DECISION SOVEREIGNTY

DECISION SOVEREIGNTY

A Governance Architecture for the Authority Layer of Critical Infrastructure Resilience

SYED ABID SHAH
Founder, Aquarian Systematic Resilience

COMPILED EDITION

ISBN- 9798176327373

Dmitri Dits,

with sincere appreciation for your thoughtful insights and shared commitment to advancing architectural resilience—may this companion volume spark many more meaningful conversations.





Syed Abid
Founder: Aquarian Systematic Resilience
26 November 2026

Table of Contents

Preface 4

Chapter 1 Decision Sovereignty 6

Chapter 2 From Data Sovereignty to Decision Sovereignty 17

Chapter 3 Decision Latency 26

Chapter 4 Human Authority Boundaries in Autonomous Infrastructure 33

Chapter 5 Federated Decision Continuity 42

Chapter 6 Restoring Authority 45

Chapter 7 Decision Sovereignty Under Real Disruption 53

Part II The ASR Decision Sovereignty Framework 62

Preface

Resilience thinking has, for the most part, been built around a single question: will the physical and technical system survive? Infrastructure operators plan for redundancy, regulators mandate continuity of service, and crisis frameworks specify recovery time objectives with increasing precision. This book is organised around a different and, its author argues, prior question: when critical infrastructure is disrupted, does the legitimate authority to decide what happens next survive alongside it?

The seven chapters that follow develop that question along its principal dimensions. Chapter 1 introduces Decision Sovereignty itself, and the Authority, Decision and Execution triangle at its centre. Chapter 2 grounds the concept theoretically, situating it within a four-stage sovereignty chain running from data to information to decision to action. Chapter 3 gives the concept a measurement discipline, decomposing the interval between detection and authorised response into a seven-stage Decision Latency Chain. Chapter 4 applies the framework to automation, defining the Autonomous Decision Boundary that separates machine-led from human-led authority. Chapter 5 extends it to federated and disconnected governance. Chapter 6 closes the loop with the Authority Recovery Protocol, the formal process by which delegated, automated or locally exercised authority is reviewed and returned to normal governance. Chapter 7 tests all six concepts against a single, real, documented event: the collapse of Pakistan's national power grid on 23 January 2023.

A final part, drawn from the research theme's capstone document, converges the six pillars into one governance architecture and

translates them into three outputs meant for direct use — a Diagnostic, a Measurement discipline, and an Architecture specification — rather than further theory.

The argument that runs through all of it is stated most plainly in Chapter 1, and is worth repeating here: a system can remain technically intact and still fail, if the authority to direct it fragments faster than the disruption itself can be contained. That failure mode is easy to miss precisely because nothing appears to be broken. This book exists to give it a name, a structure, and a set of tools.

CHAPTER 1

Decision Sovereignty

The Missing Layer in Critical Infrastructure Resilience

ABSTRACT

Resilience frameworks have historically measured whether critical infrastructure can absorb disruption and continue operating: whether assets survive, whether services resume, whether flows of resources, information and communications are restored. This paper argues that a prior and more fundamental question is frequently left implicit in that measurement: whether legitimate authority to decide what happens next survives the same disruption. Introducing Decision Sovereignty as the authority layer within the Aquarian Systematic Resilience (ASR) doctrine, this paper defines the concept, establishes its three constituent requirements of decision authority, decision integrity and decision execution, and positions it explicitly within the existing Flow Sovereignty architecture and the Six Sovereign Layers. It closes by outlining the practical consequences for how infrastructure operators, regulators and crisis planners assess systemic resilience.

1. Introduction

Modern resilience thinking has matured considerably over the past two decades. Infrastructure operators plan for redundancy, regulators mandate continuity of service, and crisis frameworks specify recovery time objectives with increasing precision. Yet across this literature, one question is asked far less often than its importance warrants:

Who retains the authority and capability to decide what happens next?

The question is easy to overlook because, under normal conditions, the answer is settled by institutional design. Authority sits where organisational charts, legal instruments and standard operating

procedures place it, and decision-making proceeds along channels that are rarely examined because they are rarely tested. Disruption changes this. When normal command centres become unavailable, communications fragment, data becomes incomplete or contradictory, and multiple actors, human and automated, begin acting independently, the settled answer can no longer be assumed. It must be established under pressure, often for the first time.

This paper argues that resilience frameworks which measure only the survival of physical assets, data and communications are measuring a necessary but insufficient condition. A system can remain technically intact and still fail, if the authority to direct that system fragments faster than the disruption itself can be contained. The result is a distinctive and under-examined failure mode: the infrastructure still exists, but governance begins to dissolve around it.

2. The Missing Layer: From Flow to Decision

The Aquarian Systematic Resilience doctrine has, to this point, examined resilience primarily as a question of continuity of flow: whether resources, information and services can continue moving through a system under stress. Flow Sovereignty and the Six Sovereign Layers ask whether the pathways along which infrastructure operates remain viable when normal conditions fail.

Decision Sovereignty asks a narrower and logically prior question. Before anything can flow, someone or something must decide that it should. Before a pathway can be used, someone must authorise its use. The movement of resources, information and services during disruption is downstream of a decision, and that decision is itself vulnerable to disruption in ways that are distinct from, and often less visible than, the physical or technical vulnerabilities that dominate conventional resilience planning.

During a major disruption, several conditions can arise simultaneously, each of which erodes the settled answer to who decides:

- normal command centres become unavailable
- communications become fragmented across surviving channels
- data becomes incomplete, delayed or contradictory
- institutional authorities become uncertain of their own standing
- multiple actors begin making independent decisions in parallel
- automated systems continue operating without renewed authorisation
- emergency powers may alter normal authority relationships without warning

None of these conditions requires the physical infrastructure to fail. Power may still flow, data may still be collected, communications equipment may still function. What fails, or risks failing, is the chain of legitimate authorisation that connects observation to action. This is the layer conventional resilience assessment tends not to examine directly, and it is the layer this paper proposes to name and define.

3. Defining Decision Sovereignty

ASR defines Decision Sovereignty as:

The continuing capacity of a legitimate governing authority to determine, authorise and direct critical actions despite disruption to conventional information, institutional, technological or command pathways.

Three elements of this definition merit particular attention. First, the emphasis on legitimate authority distinguishes Decision Sovereignty from mere decision-making capacity. A system, automated or human, may retain the technical ability to act without retaining the standing to act legitimately, and the distinction matters both operationally and institutionally: decisions taken without legitimate authority tend to be contested, reversed or distrusted even when they are technically correct, which itself becomes a further source of delay and fragility.

Second, the definition is deliberately conditional on disruption to pathways, not merely to assets. A decision authority can be intact as an institution while losing the pathways through which it would normally receive information, issue instructions or verify that instructions were followed. Decision Sovereignty is therefore as much a property of the connective tissue between authority and action as it is a property of the authority itself.

Third, the definition names three distinct verbs: determine, authorise and direct. These map onto the three constituent requirements developed below, and their separation is intentional. A system can satisfy one without satisfying the others, and understanding where the separation occurs is often the most useful diagnostic step in assessing a system's actual resilience.

4. The Authority, Decision and Execution Triangle

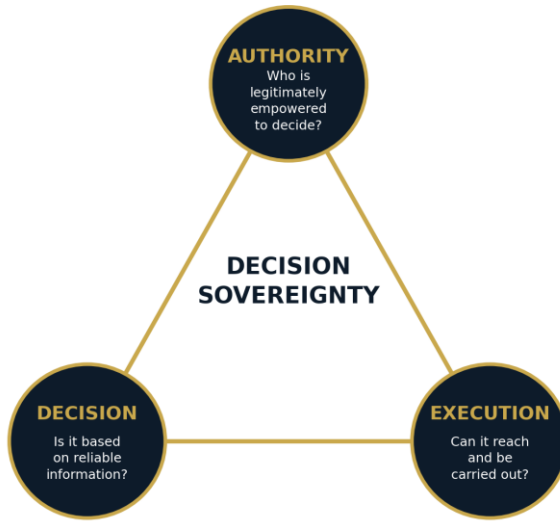

Figure 1.1 — The Authority, Decision and Execution triangle.

