temporary stop once attacks, corrected rights, and costs are counted, or if it creates inequity in access that public policy cannot accept.

2. Knowledge Foundations: From Connectivity to the Meaning of an Act

2.1 Why Is Data Replication Insufficient?

Distributed-systems research describes environments in which independent devices operate while messages between them may be delayed or lost. “Correct” does not have a single meaning in such an environment. It may mean that replicas eventually show the same value after reconnection. It may instead mean that every decision taken by a user or authority during the interval complied with a rule that should never be violated. The first goal may be achieved after the second has already failed.

Marc Shapiro and collaborators showed how replicated data types can be designed to converge after updates under specified conditions in their work on Conflict-Free Replicated Data Types (CRDTs). These are tools for guaranteeing particular merge properties. They are not, by themselves, enough to define the law governing a scarce right or the obligation not to disburse it twice. [S10](#)

2.2 What Can Be Done Without Coordination?

Consistency as Logical Monotonicity (CALM) relates certain logical classes of operation to their capacity to execute without continuous coordination while preserving consistent semantics. Joseph M. Hellerstein and Peter Alvaro present this distinction and examine its implications for distributed-software design. [S08](#)

Monotonicity is used here in its technical sense: in an appropriate logical formulation, adding new inputs to a set of inputs should not cause the program to retract previously established results in a way that violates its condition. Recording an independent event or adding a qualifying document is therefore simpler than deciding something that depends on proving an event did not occur elsewhere. Not every “addition” is safe, however; removing an effect, revoking an authorization, and balancing a resource can change the judgment.

In *Coordination Avoidance in Database Systems*, Peter Bailis and collaborators proposed a criterion tied to application invariants called **Invariant Confluence (I-Confluence)**. The practical lesson here is that replica convergence is insufficient. We must determine whether merging two locally valid states could violate a global validity condition. The paper explicitly distinguishes availability and convergence properties from guarantees that application invariants remain safe. [S09](#)

A digital state service may have an invariant that an authorization cannot be used after revocation, an entitlement share cannot be disbursed twice, or a unique record cannot be assigned to more than one party. Such invariants require legal definition before they can be represented as software conditions. A researcher cannot invent a disbursement rule and then treat every institution as bound by it.

2.3 Offline Payment Is Not a Ready-Made Operating Mode

In 2023 the Bank for International Settlements (BIS), through Project Polaris, published a handbook on offline central-bank digital-currency payments. The report treats privacy, inclusive use, device security, holding and spending limits, operating policies, and replay risks as design choices without a single solution suitable for every context. A later document from October 2023 noted that many solutions were then in experimentation or development and that few were ready for general operation. [S06S07](#)

The European Central Bank (ECB) announced consultations in August 2026 on standards for Secure Elements and embedded SIMs in planning for an offline digital-euro capability. At the present evidence cutoff, the digital euro itself **is not a currency available to the public**. A possible launch in later years remains conditional on legislation, decision, and testing. [S16S17](#)

These sources narrow and sharpen the idea. Authority moved to a device requires real management of scarcity, prevention of duplication, and verification of the authorized user. Some properties still cannot be proven locally after an update at an absent center, including some forms of revocation and simultaneous transfer across independent devices.

3. Evidence from the Real World: Three Different Units of Observation

3.1 Spain and Portugal on 28 April 2025

In an account published on 27 April 2026, the Banco de España recorded that the widespread power outage began at about 12:33 p.m. on 28 April 2025, disrupting electronic payments and ATMs while cash already in people's hands remained a direct means of exchange requiring neither power nor real-time connectivity. Some ATMs did not work despite containing cash, and the bank and cash carriers replenished machines on 29 and 30 April. [S01](#)

The European Central Bank (ECB) reported in its 2025 study that in-person card payments in affected areas fell by about 41–42% under the stated comparisons and online purchases by about 54%. CaixaBank Research data, based on a sample of transactions and comparative studies, estimated that consumer spending captured by its measures fell by about 34% on the

day of the outage, with some of the difference recovered over subsequent days. These percentages do not constitute a randomized experiment or measure all economic activity; regional differences, deferred purchases, and the scope of the institution's data affect their interpretation. S02S04

On 20 March 2026, the European Network of Transmission System Operators for Electricity (ENTSO-E) released the conclusions of its investigation into the Iberian outage, focusing on multiple interacting factors in the power system. We do not reinterpret the cause of the outage through our theory of digital authority. We use the event as a **documented external disruption** that exposed payment situations in which an absent network ordinarily had to give permission. S03

