

# Mandate: A Multi-Agent Collaboration Framework Powered by Magick.AI

Developed by Magick.AI

December 2, 2024

## Abstract

Mandate, powered by Magick.AI, is an advanced framework designed to orchestrate multi-agent workflows, resolve conflicts, and adaptively manage tasks in collaborative AI ecosystems. Built upon the Magick.AI architecture, Mandate leverages the Context Bridge, dynamic memory systems, and multi-agent collaboration tools to ensure seamless execution and resolution across distributed agents. This paper outlines the core architecture, agent interactions, and conflict resolution mechanisms while demonstrating how Magick.AI's tools transform traditional task orchestration into an adaptive and intelligent system.

## 1 Introduction

Magick.AI is a next-generation platform for building, managing, and deploying collaborative AI agents. The **Mandate framework** extends Magick.AI's core features, introducing a structured methodology for defining objectives, assigning roles, and resolving conflicts dynamically.

By integrating Magick.AI's **Context Bridge**, **memory systems**, and **multi-agent collaboration tools**, Mandate enables agents to act as specialized contributors, efficiently solving complex tasks while adhering to conflict resolution and task dependencies.

This document outlines the Mandate framework in detail, focusing on:

- System architecture and its interaction with Magick.AI agents.
- Built-in conflict resolution mechanisms and their applications.
- Strategies to mitigate failure in complex multi-agent systems.
- Use cases and future research directions for adaptive collaboration.

## 2 System Architecture and Components

The Mandate framework integrates the foundational principles of Magick.AI with advanced multi-agent collaboration features. Its architecture is designed to streamline task management, role assignment, and conflict resolution while ensuring adaptability to dynamic workflows.

## 2.1 Mandate

The **Mandate** serves as the overarching directive, encapsulating the following:

- **Objective:** Defines the overall goal or problem the system must address.
- **Agent Cast:** Specifies all Magick.AI agents involved, assigning roles for each task.
- **Tasks:** Breaks down the mandate into actionable units, each managed independently.
- **Global CRDM:** A default Conflict Resolution and Decision-Making framework that governs dispute resolution across tasks unless a task-specific framework is defined.
- **Constraints:** Establishes system-wide parameters such as deadlines, resource limits, or safety requirements.
- **Audit Log:** Maintains a detailed record of all interactions, decisions, and escalations for accountability and improvement.

The Mandate interacts directly with Magick.AI's **Context Bridge**, which:

- Dynamically injects real-time data into workflows to ensure agents remain contextually aware.
- Ensures consistency between external systems (e.g., APIs) and internal agent processes.
- Enables continuous memory synchronization, critical for multi-tasking agents.

## 2.2 Tasks and RACI Roles

Tasks are the building blocks of the Mandate. Each task defines its own structure, dependencies, and rules, including the utilization of **RACI roles**. RACI stands for:

- **Responsible:** The agent(s) who execute the task and ensure its completion.
- **Accountable:** The Orchestrator Agent overseeing the task, ensuring the outcome meets Mandate objectives.
- **Consulted:** Agents providing expertise or input during task execution.
- **Informed:** Agents updated on task progress but not directly involved.

### Role Utilization in Mandate

- **Task Management:** The Orchestrator Agent assigns RACI roles to agents based on their expertise and the complexity of the task.
- **Conflict Resolution:** RACI roles guide the CRDM pathway:
  - Responsible agents vote or escalate conflicts.
  - Consulted agents provide SME vetoes or additional validation.

– Accountable agents arbitrate unresolved issues.

- **Transparency and Accountability:** Each role is logged in the audit trail, ensuring clarity in decision-making processes and escalation paths.

Tasks are designed to be modular, allowing the Mandate to scale by defining dependencies between tasks. For example:

- Task A must be validated by the Responsible agent before Task B can begin.
- Consulted agents in Task C can override Responsible agents if SME veto conditions are met.