Decision Sovereignty introduces three requirements, each of which can fail independently:

Decision authority --- who is legitimately empowered to decide? This requires that the question of standing be answered clearly, in advance, and in a way that survives the loss of the normal institutional context in which it was originally settled.

Decision integrity --- is the decision based on sufficiently reliable information? Authority without a trustworthy operating picture produces decisions that are legitimate in form but unreliable in substance, which is frequently more dangerous than an acknowledged absence of authority, because it is less likely to be questioned.

Decision execution --- can the decision actually reach the operational layer and be carried out? A correct decision, taken by the correct authority, on the basis of reliable information, still fails if the pathway from decision to action is itself unavailable.

These three requirements form a triangle that recurs throughout this research theme:

Authority → Decision → Execution

Failure at any vertex of this triangle can produce systemic paralysis, even where the other two vertices remain fully intact. A central contribution of this paper, and of the theme it opens, is the argument that resilience assessment should examine all three vertices explicitly and separately, rather than treating decision-making as a single undifferentiated capability that either exists or does not.

5. Relationship to Flow Sovereignty and the Six Sovereign Layers

Decision Sovereignty is not presented here as a parallel or competing framework to Flow Sovereignty. It is the authority layer within it, and the Authority, Decision and Execution triangle set out above is best understood as the operational test of whether any of the Six Sovereign Layers can actually be acted upon once disruption begins.

A layer can retain full technical availability, resources present, data intact, communications functioning, and still fail this test if no legitimate authority can determine what should be done with it, or if a determination cannot reach the point of execution. Conversely, examining Decision Sovereignty without reference to the Six Sovereign Layers would produce an abstract governance theory disconnected from the physical and technical conditions that actually shape a crisis. The two are intended to be read together: Flow

Sovereignty asks whether the system can move; Decision Sovereignty asks whether anyone can legitimately tell it to.

This positioning also clarifies the relationship between Decision Sovereignty and two concepts already established elsewhere in the ASR doctrine. Commander's Intent and Human Integrity Density, developed in the governance paper on distributed AI under systemic stress, established that automation cannot independently generate legitimacy in a specific and important case, the operation of AI systems within critical infrastructure. Decision Sovereignty generalises that finding, from AI systems specifically to every actor, human or automated, that might act during disruption. The fourth paper in this theme develops that generalisation directly.

Mesh Governance and 72-Hour Autonomy, similarly, describe how a node keeps functioning when disconnected from the wider network. Decision Sovereignty asks the governance question underneath that functioning: what may a disconnected node legitimately decide on its own, and when does that delegated authority expire. The fifth paper in this theme develops that question directly.

6. Illustrative Dynamics

The failure mode Decision Sovereignty is designed to detect is easiest to see through a generic sequence, illustrative rather than drawn from any single documented event. Consider a critical facility operating normally, with a clearly designated decision authority, reliable monitoring, and functioning communication to both operational staff and higher governance. A disruption begins: a technical fault, a cyber intrusion, a natural hazard. In the first minutes, monitoring systems detect the anomaly, but the normal channel through which that detection would reach the designated authority is itself degraded by the same disruption.

At this point, three things typically happen in parallel. Automated protection systems, designed for exactly this scenario, begin taking pre-programmed action without waiting for renewed human authorisation. Operational staff on site, unable to reach the normal authority, begin making judgement calls based on local knowledge and immediate necessity. And the designated authority, once reached, begins issuing instructions based on a picture of the situation that is already several steps out of date.

None of these three actors is behaving illegitimately in isolation. But together they produce a system in which three separate decision-making processes are running simultaneously, without reconciliation, each with a partial and different picture of events. This is decision fragmentation, and it is the specific pathology Decision Sovereignty is designed to name, measure and, eventually, design against. It is materially different from asset failure, and it can occur even where every physical asset in the system continues to function exactly as designed.

7. Implications for Practice

Treating Decision Sovereignty as a distinct layer of resilience carries several practical consequences for infrastructure operators, regulators and crisis planners:

Resilience assessments should examine the Authority, Decision and Execution triangle explicitly, rather than assuming that decision-making capability is intact wherever physical and technical systems are intact.

Decision authority should be pre-established and pathway-independent, meaning it should not rely on any single communication channel or command centre to remain legible during disruption.

Automated systems should have explicit, documented boundaries on independent action, developed further in the fourth paper of this theme, so that speed of automated response does not substitute for legitimacy of automated response.

Disconnected nodes should have pre-defined delegated authority with clear expiry conditions, developed further in the fifth paper of this theme, so that local decision-making during isolation does not become a permanent or unreviewed departure from normal governance.

Post-crisis review should explicitly reconcile parallel decision processes, developed further in the sixth paper of this theme, treating the recovery of normal authority as a distinct task from the recovery of normal operations.

8. Conclusion

Infrastructure resilience has been substantially advanced by frameworks that ask whether physical and technical systems survive disruption. This paper has argued that a further question deserves equal and explicit attention: whether legitimate authority to decide survives alongside them. Decision Sovereignty names this question, defines it through the Authority, Decision and Execution triangle, and positions it as the authority layer within the existing Flow Sovereignty architecture rather than as a competing framework.

The six papers that follow in this research theme develop the concept along its principal dimensions: its theoretical grounding relative to data and information sovereignty, its measurement through decision latency, its application to autonomous infrastructure, its extension to federated and disconnected governance, its resolution through authority recovery, and its grounding in a real, previously documented disruption event.

Together, they are intended to establish Decision Sovereignty as a distinct and necessary layer of resilience assessment, not a restatement of work already published under Flow Sovereignty, but its logical completion.

***"

CHAPTER 2

From Data Sovereignty to Decision Sovereignty

The Next Frontier of Digital Governance

ABSTRACT

Digital governance discourse over the past decade has organised itself substantially around data sovereignty: questions of ownership, storage location, access rights and jurisdictional control over data. This paper argues that data sovereignty, while necessary, is an incomplete foundation for governance, because data itself does not act. It is interpreted, and that interpretation is acted upon by people and systems whose authority to act is a separate and frequently unexamined question. Building on the Decision Sovereignty concept introduced in the first paper of this theme, this paper proposes a four-stage sovereignty chain, running from data sovereignty through information sovereignty to decision sovereignty and finally to action sovereignty, and argues that a state or organisation may possess full sovereignty over its data while lacking sovereignty over the decisions ultimately made with it. The distinction has direct application to national digital strategy, critical infrastructure governance and the assessment of dependency on foreign technology providers.

1. Introduction

Data sovereignty has become one of the organising concepts of contemporary digital policy. Governments legislate on where data may be stored. Regulators specify who may access it and under what conditions. Enterprises restructure infrastructure to satisfy jurisdictional requirements. The underlying assumption, rarely stated but widely held, is that control over data is the decisive form of digital control: secure the data, and the governance question is substantially answered.

This paper argues that the assumption is incomplete in a way that becomes consequential precisely during the conditions this research

theme is concerned with: disruption, crisis and systemic stress. Data, however securely held, does not by itself determine anything. It must be interpreted, and interpretation requires systems, models and institutional processes that may sit entirely outside the jurisdiction that holds sovereignty over the underlying data. The interpreted picture must then inform a decision, taken by an authority whose legitimacy and standing are separate questions again. And the decision must be executed by systems or actors capable of carrying it out. At each step, sovereignty over the previous step does not guarantee sovereignty over the next.