3.2 The British Library, 2023–2025

After the October 2023 cyberattack on the British Library, its buildings and reading facilities remained available while research, catalogue, and online-service capabilities were disrupted. The institution published its own review in March 2024 and progressive service-restoration updates in 2025. The available materials do not establish that any special “local authority” model was in use. They show only **that continuity in one part of an institution's mission can coexist with substantial, publicly acknowledged incapacity in another**. S14S15

3.3 Saudi Arabia as the Site of a Question, Not an Assumed Failure

The Saudi Central Bank (SAMA) regulatory site includes a Business Continuity Management Framework published in 2017 and 2021 instructions on alternative power arrangements for banks, including generator testing, capacity, and maintenance. The Digital Government Authority (DGA) also publishes regulatory materials and reference guidance for risk and continuity management in digital government, including version 2.0 of its controls and a later guide. S11 (related 2025 guide)S12S13

This research program has no right to access internal failure or authorization logs. We do not infer that a sensitive Saudi application lacks local delegation or possesses any form described here. The existence of continuity regulation demonstrates institutional attention to the subject, but it does not establish which acts can legally and operationally be finalized without connectivity or the results of response exercises.

Table of Inference Limits

Evidence group	Observation established by the source	Function in this research	Claim it does not authorize
Iberia 2025	Widespread power loss and disruption of some payment channels; cash held by users remained exchangeable	Illustrates the independence of some forms of settlement from real-time authorization networks	That cash served every user or all digital payments failed equally
British Library 2023–2025	Study spaces continued while research and catalogue services suffered prolonged disruption	Shows unequal levels of capability and operation that an institution can provide	That partial operation removed the attack's consequences or a local-delegation document was in use
Saudi texts	Formal frameworks for management, testing, and alternatives in the specified sectors	Defines an institutional basis on which research collaboration could be requested	That offline execution units have undergone published trials or proven successful

4. Conceptual Model: A Decision Is Information, Authority, and Evidence

Let a service action be a , the party wishing to complete it i , and the time t . A local decision requires at least four forms of validation:

1. **Authorization:** The authority of the body, device, or employee to perform a specified act within a recognized mandate.
2. **Evidence Integrity:** The source, signature, and integrity of the copy or request, and the validity of verification keys under the relevant policy.
3. **Freshness:** The issuance time of the data, the risk that it changed or was revoked during disruption, and the tolerable duration of that uncertainty.
4. **Invariant Safety:** Assurance that merging results from other sites will not cause the act to create a duplicate right, debt, or prohibited change.

In a proposed design framework, permission for a local act can be represented by this logical condition:

$$\text{ALLOW}(a, i, t) = A(a, i, t) \wedge E(a, i, t) \wedge F(a, t) \wedge I(a, i, t).$$

A zero value for any condition means that **finalization** should stop under this rule; provisional intake, explanation, or referral may remain possible through another path. This is a conceptual formulation, not a complete legal decision engine. Each function must be defined and tied to regulatory references, tests, and data published to the necessary extent.

We do not equate a signature with the validity of every element of evidence. Under its cryptographic model, a signature can establish that a message was not changed and was issued by a specified key. It cannot alone establish that the holder still owns the right, that the key was not revoked while connectivity was absent, or that the message was not used elsewhere. This is not a flaw that must appear in every design. Shorter-lived delegation, distributed allocations, protected hardware, value and use limits, and audit logs may address it, but each remedy introduces cost and new trust conditions.

4.1 A Necessary Distinction Between Decision Safety and Continuity of Conversation

We distinguish four outcome types:

- **Completed Valid:** A final, valid act within a declared and auditable delegation.
- **Accepted Provisional:** A request received with a receipt stating that review or reconciliation remains pending.
- **Deferred or Escalated:** A request suspended with a reason, a clear status, and a review time or assistance channel.
- **Rejected or Unsafe:** A correct refusal because authority or evidence is absent, or an unsuccessful attempt whose approval would have violated a condition.

These four types must not be collapsed into a count of “transactions completed.” Fairness also requires that provisional intake not always be treated as failure. It may preserve the claimant’s place and submission time where policy gives the receipt a real, though non-final, effect. That question remains subject to the legal context and the nature of harm caused by delay.

5. Preserving a Scarce Right: From Impossible Coordination to Conditional Division

We begin with a deliberately simple, auditable model. A limited total quantity B of a right exists. Before the disruption, an authorized central body gives independent sites disjoint allocations q_i . Each site may finalize uses c_i only within its allocation. If:

$$q_i \geq 0, \quad c_i \leq q_i, \quad \sum_i q_i \leq B,$$

then:

$$\sum_i c_i \leq \sum_i q_i \leq B.$$

This argument is valid **inside the model**, conditional on allocations neither being forged nor duplicated, local state being preserved honestly, and every participant agreeing to the rule of use. It is not a security proof for a real system. It does not examine lost devices, cloned hardware, shared keys, replay attacks, revocation of a prior delegation, or human decisions that override policy. These conditions are not minor details: violating any one of them may break the bound itself.

The mechanism resembles escrow allocation: permission need not be requested for every movement because part of the right to act now has known limits and is exclusively reserved at one location. The benefit has a price: **allocations are distributed before actual demand during the disruption is known**. Site A cannot simply move an unused right from B, because that would require connectivity or a new protocol ensuring that the transfer was not repeated.

5.1 Why Is a Stale Balance at Two Sites Insufficient?

If, before disconnection, both sites retain the same copy saying “three units remain” and each allows itself to spend three, the result is six. Both sites act honestly against their local copy, yet together they violate the total. This is a **counterexample** enabled by absent coordination and an undivided copy of a scarce resource. It must be kept separate from real money: financial systems may have other rules, settlement mechanisms, hardware, and safeguards that this model does not simulate.

5.2 Why Does a Valid Signature Not Imply Unique Redemption?

We executed a separate example with one copyable digital token, τ_1 , bearing a valid issuer signature and known to two disconnected points. If neither point can see the other’s usage log and the token is not bound to unclonable hardware or an exclusive place of use, each point may accept it once. The signature is valid twice, but the invariant that the token is used **only once globally** fails. `data/duplicate_token_counterexample.json` records this state explicitly.

The claim is not that every signature or secure digital wallet has the same vulnerability. The model shows why a claim that “a signature prevents offline double spending” must specify its hardware design, balance binding, allocation, or duplication constraints. Project Polaris distinguishes solutions and risks according to their architecture and gives Secure Elements and operational limits a central place. [S06S07](#)

6. Completed Computational Experiment: 25 States $\times$ Five Policies

6.1 Scope of Exhaustion and Limits of the Claim

We completed a programmatic enumeration of a **predefined state set**, not a Monte Carlo simulation or a sample of institutions. In each state, the independent request count at sites A and B belongs to $\{0, 1, 2, 3, 4\}$, producing $5 \times 5 = 25$ pairs. The total authorization budget is $B=3$ synthetic units. Every pair is evaluated under five policies. The model contains no occurrence probabilities for those pairs, population weights, outage durations, production cryptography, or models of human behavior.

6.2 The Five Policies

P0 — Central Stop: All final approvals wait for central connectivity; nothing is executed during isolation. Another institution might still provide information or assistance, but that lies outside the numerical calculation.

P1 — Provisional Intake: The system records every intention as provisionally received. None counts as finally completed, and each awaits review after reconnection.

P2 — Escrow (1,2): One unit is reserved for A and two for B. Each site has a limit it cannot exceed offline; excess requests remain in a review queue.

P3 — Escrow (2,1): The exact reverse, testing the effect of giving one region twice the share of the other before demand is known.

P4 — Blind Local Commit: Each site executes up to three units as though it alone were consuming the total resource. The policy is used to expose a violation and is not presented as a professionally acceptable option.


The connected reference benchmark can finalize $\min(3, d_A + d_B)$ units in each state after observing both demand values. The model does not calculate waiting time, operating cost, or verification time and therefore cannot support a financial comparison or describe the duration of a real outage.

6.3 Reproducible Results

Policy	Test states	States exceeding budget	Largest excess in model units	Total finalizations across 25 states	Total provisional intents	Safe-finalization shortfall versus connected reference
P0 Central Stop	25	0	0	0	0	65
P1 Provisional Intake	25	0	0	0	100	65
P2 Escrow (1,2)	25	0	0	55	45	10
P3 Escrow (2,1)	25	0	0	55	45	10
P4 Blind Local Commit	25	13	3	90	0	0*

For P4, zero in the shortfall column does not mean safe operation. The policy exceeds the total right in 13 states, and some approvals are invalid. All totals aggregate across **mutually exclusive alternative states**; they are not transactions or people observed together in time. The value 13/25 must not be interpreted as a 52% probability of failure.