2. The Sovereignty Chain



Figure 2.1 — The sovereignty chain, from data control to executed action.

ASR proposes that digital governance be understood not as a single question of data control, but as an evolutionary chain of four distinct sovereignties, each building on the last and each capable of failing independently of the others:

Data Sovereignty

Who controls the data?

↓

Information Sovereignty

Who controls interpretation and knowledge?

↓

Decision Sovereignty

Who has legitimate authority to decide?

↓

Action Sovereignty

Who can execute the decision?

Data sovereignty is the most extensively legislated and discussed of the four, and this paper does not dispute its importance. But it is properly understood as the raw material stage of the chain: it establishes control over the record, not over what the record is taken to mean.

Information sovereignty concerns who performs the interpretation of that data, through analytics, models, dashboards and human expertise, and therefore who shapes the picture of reality on which subsequent decisions will be based. A jurisdiction may hold its data locally while depending entirely on foreign-built analytical models, foreign-trained artificial intelligence systems, or foreign-employed technical expertise to turn that data into usable knowledge. In such cases, the party that interprets the data exercises a form of

influence over outcomes that data sovereignty, narrowly defined, does not capture.

Decision sovereignty, as established in the first paper of this theme, concerns who holds the legitimate authority to determine, authorise and direct action on the basis of that interpreted information. This is the layer at which governance either holds or fragments under stress, and it is the central concern of this research theme as a whole.

Action sovereignty concerns whether the decision, once legitimately taken, can actually be executed by systems and actors under the deciding authority's effective control. A decision authorised correctly but executed through infrastructure, platforms or personnel outside that authority's control is only nominally sovereign; the final step of the chain has been outsourced along with the infrastructure.

3. Possessing Data Without Possessing Decisions

The practical significance of separating these four sovereignties becomes clear when a state or organisation is assessed against all four simultaneously rather than against data sovereignty alone. A state may possess its data in the fullest legal and physical sense, stored within national borders, subject to national law, accessible only through nationally controlled systems, and still depend substantially upon:

- foreign cloud infrastructure for processing at scale
- foreign-developed artificial intelligence models for interpretation
- foreign-operated communications backbones for transmission

- foreign platforms for coordination and decision support
- external technical operators for system maintenance and incident response

In each of these cases, data sovereignty is intact while information sovereignty, decision sovereignty, or action sovereignty may not be. The state possesses its data without necessarily possessing the decisions ultimately made from it, and this gap is frequently invisible under normal operating conditions, precisely because normal conditions do not test where the chain would actually break.

This is not a claim that dependency on foreign providers is inherently illegitimate or unmanageable; most digital infrastructure globally involves some degree of cross-border dependency, and the sovereignty chain is a diagnostic tool rather than a prescription for self-sufficiency. Its purpose is to make an otherwise implicit dependency explicit, so that it can be assessed, managed and, where necessary, deliberately accepted rather than discovered under crisis conditions.

4. Why the Distinction Matters Under Disruption

The four-stage chain is of limited additional value under normal operating conditions, where all four sovereignties are typically exercised, formally or informally, by an overlapping set of actors and the seams between them are rarely tested. Its value becomes apparent under disruption, which is precisely when the seams are most likely to be exposed.

A cross-border technical dependency that is invisible in ordinary operation can become decisive during a crisis: a foreign-operated analytical platform may become unavailable at the exact moment its interpretive function is most needed; a foreign jurisdiction's export

controls, security posture or commercial priorities may constrain access to a model or platform precisely when domestic demand for it peaks; a communications backbone operated outside national control may be deprioritised, congested or unavailable during a regional or global incident affecting multiple customers simultaneously.

In each case, the state or organisation affected retains its data sovereignty throughout. What it loses, temporarily or permanently, is one or more of the three sovereignties built upon that data. National security and critical infrastructure discussions that assess digital dependency only in terms of data location will systematically underestimate this exposure.

5. Implications for Digital Governance

Treating digital sovereignty as a four-stage chain rather than a single condition carries several practical implications:

Sovereignty assessments should be conducted separately at each stage, rather than inferring information, decision and action sovereignty from data sovereignty alone.

Dependency mapping should trace the full chain for critical functions, identifying at which stage, if any, control passes outside the deciding authority's jurisdiction.

Contingency planning should address interpretation and execution, not only data access, since a data breach and an interpretation-layer outage produce different failure modes requiring different responses.

Procurement and platform decisions should be evaluated against the chain as a whole, recognising that a technically excellent platform

which sits outside the deciding authority's control transfers a portion of sovereignty along with its convenience.

6. Positioning Within the Research Theme

This paper's principal contribution to the wider research theme is theoretical grounding: it situates Decision Sovereignty, introduced in the first paper, within a broader and already familiar policy vocabulary, connecting it to debates on data localisation, digital sovereignty and technological dependency that are already active in government and industry. Decision sovereignty is not presented here as a replacement for those debates but as their logical continuation, the layer those debates most often gesture toward without naming directly.

The remaining papers in this theme build outward from this positioning. The third paper develops decision latency as a means of measuring how quickly the chain, once established, can actually operate under stress. The fourth and fifth papers examine how decision sovereignty is preserved when action sovereignty is delegated to autonomous systems and disconnected nodes respectively. The sixth paper addresses how decision sovereignty, once fragmented by crisis, is formally restored.

7. Conclusion

Data sovereignty has rightly occupied a central place in digital governance policy, but it answers only the first of four questions that determine whether a state or organisation genuinely governs its own digital and infrastructural future. Information, decision and action sovereignty each represent a further stage at which control can be exercised, or quietly ceded, independently of data sovereignty itself. Recognising the chain in full allows dependency to be assessed accurately rather than assumed away, and positions decision

sovereignty, the subject of this research theme, as the hinge on which the practical value of the entire chain ultimately turns.

***"

CHAPTER 3

Decision Latency

A New Measure of Systemic Resilience

ABSTRACT

Conventional resilience metrics, downtime, recovery time, availability, measure how long a system takes to return to normal operation. This paper proposes a complementary and, in several respects, more diagnostic measure: decision latency, the time elapsed between the detection of a disruptive event and the execution of a legitimately authorised response to it. Two systems experiencing identical physical disruption can produce materially different outcomes depending on how quickly the Authority, Decision and Execution triangle established in the first paper of this theme actually completes under stress. This paper defines a seven-stage Decision Latency Chain, introduces the concept of an ASR Decision Latency Index, and sets out the measurement discipline required before that index can be used as a scoring tool rather than a diagnostic framework.

1. Introduction

Resilience is conventionally measured through outcome metrics: how long a system was unavailable, how quickly service was restored, what proportion of capacity survived. These metrics are necessary and well established. They are also, this paper argues, incomplete in a specific way: they measure the consequence of a disruption without directly measuring the process that determines how severe the consequence becomes.

Consider two systems subject to an identical physical disruption of the same type, scale and duration. System A detects the disruption, escalates it to the appropriate authority, receives an authorised response, and executes that response within minutes. System B experiences an identical physical event but takes hours to complete the same sequence. The physical damage caused by the initiating event may be indistinguishable between the two systems. The

consequences, in terms of cascading failure, secondary damage and operational disruption, are unlikely to be.

Decision latency is a leading indicator of systemic failure.

This paper argues that the interval between detection and authorised execution deserves to be measured directly, as a distinct property of a system's resilience architecture, rather than inferred only after the fact from how badly an incident turned out.

2. The Decision Latency Chain



Figure 3.1 — The seven-stage Decision Latency Chain.