Twenty-five EXHAUSTIVE, SYNTHETIC request states



Two disconnected sites | global right = 3 units | no frequencies or probabilities inferred

Two maps of budget violations and forgone opportunities under split delegation

Figure 2 — Each axis shows the demand at one site, producing 25 cells. The first map shows the amount by which the right is exceeded; the second shows safe-finalization opportunities forgone under a fixed 1+2 split relative to complete connected knowledge. The figures come directly from the code and contain no Saudi data.

6.4 A Counterexample and a Benefit That Must Not Be Forgotten

If site A requests three units and B requests none, the (1,2) allocation completes one unit. A therefore loses two opportunities for safe completion that the center could have realized without exceeding the budget. This is **the price of planning before observing demand**. It may be reduced by reallocating quotas while connectivity is available before a crisis, using finer divisions, or establishing a limited channel for transferring authority where legal and technical safeguards allow. Scarcity cannot be solved by copying the balance.

In the reverse example, where each site requests two units, blind local commitment completes four despite a total right of three. Disjoint escrow preserves the limit, but may deny or defer one request. The correct choice depends on the cost of inequity and delay, declared priority rules, and the nature of the right—not merely on maximizing a “completed” column in a dashboard.

6.5 What Was Actually Verified, and What Was Not

The code ran and produced `model_check_states.csv` with 125 rows, `model_check_summary.csv` for five policies, an assumptions file, a log of the duplicated-voucher violation, and two original figures. The program includes computational assertions: no honestly enforced escrow design exceeds the budget in the defined states; the row count is 125; and the blind policy produces actual violations. **This is a test of a small software model**. It does not test devices, payment protocols, key custody, or comprehensive institutional assurance.

We did not conduct sensitivity tests by adding population segments or inventing a probability distribution. Applied research would require rerunning the work with authorized real demand allocations, differential priority costs, storage and battery limits, connectivity loss, key-revocation states, stolen wallets, and audit integrity.

7. Granting Authority Is Not the Same as Enabling a Device

7.1 Revocation and Time Limits

A signed document may be valid when issued and the right it represents may be revoked during disruption. Consider two possible system histories: in H_0 , no revocation occurs; in H_1 , revocation occurs at a center the local device cannot hear. If all inputs visible to the device are identical in both histories, a local algorithm alone cannot determine which history is true. The system then needs a declared policy specifying the acceptable age of evidence, an act that remains valid if state changes, a connectivity path, a more tightly time-bounded authority, or a stop on final execution.

Machine learning cannot overcome this ignorance merely by estimating the probability of revocation. A prediction does not make an unauthorized decision authorized. An artificial-intelligence (AI) model can help prioritize review cases or detect contradictions after reconnection, but it creates neither legal authority nor evidence of information it cannot observe.

7.2 A Capability Its Holder Cannot Use During the Crisis

A user may carry a digital document that can be displayed offline yet still need someone to accept it under law and institutional practice. A payment card may verify some data locally while the merchant's terminal cannot operate without electricity. The entire system must therefore be tested: issuer, device, employee, beneficiary evidence, complaint channel, local power, and disposition of the receipt.

7.3 Privacy and Discrimination

Offline capability may help users in places with weaker connectivity, but it may exclude people without a modern device, valid document, or preloaded balance. An emphasis on synchronizable audit logs may also move more personal data than necessary. The design needs data minimization, parallel human routes, and measurement of performance differences among appropriate regions and groups under privacy protection and authorized participation. A general average speed must not conceal that a user with limited means was left without a route.

7.4 Not Every Effect Can Be Repaired

A queue can be reordered, and two notes on a register can be merged. It may be impossible to undo delivery of a unique good to an unauthorized person, an unsafe medical decision, or a legal transfer made by someone without final authority. “It can be synchronized later” describes an informational capability; it does not always mean that social and legal conditions can be restored to their former state.

8. Design for an Authorized Saudi Field Study

8.1 Sample and Unit of Record

We propose voluntary, authorized partnerships with institutions whose service functions can be exercised in isolated drills or separate test environments without disrupting a real public service. The unit of observation would be **a class of action within an authorized test window**, not a server name or the address of a sensitive facility. Acts would be classified by legal effect: read, intake, reserve a resource, reversible finalization, or high-impact finalization that must remain central.

8.2 Variables to Collect

Variable	Required operational definition	Possible authorized source	Protection or interpretation constraint
Action type	Is it an inquiry, intake, issuance, disbursement, or revocation?	Service policy and specification	Researchers do not change actual categories of authority
Local delegation	The authorizing document and its temporal, qualitative, and numerical limits	Approved de-identified institutional documents	Key secrets and storage locations are not published
Evidence freshness	Time of the most recent validation of data and authority	Aggregated local test logs	Individual log paths are not published
Valid finalization	A final decision satisfying the policy selected before the experiment	Legal and operational review of outputs	Must not be equated with provisional intake
Provisional receipt	A preserved request with evidence of time, tamper resistance, and a follow-up context	Experimental or synthetic records	Acceptance of the right is not declared before reconciliation
Conflict	A difference between local decisions or between one decision and a later register	Synthetic events in a protected environment	Classifying real harm requires specialist approval
Review and remedy	Time to settle disagreement and method for notifying the beneficiary	Tests of institutional procedures	The assumption that every error is fixed automatically is rejected
Equity	Differences in the ability of appropriate groups to submit or complete a request	Aggregated measures with privacy protection	Sensitive traits are not inferred without necessity and legal basis
Operating cost	Costs of training, hardware, review, delay, and human assistance	Aggregated internal financial data	Only homogeneous costs may be compared

8.3 Comparison Design

For every homogeneous action class, the test environment would compare policies specified in advance: stop finalization, provisional intake, local delegated allocations where lawful and feasible, and perhaps an alternative connectivity route. Synthetic scenarios could be randomized inside the authorized test, with equivalent request types, allocations, and assumed disruption durations. We do not recommend deliberately exposing real users or real legal rights to a violation for experimental purposes.

Primary outcomes: the number of valid acts completed within legal limits; the number of detected and prevented attempts to violate an invariant; time to reconcile waiting receipts; the number of cases that cannot be resolved without human intervention; aggregate access differences; and total cost. Acceptable thresholds should be set **before** reviewing the comparison results, and every deviation from the analysis plan should be recorded.

An acceptable competing result: The study may find that limited backup connectivity, stronger human-service channels, or a policy of stopping certain acts provides greater benefit and safety than distributed delegation. If so, the evidence should support that recommendation. The research should not be forced to prove the value of its author's preferred model.

8.4 Adversarial and Safety Testing

The following conditions would be tested in an authorized defensive environment: stale evidence; a document revoked during disconnection; a replayed token; a lost device; an untrusted clock; different update-arrival orders; a provisionally accepted receipt that cannot later be honored; partial reconnection that does not cover every register; and repeated disruption during reconciliation. A public research report need not disclose key material, critical topology, or procedures that could help target specific facilities or services.

9. Offline Authority Profile (OAP): A Design Proposal

We propose a structured institutional record for each sensitive act, rather than a mere appendix to an emergency operating plan. It answers seven connected questions: **the right, the authority holder, the evidence, its time horizon, use limits, stop rules, and the means of reconciliation and appeal**. A schematic representation follows:

Proposed field	Minimum content	Failure criterion requiring attention
authority_source	Basis of delegation and the law or policy allowing the act	Undefined authority or action outside the institution's jurisdiction
action_class	Inquiry, intake, reservation, finalization, referral	Labeling a final act as a "provisional request," or the reverse
evidence_policy	Source, signature, data validity, update time	Validation of an expired signature or revoked evidence
scarcity_guard	Exclusive bound preventing excess use of a non-repeatable resource or right	Clonable tokens or allocations
stop_rule	Conditions of insufficient evidence, conflict, or inequity	Continuing finalization despite violation of a binding invariant
reconnect_policy	Event log, reconciliation, audit, notice, appeal	Silent merging that obscures harm and the affected user
inclusion_path	Human or physical assistance and a fair deadline for deferred rights	Restricting continuity to one device class or user group

This profile is **an initial proposal for research and review**, not a government standard or the result of a workshop with real institutions. It requires development with legal, identity, and security specialists, and engagement with existing standards and Saudi frameworks without claiming affiliation with or approval by any body.