Building on the Authority, Decision and Execution triangle established in the first paper of this theme, decision latency can be decomposed into a seven-stage chain, each transition representing a point at which delay, and therefore risk, can accumulate:

Detection → Interpretation → Escalation → Authorisation → Decision → Communication → Execution

Detection --- the interval before an anomaly is registered as significant by any monitoring system or observer.

Interpretation --- the interval required to convert a raw signal into an assessed situation, connecting directly to the information sovereignty question raised in the second paper of this theme.

Escalation --- the interval required for the assessed situation to reach an actor with the standing to decide.

Authorisation --- the interval required for that actor to confirm their own standing to decide, which can itself be delayed where authority is contested or unclear.

Decision --- the interval required to select a course of action once authorised to do so.

Communication --- the interval required for the decision to reach the point of execution intact and unambiguous.

Execution --- the interval required for the decision, once received, to be physically or operationally carried out.

A system's total decision latency is the sum of these seven intervals, but the sum alone understates the diagnostic value of the chain. In practice, one transition is typically responsible for the majority of total delay, and identifying which one is usually more useful for resilience design than the aggregate figure itself.

The weakest transition determines the effective response time.

A system with an excellent detection capability and a fast execution capability can still perform poorly overall if authorisation is routinely delayed by unclear or contested standing to decide, a failure mode invisible to metrics that measure only total elapsed time.

3. Toward an ASR Decision Latency Index

The seven-stage chain suggests the possibility of a composite ASR Decision Latency Index (DLI), intended to allow comparison of decision-making speed across systems, sectors and time, in the way

that recovery time objectives already allow comparison of restoration speed. A candidate index would examine, at minimum:

- normal decision latency, measured under routine operating conditions
- crisis decision latency, measured under simulated or actual disruption
- authority-transfer latency, the added delay when a decision must move to a different or emergency authority
- communication latency, isolated from decision-making time itself
- execution latency, isolated from the time required to reach the decision
- recovery-to-normal-authority latency, connecting forward to the authority recovery process addressed in the sixth paper of this theme

Such an index would be particularly relevant to sectors where the gap between normal and crisis decision latency tends to be largest and most consequential: ports and logistics networks, energy transmission and distribution, transportation control systems, emergency management coordination, and government continuity operations.

4. Measurement Discipline

A composite index of this kind carries an obligation that this paper treats as a precondition for its practical use rather than an afterthought. ASR's existing measurement instrument, the Net Fragility Matrix, earned its credibility within the doctrine through a defined weighting methodology, documented back-test rationale, stated inter-rater reliability thresholds and a formal validation appendix. The Decision Latency Index should be held to the same standard before it is presented as a scoring or comparison tool.

This paper is deliberately scoped as conceptual

It defines the seven-stage chain and the dimensions a DLI would need to examine. It does not publish weightings, scoring thresholds or comparative results, because none has yet been validated against back-tested incident data in the manner ASR-VAM was. Two paths are available before the index can be used operationally:

- a dedicated technical companion paper, following the ASR-VAM precedent, presenting worked calculations, a stated weighting methodology and back-test rationale against documented incidents; or
- continued use of the seven-stage chain as a qualitative diagnostic tool, identifying the weakest transition in a given system without assigning it a numerical score, until such validation work is complete.

Publishing a numerical index ahead of this validation work would risk exactly the outcome ASR-VAM's appendices were built to avoid: a measurement that appears rigorous without being reproducible. The seven-stage chain itself, by contrast, is immediately usable as a structured diagnostic, independent of whether or when a validated numerical index follows it.

5. Applications

Even in its present, pre-validation form, the Decision Latency Chain offers practical value to infrastructure operators and crisis planners. It provides a shared vocabulary for post-incident review, allowing an organisation to ask, at each of the seven stages, where delay actually occurred, rather than treating response time as an undifferentiated total. It also provides a design checklist for new systems: each of the seven transitions can be examined at the design stage for single points of failure, undefined authority, or communication pathways that depend on infrastructure likely to be degraded by the same disruption the system is meant to respond to.

It further connects directly to the Autonomous Decision Boundary developed in the fourth paper of this theme. Automated systems are, in effect, an attempt to compress the Decision Latency Chain by removing or shortening several of its transitions. The chain developed here gives that compression a precise vocabulary: an autonomous system does not simply act faster in general, it removes or automates specific, nameable stages of the chain, and each removal is a governance decision in its own right, not merely a technical one.

6. Conclusion

Outcome-based resilience metrics answer how badly a system was affected by disruption. Decision latency asks a prior question: how long did it take the system's governance, not merely its infrastructure, to respond legitimately once disruption began. The seven-stage Decision Latency Chain developed in this paper makes that interval visible and decomposable, identifying the weakest transition rather than treating response time as an undifferentiated total. A validated numerical index remains future work, to be developed with the same discipline that established the Net Fragility Matrix. The chain itself, however, is available immediately as a diagnostic and design tool, and forms the measurement foundation the remaining papers in this theme draw upon.

***"

CHAPTER 4

Human Authority Boundaries in Autonomous Infrastructure

Commander's Intent Made Operational

ABSTRACT

As critical infrastructure becomes increasingly autonomous, governance discussion has tended to concentrate on a single question: how autonomous should a given system be. This paper argues that the more consequential question is a different one: where does legitimate human authority begin and end within a system that is partly autonomous. Building on the Decision Sovereignty concept developed in this theme, and directly extending the Commander's Intent and Human Integrity Density framework established in ASR's earlier governance paper on distributed artificial intelligence, this paper introduces the Autonomous Decision Boundary: an explicit, documented structure separating detection, recommendation, automated action, human authorisation, institutional override and emergency authority. It argues that automation may legitimately increase the speed of a decision, but cannot by itself establish the legitimacy of that decision, and that accountability for autonomous action in critical infrastructure must therefore be designed in advance rather than discovered after an incident.

1. Introduction

Autonomous systems are increasingly embedded in the operation of critical infrastructure: intrusion detection that triggers automatic isolation, grid protection systems that shed load without waiting for a control room instruction, logistics platforms that reroute shipments algorithmically in response to disruption. The conventional framing of the governance question this raises is a question of degree:

How autonomous should the system be?

This paper argues that the question of degree, while not unimportant, is secondary to a question of structure. An autonomous system may be highly capable and still operate within a clearly bounded and accountable relationship to human authority, or it may be modestly capable and still operate in a way that obscures who is actually accountable for what it does. The relevant question this paper poses instead is:

Where does legitimate human authority begin and end within an autonomous system?

2. Technical Capability Is Not Legitimate Authority

An autonomous infrastructure system may be technically capable of detecting and responding to a wide range of conditions without human involvement:

- a cyberattack in progress
- abnormal traffic or throughput patterns
- grid or network instability
- a supply-chain interruption
- physical intrusion at a facility
- equipment failure or degradation

In each case, the system may be capable of responding faster and more consistently than a human operator. But technical capability to act does not, by itself, establish legitimate authority to act, in the sense developed in the first paper of this theme. Decision Sovereignty requires that action be taken by an actor with standing to take it, on the basis of sufficiently reliable information, in a way that can actually reach execution. An autonomous system can satisfy the third of these

requirements very well, and the second reasonably well, while the first, standing, remains a governance question that the system's technical design does not itself answer.

The distinction matters because the two are easily conflated in practice. A system that responds quickly and correctly to a wide range of conditions tends to accumulate implicit trust, and with it, an informal expansion of what it is allowed to decide, that is rarely reviewed as a deliberate governance choice. This paper argues that the boundary of an autonomous system's authority should be an explicit design decision, not an emergent property of how well the system happens to perform.