10. Economic and Strategic Consequences Through 2035

We do not tie economic effect merely to "hours of service availability." An institution may execute fifty offline acts and then bear a high cost to resolve five irreversible errors. It may defer only a few acts, but those acts may be so important that the social cost of waiting multiplies. The paper therefore proposes a **Multi-Objective Decision Frontier** spanning valid completion, exclusivity, duration of delay, potential harm, cost, and access disparities. We do not construct a national composite score before agreeing on weights and on whether the comparison is legitimate.

Four conditional research and policy options appear on the Saudi horizon:

1. Full centralization of final acts, with provisional-intake channels. This may be best where exclusivity cannot be guaranteed or law does not permit delegation. It requires clear receipt standards, preservation of requests and their timing, and appeal mechanisms.

2. Preallocated, disjoint authority for tightly specified acts. This may help where bounds are manageable and audits are persuasive, while imposing the cost of frozen capacity during disruption.

3. Delegation based on short-lived signed evidence. This may be suitable for displaying information or procedures with limited effect, but it faces recent-revocation risk and insecure key custody. A design does not become lawful merely because the mathematics can verify its signature.

4. Development of national products for authority-profile and reconciliation management.

This is a possible field for software that tests and models acts and entitlements. A market is proven only when a product can be sold, repeatedly secured and operated for customers, and purchased again because of its results—not when a national map is drawn with arrows.

These are not forecasts with assigned probabilities or official outcomes. A research body could create a future prediction register if it specifies a measurable event, resolution date, outcome source, and probability rules, without backdating its issue date.

11. Scientific Limitations: Where Does This Research's Right to Speak End?

- 1. Measurement:** The paper contains no data on Saudi services, availability, delegated-authority shares, incidents, or the laws that apply to them. The toy result for the number three cannot be transferred to any real amount or register.
- 2. Computation:** The exhaustive search covers 25 demand states, not the state space of real systems with thousands of roles, keys, devices, revocations, coercive situations, and human errors.
- 3. Security:** The model does not prove resistance to compromise of hardware, keys, messages, or payment channels. It includes a copyable-token example specifically to show the limit of signature validity.
- 4. Social representation:** Each state in the example is only a pair of numbers. Differences in people's capabilities, the importance of the act, and rights of appeal and remedy appear only as requirements for a proposed study.
- 5. History:** Neither the Iberian episode nor the British Library case can establish the quality of a particular local design. The sources record service disruption or partial continuity; they do not measure the effect of the proposed Offline Authority Profile.
- 6. Bias:** Stop and provisional-intake policies are included as genuine alternatives. The blind policy is deliberately designed to generate a violating example, so the comparison does not claim neutrality among deployed commercial solutions.

7. **Legitimacy:** The authority of people and institutions is a legal and administrative matter. An experimental model cannot issue permission in place of an authorized body. Defining permitted acts must precede technical research into their bounds.

12. The Next Decisive Study

The most useful next step is not to build a vast platform advertised as working offline. It is to choose **one public act** that can legally be received provisionally, another that can be completed under a valid delegated allocation if policy permits, and a third that must stop. An authorized test environment would give each explicit legal outcomes and measure what happened to the request, evidence, right, and applicant **during disruption and after reconnection**.

The decisive paper we seek would be able to say that, under specified forms of disconnection and stated conditions, architecture X increased the valid acts completed relative to policy Y without violating approved invariants, and that its cost and social limits were a specified amount. It may find that the proposed design failed or that a simpler alternative performed better. Accepting such a result is part of the research method, not an editorial embarrassment to conceal.

13. Reproduction, References, and Data

Reproduction: Save the directory files, install the dependencies listed in `analysis/requirements.txt`, and run:

```
python3 analysis/reproduce.py
```

The program will reproduce `data/model_check_states.csv`, `data/model_check_summary.csv`, `data/model_assumptions.json`, `data/duplicate_token_counterexample.json`, and two figures in SVG and PNG. The package contains the complete code and synthetic data and does not depend on private operational datasets or access keys.

Availability of real-world sources: Articles and published reports are linked to their originating institutions and researchers. We do not redistribute their protected files or claim a license to underlying bank transaction data. The source register records the access state during preparation of the original package; that is a historical audit note, not a current restriction on document extraction. Every computational value in the model is based on synthetic variables declared in the code.