3. The Autonomous Decision Boundary

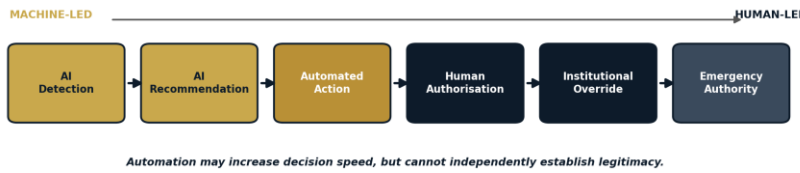

Figure 4.1 — The Autonomous Decision Boundary, from machine-led to human-led authority.

ASR proposes that every autonomous critical-infrastructure system should have an explicitly defined and documented boundary structure, separating six distinct positions along a single continuum:

AI Detection → AI Recommendation → Automated Action → Human Authorisation → Institutional Override → Emergency Authority

AI Detection --- the system identifies and characterises an anomaly, without taking or recommending action.

AI Recommendation --- the system proposes a course of action to a human decision-maker but does not execute it.

Automated Action --- the system executes a pre-authorised category of response without seeking authorisation in the moment, because that category of response has already been authorised in advance.

Human Authorisation --- a designated human authority must approve the specific action before it is executed.

Institutional Override --- a higher institutional authority retains the standing to reverse or halt any of the preceding categories of action.

Emergency Authority --- a distinct, pre-defined authority activates specifically under declared emergency conditions, with its own scope and expiry.

The Autonomous Decision Boundary is the documented answer to which of these six positions applies to which category of decision, for a given system, in advance of any incident. Two systems with identical technical capability can have very different boundary structures, and this paper argues that the boundary structure, not the technical capability, is the primary object of governance review.

Automation may increase decision speed, but it cannot independently establish legitimacy.

4. Relationship to Commander's Intent and Human Integrity Density

The Autonomous Decision Boundary developed in this paper is not a freestanding governance proposal. It is the operational mechanism for Commander's Intent, the governance principle established in ASR's earlier paper, "Flow Sovereignty and Commander's Intent: A Technical Governance Framework for

Distributed AI Under Systemic Stress." That paper established, in the specific context of AI operating within distributed critical infrastructure, that automation cannot independently generate legitimacy, and introduced Human Integrity Density as a means of assessing the degree to which human judgement remains meaningfully embedded in an AI-assisted decision process under systemic stress.

This paper generalises that finding from AI systems specifically to any autonomous or semi-autonomous actor operating within critical infrastructure, and translates the governance principle into a specific, testable structure. Where Commander's Intent establishes that human judgement must remain present in substance, the Autonomous Decision Boundary specifies exactly where along the detection-to-execution continuum that judgement is required to sit for a given category of decision, and where it may legitimately be pre-authorised in advance. Human Integrity Density can be understood as a measure of how faithfully a system's actual operation adheres to its documented Autonomous Decision Boundary under stress, rather than drifting toward automated action beyond what was authorised.

This paper should be read as an extension of that earlier governance paper rather than as a separate treatment of the same territory, and the two are intended to be cited together in any application of either.

5. Accountability Under Autonomous Action

ASR's central contribution to the broader AI governance discussion is not primarily a claim about whether autonomous systems can be trusted. It is a narrower and more specific question:

Who remains accountable when AI becomes operationally decisive?

The Autonomous Decision Boundary gives that question a structural answer. Accountability for a given autonomous action traces to whichever position on the boundary authorised that category of action in advance, whether that is a human authoriser acting in the moment, an institutional authority that pre-approved a category of automated action, or an emergency authority whose scope was defined before the emergency began. Where no position on the boundary can be identified as having authorised a given action, this paper treats that as a governance failure independent of whether the action itself produced a good or bad outcome. A system that acts correctly without a traceable authorisation has still exposed a gap that a system acting incorrectly under a properly documented boundary has not.

6. Design Implications

Treating the Autonomous Decision Boundary as an explicit design artifact carries several practical consequences:

Boundary documentation should precede deployment, not be reconstructed after an incident to explain what the system was allowed to do.

Categories of automated action should be pre-authorised explicitly and narrowly, rather than inferred from the system's general mandate to protect or maintain operations.

Institutional override should be technically as well as procedurally real, meaning a higher authority must retain the actual capability, not merely the nominal right, to halt or reverse automated action.

Emergency authority provisions should have explicit expiry conditions, connecting directly to the Authority Recovery process developed in the sixth paper of this theme, so that emergency scope does not persist informally after the emergency ends.

The boundary should be reviewed against actual decision latency data, drawing on the Decision Latency Chain developed in the third paper of this theme, since a boundary that routes too many decisions through slow human authorisation may be as much a resilience risk as one that automates too much.

7. Conclusion

The governance question posed by increasingly autonomous critical infrastructure is less usefully framed as how much autonomy a system should have than as where, precisely, legitimate human authority sits along the path from detection to execution. The Autonomous Decision Boundary developed in this paper makes that structure explicit and documentable, and operationalises the Commander's Intent principle established elsewhere in the ASR doctrine for the general case of autonomous action in critical infrastructure. Automation can shorten the Decision Latency Chain considerably. It cannot, on its own, answer the question of who was entitled to decide, and that question must be settled in advance, not reconstructed after the fact.

***"

CHAPTER 5

Federated Decision Continuity

Decision Authority in Federated and Networked Governance

Federated Decision Continuity specifies, for every node in a networked or federated system, what it may decide independently, what requires higher authority, and when that delegated authority activates and expires, so that a disconnected node can keep deciding legitimately, not merely keep operating.

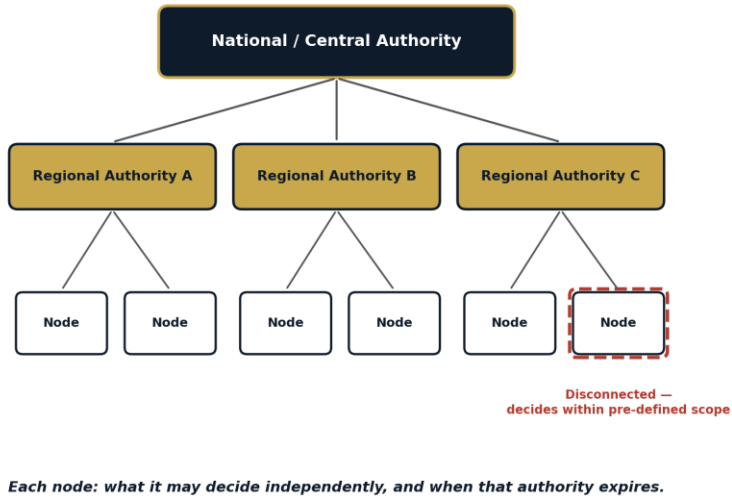

Figure 5.1 — Federated Decision Continuity across a networked authority structure.

What the Rest of the Theme Establishes About This Concept

The first paper in this theme, in positioning Decision Sovereignty relative to ASR's existing doctrine, observes that Mesh Governance and 72-Hour Autonomy — which describe how a node keeps functioning when disconnected from the wider network — leave

open a governance question underneath that functioning: what may a disconnected node legitimately decide on its own, and when does that delegated authority expire. Federated Decision Continuity is the paper in this theme that answers that question directly.

The capstone framework document positions the concept in the same terms: Mesh Governance and 72-Hour Autonomy describe a node's operational capacity to function while disconnected; Federated Decision Continuity gives that capacity its authority dimension, describing what the node may legitimately decide while it remains disconnected — the difference between a node that keeps running and a node that keeps governing itself legitimately.

The sixth paper in this theme, on Authority Recovery, connects forward to this concept from the other direction: its Authority Reconciliation stage is described as "the formal process through which independently taken decisions across a federated network are integrated once connectivity, and with it central visibility, returns" — the point at which decisions made independently under Federated Decision Continuity are reconciled back into a single account once the disruption that necessitated them has passed.

The applied case study in Chapter 7 finds a real-world instance of the pattern this concept is designed to address: the marked difference in restoration times across Pakistan's regions following the January 2023 grid collapse is consistent with a federated restoration process, in which regional transmission and distribution entities restored their own sections of the network at different rates rather than following a single, centrally sequenced plan — the public record does not establish whether that pattern reflected a pre-defined federated structure of the kind this concept recommends, or an improvised division of effort.

CHAPTER 6

Restoring Authority

Decision Recovery After Systemic Disruption

ABSTRACT

Crisis management frameworks devote extensive attention to the recovery of infrastructure: restoring assets, resuming services, returning systems to normal operating parameters. Considerably less attention is given to a parallel and equally necessary recovery: the recovery of normal governing authority itself. During disruption, decision authority may be temporarily delegated to emergency authorities, exercised independently by autonomous systems operating within pre-authorised boundaries, or taken locally by disconnected nodes under federated decision continuity arrangements. This paper addresses what happens next: how emergency, automated and locally delegated decisions are reviewed, how conflicting decisions taken in parallel are reconciled, and how normal institutional authority is formally and verifiably restored. It introduces the Authority Recovery Protocol as the closing stage of the Decision Sovereignty framework developed across this research theme, and argues that authority recovery is a distinct task from operational recovery, requiring its own explicit process rather than being assumed to follow automatically once services resume.

1. Introduction

Most established crisis and continuity frameworks are organised around a recognisable arc: detection, response, stabilisation, recovery, return to normal operations. The word recovery, in this arc, is almost always understood to mean the recovery of infrastructure and services, restoring power, restoring communications, restoring throughput to pre-incident levels. This paper argues that a second and separate recovery is required, and is frequently left implicit or assumed rather than deliberately managed:

The recovery of normal authority.

The preceding papers in this theme have established several mechanisms through which decision authority can legitimately shift during disruption: emergency authority, activated under declared crisis conditions and developed as one position on the Autonomous Decision Boundary in the fourth paper of this theme; automated action, taken by autonomous systems operating within pre-authorised categories, also developed in the fourth paper; and locally delegated authority, exercised by disconnected nodes under Federated Decision Continuity, developed in the fifth paper of this theme. Each of these mechanisms is designed to preserve legitimate decision-making under conditions where normal authority cannot function. None of them is designed to persist once those conditions end, and each therefore requires an explicit process for winding down.

2. The Problem of Unreviewed Authority

Suppose a local commander, an emergency authority, or an autonomous system receives temporary decision powers during a disruption, as anticipated by the governance structures developed elsewhere in this theme. Several questions become pressing once the disruption itself has stabilised, and this paper argues that established frameworks answer them inconsistently, if at all:

- What happens when the emergency conditions that triggered delegated authority end?
- Who verifies what decisions were actually made during the period of delegation?
- Who reviews those decisions for legitimacy, accuracy and consistency with authorised scope?
- How is authority formally transferred back to normal governance, rather than simply allowed to lapse?

- How are decisions that conflict with one another, taken independently by different actors during the same period, reconciled?
- How is institutional trust in the normal governance structure restored, particularly where delegated authority was exercised visibly and directly by the public or by subordinate institutions?

Where these questions are not answered explicitly, delegated authority tends to persist informally past the point at which it was actually needed, either because no clear process exists to end it, or because the actors who held it during the crisis have little incentive to relinquish it without being asked. This is a governance risk distinct from, and in some respects more durable than, the original disruption itself.

3. The Authority Recovery Protocol



Figure 6.1 — The six-stage Authority Recovery Protocol.

ASR proposes a six-stage Authority Recovery Protocol as the standard closing sequence for any episode in which decision authority was delegated, automated or exercised locally under disruption:

Emergency Delegation → Operational Autonomy → Stabilisation → Verification → Authority Reconciliation → Normal Governance Restoration

Emergency Delegation --- the point at which authority is transferred to an emergency, automated or local actor, ideally under conditions and scope defined in advance.

Operational Autonomy --- the period during which that actor exercises delegated authority to manage the disruption.

Stabilisation --- the point at which the conditions that necessitated delegation no longer apply, marking the intended end of the delegation period rather than an informal wind-down.

Verification --- a formal review of decisions taken during the delegation period, checking them against the scope that was actually authorised.

Authority Reconciliation --- resolution of any conflicts between decisions taken independently by different actors during the same period, connecting directly to the reconciliation process described in the fifth paper of this theme.

Normal Governance Restoration --- the explicit, documented act of returning full decision authority to normal institutional structures, closing the delegation episode formally rather than allowing it to lapse.

The protocol is deliberately sequential and deliberately explicit at each stage. Its purpose is to ensure that a delegation of authority, however necessary and well-designed at the point it was granted, does

not become a permanent or unreviewed departure from normal governance simply because no one was assigned the task of formally ending it.

4. Relationship to the Preceding Papers in This Theme

The Authority Recovery Protocol closes the loop opened by the second, fourth and fifth papers in this theme. Where the second paper established that decision sovereignty can be lost independently of data sovereignty, this paper addresses how it is regained. Where the fourth paper defined the Autonomous Decision Boundary's Emergency Authority position, this paper defines what happens when that position's activating conditions end. Where the fifth paper established Federated Decision Continuity for disconnected nodes, this paper's Authority Reconciliation stage is the formal process through which independently taken decisions across a federated network are integrated once connectivity, and with it central visibility, returns.

Verification, in particular, depends on the Decision Latency Chain developed in the third paper of this theme: a meaningful review of decisions taken under delegated authority requires reconstructing not only what was decided, but how quickly and through what chain of detection, escalation and authorisation each decision was reached, in order to assess whether the delegated actor operated within its intended scope and tempo.

5. Institutional Trust as a Recovery Objective

Authority Recovery has a dimension that is easy to treat as secondary but is, in practice, often decisive for how a disruption is judged after the fact: the restoration of institutional trust. Where emergency, automated or locally delegated authority was exercised

visibly, directly affecting the public or subordinate institutions, the formal act of Normal Governance Restoration also functions as a public signal that oversight has resumed and that the decisions taken during the delegation period have been reviewed rather than simply left unexamined.

A disruption that is resolved operationally but whose authority recovery is handled informally or invisibly can leave institutions and the public uncertain whether emergency conditions have genuinely ended, an uncertainty that itself becomes a source of continued instability even after the original disruption is fully resolved. Treating Verification and Normal Governance Restoration as explicit, documented and where appropriate publicly visible stages addresses this risk directly.

6. Design Implications

Treating authority recovery as a distinct, designed process carries several practical consequences:

Every delegation of emergency, automated or local authority should have a corresponding, pre-defined recovery process, not a general assumption that authority will revert once things return to normal.

Stabilisation conditions should be defined in advance and objectively verifiable, so that the end of a delegation period is not itself a contested judgement call.

Verification should be resourced and scheduled as a standard part of incident closure, rather than conducted only where a decision taken under delegation is separately challenged or disputed.

Authority Reconciliation should have a designated owner, particularly in federated structures where multiple nodes may have taken independent decisions during the same period.

Normal Governance Restoration should be a documented and, where relevant, communicated event, not an informal fade-out of emergency arrangements.

7. Conclusion

Infrastructure recovery and authority recovery are frequently treated as though the second follows automatically from the first. This paper has argued that they are distinct tasks, and that authority recovery, left unmanaged, tends to persist informally past the point of genuine necessity, eroding institutional trust even after the original disruption has been resolved. The Authority Recovery Protocol developed here, Emergency Delegation through Stabilisation, Verification, Reconciliation and Normal Governance Restoration, closes the Decision Sovereignty framework developed across this research theme, giving delegated, automated and locally exercised authority a defined and reviewable end, not merely a defined beginning.