Verified References in Order of Citation

1. **S01** Banco de España. *How do we pay when the power goes out? Spain's blackout experience*. 27 April 2026. Original institutional blog. <https://www.bde.es/wbe/en/noticias-eventos/blog/como-pagamos-cuando-todo-se-apaga-la-experiencia-del-apagon-en-espana.html>
2. **S02** European Central Bank. *Keep calm and carry cash: lessons on the unique role of physical currency across four crises*. ECB Economic Bulletin, 2025. https://www.ecb.europa.eu/press/economic-bulletin/articles/2025/html/ecb.ebart202506_02~1a773e2ca3.en.html
3. **S03** ENTSO-E. *Final report of the expert panel on the blackout in Spain and Portugal*. Published 20 March 2026. <https://www.entsoe.eu/news/2026/03/20/entso-e-publishes-expert-panel-final-report-on-28-april-2025-blackout-in-spain-and-portugal/>
4. **S04** CaixaBank Research, Josep Mestres Domènech and Zoel Martín Vilató. *The economic impact of the blackout in detail*. 18 June 2025. <https://www.caixabankresearch.com/en/sectoral-analysis/consumption/economic-impact-blackout-detail>
5. **S05** Banco de España. *Situation of payment infrastructures and the banking sector in Spain after the blackout*. 29 April 2025. <https://www.bde.es/wbe/en/inicio/noticias/monitorizacion-del-impacto-del-apagon-el-sistema-financiero.html>
6. **S06** BIS Innovation Hub. *Project Polaris: handbook for offline payments with CBDC*. 11 May 2023. <https://www.bis.org/publications/project-polaris-handbook-offline-payments-cbdc>
7. **S07** BIS Innovation Hub. *Project Polaris: high-level design guide for offline payments*. 26 October 2023. <https://www.bis.org/publications/project-polaris-high-level-design-guide-offline-payments>
8. **S08** Joseph M. Hellerstein and Peter Alvaro. *Keeping CALM: When Distributed Consistency is Easy*. Author-hosted 2019 preprint; subsequent publication context noted separately in source ledger. <https://arxiv.org/html/1901.01930v2>
9. **S09** Peter Bailis and collaborators. *Coordination Avoidance in Database Systems*. Proceedings of the VLDB Endowment, volume 8, 2014. <https://www.vldb.org/pvldb/vol8/p185-bailis.pdf>
10. **S10** Marc Shapiro, Nuno Preguiça, Carlos Baquero and Marek Zawirski. *Conflict-Free Replicated Data Types*. SSS 2011, LNCS 6976. DOI: 10.1007/978-3-642-24550-3_29. <https://inria.hal.science/inria-00609399v2/document>
11. **S11** Digital Government Authority. *Controls of Risk and Business Continuity Management for Digital Government, V2.0; Guideline of Risk and Business Continuity Management for Digital Government, V1.0* (2025). <https://dga.gov.sa/sites/default/files/2024-02/Controls%20Of%20Risk%20and%20Business%20Continuity%20Management%20For%20Digital%20Government%20V2.0.pdf>

12. **S12** Saudi Central Bank. *Business Continuity Management Framework*, introduction, 27 February 2017, in force. <https://rulebook.sama.gov.sa/en/11-introduction-bcm-framework>
13. **S13** Saudi Central Bank. *Providing Backup Generators in Branches and Remittance Centers*, circular of 18 May 2021. <https://rulebook.sama.gov.sa/en/providing-backup-generators-branches-and-remittance-centers>
14. **S14** British Library. *British Library Cyber Incident Review*. 8 March 2024, primary post-incident account.
<https://www.bl.uk/files/v5dwkion/production/99206a2d1e9f07b35712b78f7d75fbb09560c08d.pdf/british-library-cyber-incident-review-8-march-2024.pdf?dl=>
15. **S15** British Library. *Restoring our services – November 2025 update*. 21 November 2025.
<https://www.bl.uk/stories/news/restoring-our-services-november-2025-update>
16. **S16** European Central Bank. *Call for expression of interest to provide feedback on the list of standards relevant for the offline digital euro*. 18 August 2026.
<https://www.ecb.europa.eu/press/intro/news/html/ecb.mipnews260818.en.html>
17. **S17** European Central Bank. *FAQs on the digital euro*. Updated 17 August 2026; conditions for development, legislation and potential 2029 issuance.
https://www.ecb.europa.eu/euro/digital_euro/faqs/html/ecb.faq_digital_euro.en.html