***"

CHAPTER 7

Decision Sovereignty Under Real Disruption

The Pakistan Grid Collapse, 23 January 2023

A note on sourcing: this chapter applies the six Decision Sovereignty concepts developed across this book against the public factual record of the 23 January 2023 national grid collapse. It draws only on matters of public record, reflected in government statements and independent reporting at the time, and does not introduce proprietary data, internal ASR case file material, or claims beyond what is publicly documented.

ABSTRACT

The six theoretical papers that precede this one in the research theme develop Decision Sovereignty, its constituent concepts, and its measurement, largely as general propositions. This paper grounds them in a single, publicly documented event: the collapse of Pakistan's national power grid on 23 January 2023, in which a frequency disturbance originating in the southern transmission network cascaded into a near-total nationwide blackout affecting the country's more than 220 million people. Applying the Decision Sovereignty, Decision Integrity, Decision Latency, Autonomous Decision Boundary, Federated Decision Continuity and Authority Recovery concepts to the public record of this event, the paper argues that a purely asset-and-restoration reading of the incident, which dominated public and official commentary at the time, captures only part of what determined its course and its consequences.

1. Introduction

On the morning of 23 January 2023, at approximately 7:34 local time, a frequency disturbance in Pakistan's southern transmission network triggered a cascading failure that brought down the national grid, leaving the great majority of the country's population without electricity. Islamabad and Rawalpindi saw power restored within roughly eight hours; Lahore and Karachi took approximately sixteen; some rural areas remained without power for between one and three

days. Official estimates placed the economic cost in excess of PKR 100 billion, with specific losses reported in export-facing sectors such as textiles. The event was Pakistan's second near-complete national grid failure within two years.

Public commentary on the incident, consistent with the pattern this research theme has argued is typical of resilience assessment generally, concentrated overwhelmingly on the physical and operational dimension: the technical cause of the initial disturbance, the sequence of the cascading failure, and the speed of restoration by city and region. This paper does not dispute the importance of that analysis. It argues that a second, largely unexamined dimension of the same event, the decision architecture that governed the response, is equally instructive, and is where the ASR Decision Sovereignty framework offers analysis the conventional asset-restoration account does not.

2. Decision Sovereignty: Who Was Empowered to Decide the Restoration Sequence

Public statements attributed responsibility for the response to the federal Energy Ministry and its officials, who communicated restoration timelines and technical explanations to the public within hours of the collapse. The visible authority structure was therefore federal and ministerial. Restoration itself, however, necessarily proceeded through a technical hierarchy of transmission and distribution companies operating at regional and local level, entities whose relationship to the ministry's public statements and timelines is not clarified in the public record.

This is precisely the gap Decision Sovereignty is designed to make visible. A federal ministry announcing a national restoration timeline is a statement of political and institutional authority; the actual

sequencing of which grid station is restored before another is, in practice, a technical decision taken considerably further down the hierarchy, closer to the fifth paper of this theme's federated structure than to the first paper's single decision authority. The public record does not establish how, or how clearly, standing to make that specific sequencing decision was distributed between federal authority and the regional transmission and distribution companies actually executing restoration.

3. Decision Integrity: The Reliability of the Early Operating Picture

Initial public statements attributed the disturbance to a voltage or frequency fluctuation originating in southern Pakistan, with the specific technical explanation evolving over the following hours and days as internal assessments developed. This pattern, in which the first public account of a cascading grid failure is subsequently revised as more complete technical data becomes available, is common to complex infrastructure incidents generally, and is not unique to this event.

From a Decision Sovereignty perspective, the relevant observation is not that the initial account was later refined, which is expected, but what that evolution implies for decision integrity in the first hours: restoration decisions taken immediately after the collapse were necessarily made against an incomplete and evolving operating picture, before the fuller technical account of the cascade's origin and path had been established. The public record does not indicate what contingency existed for revising early restoration sequencing decisions as that picture improved.

4. Decision Latency: Reconstructing the Response Chain

The public timeline permits a partial reconstruction of the Decision Latency Chain developed in the third paper of this theme. Detection appears to have been rapid: the frequency drop was identified and communicated publicly within a comparatively short window of the event beginning. The Energy Ministry's public communication, offering an initial technical explanation and a restoration estimate, followed within hours, suggesting a short interval across detection, interpretation and the initial public escalation stage.

The more instructive latency, from a decision-sovereignty perspective, sits later in the chain: the considerable variation in restoration time across regions, roughly eight hours for the capital region, approximately double that for the two largest cities, and substantially longer again for some rural areas, indicates that execution latency, not detection or initial authorisation latency, was the dominant factor in how long the disruption's consequences actually lasted for most of the population. This is consistent with this theme's general argument that the weakest transition in the chain, not the chain's total length, determines effective response time, and here that weakest transition appears to sit in execution and communication across a federated technical hierarchy rather than in the initial decision to respond.

5. Autonomous Decision Boundary: Automated Protection Action

Reporting on the incident's technical cause, a frequency disturbance cascading through the transmission network, is consistent with the involvement of automated protection systems, which are a standard feature of modern grid infrastructure and are

designed to isolate sections of a network automatically when frequency or voltage parameters fall outside safe limits, precisely in order to prevent a more extensive cascade. The public record does not detail the specific automated protection actions taken during this event, but the pattern of a fast-moving cascade originating from a localised disturbance is broadly consistent with automated systems acting, as designed, faster than any human authorisation chain could have intervened.

This paper does not have access to the internal technical detail that would allow a full application of the fourth paper's Autonomous Decision Boundary to this specific event. It notes the general point that where automated protection systems acted to isolate parts of the network during the cascade, they were very likely operating within the Automated Action position on that boundary, pre-authorised in advance rather than authorised in the moment, and that a fuller version of this paper, drawing on internal technical data, could assess whether that boundary was clearly defined and whether it held under the specific conditions of this event.

6. Federated Decision Continuity: Independent Restoration Across Regions

The marked difference in restoration times across regions, cities restored within roughly eight to sixteen hours against rural areas requiring one to three days, is consistent with a federated restoration process in which different regional transmission and distribution entities restored their own sections of the network at different rates, rather than a single, centrally sequenced restoration applied uniformly across the country. This is exactly the pattern the fifth paper of this theme's Federated Decision Continuity model is designed to address: independent nodes, here regional grid operators, making and executing their own restoration decisions during a period

when full national coordination was, by the nature of a nationwide collapse, difficult to sustain in real time.

The public record does not indicate whether this regional variation reflected a pre-defined decision scope of the kind this theme recommends, in which each regional operator's independent restoration authority and its limits were established in advance, or whether it reflected an improvised division of effort driven by differing local technical conditions. The distinction matters for resilience design even though both produce broadly similar restoration timelines: a pre-defined federated structure is repeatable and improvable in a way that an improvised one is not.

7. Authority Recovery: Reasserting Normal Governance

Public reporting on the aftermath of the collapse focused principally on the technical cause, the economic cost, and the political scrutiny that followed regarding the government's handling of the event, consistent with prior incidents of this kind in Pakistan. This is, in the terms of the sixth paper of this theme, largely a Verification-stage response: a review, conducted substantially in public and political fora, of what happened and why.

The public record does not indicate whether a formal Authority Reconciliation stage followed, addressing how independently taken regional restoration decisions were integrated into a single national account, or whether Normal Governance Restoration was marked as an explicit, documented event distinct from the simple fact that power had been restored nationally. This absence is itself an instructive data point for this research theme: it suggests that, at least in the public record, operational recovery and authority recovery were not clearly distinguished from one another in the handling of

this incident, exactly the conflation the sixth paper of this theme argues resilience frameworks should avoid.

8. Conclusion

Applied to the public record of the 23 January 2023 national grid collapse, the six Decision Sovereignty concepts developed across this research theme surface a set of questions that a conventional asset-and-restoration account of the incident does not directly address: how clearly restoration authority was distributed between federal and regional actors, how reliable the operating picture available to decision-makers was in the first hours, where in the response chain the greatest delay actually occurred, how automated protection systems' independent action was bounded, how regional restoration variation reflected a designed or improvised federated structure, and whether authority recovery was handled as a distinct task from operational recovery.

This paper has answered these questions only partially, constrained by the limits of the public record. A fuller version, incorporating the internal technical and institutional detail held in the ASR case study library, would allow each of the six concepts to be applied with greater precision, and would give the Decision Sovereignty Framework developed across this theme its first fully worked applied example, grounded in a real, well-documented national event rather than in generic illustration alone.

***"

PART II

The ASR Decision Sovereignty Framework

The Six Pillars, Converged Into One Governance Architecture

The seven chapters that precede this one develop Decision Sovereignty concept by concept: its definition, its theoretical grounding, its measurement, its application to autonomous systems, its extension to federated and disconnected governance, its resolution through authority recovery, and its grounding in a real, documented disruption. This closing part restates those six pillars as a single, converged framework, and translates them into three practical outputs intended for direct use by infrastructure operators, regulators and crisis planners: a Diagnostic identifying where decision flow fails, a Measurement discipline quantifying how much decision latency exists, and an Architecture specifying what governance structure preserves legitimate decision flow under disruption.

Continuity of Flow

A critical system is not truly resilient merely because its physical assets, data and communications remain available. It is resilient when legitimate decisions can still be made, authorised and executed when normal conditions fail.

The thing that must remain continuous under disruption is not merely electricity, data, transport or supply. It is the legitimate flow from information to decision to authorised action, and the six pillars developed across this book are the six points at which that flow can be preserved, or can fail.

Positioning Within the ASR Architecture

The Decision Sovereignty Framework is not a parallel doctrine to Flow Sovereignty. It is the authority layer within it. Flow Sovereignty and the Six Sovereign Layers ask whether resources, information and services can continue to move under disruption. The Decision Sovereignty Framework asks the prior question of who has standing to decide that they should, and whether that decision can reach execution.

Two concepts developed elsewhere in the ASR doctrine are generalised, not duplicated, by this framework. Commander's Intent and Human Integrity Density, developed for the specific case of AI operating within distributed infrastructure, are generalised by the Autonomous Decision Boundary to any autonomous or semi-autonomous actor. Mesh Governance and 72-Hour Autonomy, which describe a node's operational capacity to function while disconnected, are given their authority dimension by Federated Decision Continuity, which describes what that node may legitimately decide while disconnected.

Resilience is not only the ability of infrastructure to continue functioning under disruption; it is the ability of legitimate decisions to continue flowing through the system when normal pathways fail.

Where decision flow breaks, infrastructure resilience eventually becomes operational paralysis.

Three Practical Outputs

The framework converges into three outputs, each drawing on specific pillars above and each intended for direct practitioner use rather than further theoretical elaboration.



Figure 8.1 — The six pillars converging into the ASR Decision Sovereignty Framework and its three practical outputs.

Diagnostic — Where does decision flow fail?

The Diagnostic applies the Authority, Decision and Execution triangle as a structured assessment method, walking a given system, facility or network through each of the six pillars in turn and identifying where its decision architecture is undefined, untested or dependent on a single point of failure.

- Authority mapping: is decision authority for each critical function pre-established, documented and pathway-independent?
- Integrity check: what is the reliability of the operating picture available to decision-makers in the first hours of disruption?
- Boundary review: for each autonomous or automated system, is the Autonomous Decision Boundary explicitly documented across all six of its positions?

- Federation review: for each networked or distributed node, is its Federated Decision Continuity scope defined in advance, with clear activation and expiry conditions?
- Recovery readiness: does an Authority Recovery Protocol exist for every category of emergency, automated or delegated authority the system can activate?

The output of the Diagnostic is not a score but a map: a documented account of which vertices of the Authority, Decision and Execution triangle are secure for a given system, and which are assumed rather than designed.

Measurement — How much decision latency exists?

The Measurement output applies the seven-stage Decision Latency Chain — Detection, Interpretation, Escalation, Authorisation, Decision, Communication, Execution — to a specific system or incident, identifying the weakest transition rather than only the total elapsed time.

- Baseline measurement: decision latency under normal, non-crisis operating conditions, establishing what the chain looks like when nothing is wrong.
- Stress measurement: decision latency under simulated or actual disruption, isolating which of the seven transitions degrades most under stress.
- Authority-transfer measurement: the added latency when a decision must move to an emergency, automated or federated authority.
- Recovery measurement: the latency of the Authority Recovery Protocol itself, from stabilisation through to formal normal governance restoration.

This output remains a measurement discipline rather than a published scoring index until a dedicated technical companion validates weighting and thresholds against back-tested incident data,

following the precedent set by the Net Fragility Matrix and the ASR-VAM Assessor Handbook.

Architecture — What governance structure preserves legitimate decision flow?

The Architecture output is the design specification a system, facility or network should hold before disruption occurs, drawing together the structural elements defined across the six pillars:

- A documented decision authority map, naming who is empowered to decide for each critical function, independent of any single communication or command pathway.
- A documented Autonomous Decision Boundary for every autonomous or automated system, specifying its position across AI Detection, AI Recommendation, Automated Action, Human Authorisation, Institutional Override and Emergency Authority.
- A documented Federated Decision Continuity scope for every distributed or disconnectable node, specifying what it may decide independently, what requires higher authority, and when that delegation activates and expires.
- A documented Authority Recovery Protocol for every category of delegated or automated authority, specifying Stabilisation conditions, Verification responsibility, Authority Reconciliation ownership and the act of Normal Governance Restoration.

Together, these four elements form the Architecture output's core deliverable: a single governance document, specific to a given system or network, that can be produced, tested through disconnection and disruption exercises, and revised, in the same way physical and technical resilience plans already are.

Using the Three Outputs Together

The three outputs are designed to be applied in sequence for a new engagement, and revisited cyclically thereafter. The Diagnostic identifies where a system's decision architecture is weakest. The Measurement discipline quantifies how severe that weakness is in practice, in terms of actual response time under stress. The Architecture output then specifies the governance structure required to close the gap the first two outputs identified. Applied together, they move the Decision Sovereignty Framework from a set of theoretical concepts to a repeatable engagement methodology, comparable in structure to the diagnostic, measurement and design sequence ASR-VAM already established for asset-level vulnerability.

Closing

The ASR Decision Sovereignty Framework converges six interlocking pillars into a single governance architecture, sitting as the authority layer within ASR's existing Flow Sovereignty doctrine. Its purpose is to make an implicit and often untested question explicit: whether legitimate authority to decide survives disruption alongside the infrastructure it governs. The Diagnostic, Measurement and Architecture outputs set out in this part are intended to carry that question out of theory and into practice, giving infrastructure operators, regulators and crisis planners a repeatable means of identifying, measuring and designing against decision fragmentation before it is discovered under crisis conditions.



"Resilience is Architecture, not reaction."

Syed Abid Shah, Founder, Aquarian Systematic Resilience

AI Attribution

Generative artificial intelligence (Claud & OpenAI ChatGPT) was used as a research-support and editorial tool in the preparation of this Book. Its assistance included helping to organise material, improve structure and clarity, synthesise information supplied by the author, and refine the presentation of the analysis.

The substantive ideas, research questions, conceptual framework, interpretation, arguments, and conclusions are those of **Syed Abid Shah**. The author reviewed, revised, and takes responsibility for the final text. AI was not treated as an independent scholarly authority or substitute for the author's judgment.