## 2.3 Conflict Resolution and Decision-Making Frameworks (CRDMs)

Each task in Mandate is governed by a CRDM, which determines how disputes are resolved and decisions are made. These frameworks include:

1. **Majority Voting:** Suitable for low-stakes decisions where all agents have equal input.
2. **Weighted Voting:** Assigns weights to agent votes based on expertise or role priority.
3. **SME Veto:** Allows Consulted agents with subject matter expertise to override decisions.
4. **Arbitration by Orchestrator:** Enables the Accountable agent to make the final decision in unresolved disputes.
5. **Consensus Decision-Making:** Requires unanimous agreement among all agents, ideal for high-stakes tasks.

These CRDMs leverage Magick.AIs **memory systems** to store past decisions and outcomes, enabling agents to:

- Adapt CRDM logic based on historical context.
- Identify patterns in task outcomes to improve decision-making efficiency.

## 2.4 Agent Memory and Context Bridge Integration

Magick.AIs memory systems enable agents to operate contextually and collaboratively:

- **Short-Term Memory:** Retains task-specific data for active workflows.
- **Long-Term Memory:** Stores historical interactions, allowing agents to learn from past tasks.
- **Episodic Memory:** Captures event-specific data to facilitate escalation and arbitration.

The **Context Bridge** ensures real-time synchronization of:

- Task dependencies across agents.
- Dynamic updates from external systems, such as APIs or environmental data streams.

## 2.5 Audit Trail and Transparency

The audit trail is central to the Mandate framework, logging:

- Agent decisions and their rationale.
- Conflicts, escalations, and CRDM pathways used.
- Task outcomes and any deviations from planned objectives.

This transparency fosters continuous improvement and mitigates hindsight bias during post-incident analyses.

## 3 Workflow Functionality and Failure Mitigation

The Mandate framework orchestrates a dynamic workflow that integrates RACI roles, CRDMs, and Magick.AIs core systems (memory and Context Bridge). This ensures seamless task execution, real-time decision-making, and proactive failure mitigation. The following describes how the system operates end-to-end while interrelating key concepts.

### 3.1 End-to-End Workflow

The Mandate workflow is divided into the following stages:

#### 1. Mandate Initialization:

- The Architect Agent defines the overarching goal of the Mandate.
- Tasks are created with associated dependencies, constraints, and RACI roles.
- The Global CRDM is assigned, providing a default fallback for all tasks.

#### 2. Task Delegation and Execution:

- Orchestrator Agents assign Responsible, Consulted, and Informed agents for each task.
- Responsible agents execute the task using memory systems to reference historical data and adapt to real-time inputs.
- Consulted agents validate intermediate outputs or provide SME expertise when invoked by CRDM rules.

#### 3. Conflict Resolution and Escalation:

- When disputes arise, the task-specific CRDM governs resolution.
- Unresolved conflicts are escalated to the Orchestrator or Architect Agent via the escalation pathway.

#### 4. Completion and Audit:

- Checklist validation is performed by Responsible agents, with oversight from Consulted agents if necessary.
- All decisions, escalations, and outcomes are logged in the Audit Trail for future reference.

## 3.2 Callbacks to RACI Roles and CRDMs

The RACI framework interrelates closely with CRDMs to ensure clear decision-making at every step:

- **Responsible agents** actively execute tasks, invoking CRDM logic when conflicts arise.
- **Consulted agents** provide SME vetoes during critical decision-making stages, ensuring task-specific expertise is prioritized.
- **Accountable agents (Orchestrators)** resolve escalations when lower-level CRDMs fail to achieve consensus.
- **Informed agents** remain updated on task outcomes, ensuring cross-task synchronization and reducing redundant work.

### Example: RACI Roles in Weighted Voting CRDM

In a high-stakes task, the weighted voting CRDM uses:

- **Responsible agents** to provide initial inputs weighted equally.
- **Consulted agents** to override the majority if SME expertise detects a flaw.
- **Orchestrator agents** to arbitrate remaining disputes using historical task performance data.

## 3.3 Failure Mitigation in Mandate

The Mandate framework incorporates strategies to mitigate common failure points in complex systems:

### 1. Proactive Failure Monitoring:

- The Context Bridge detects anomalies in real-time data streams.
- Memory systems identify latent errors from proto-failures, flagging potential risks before escalation.

### 2. Decoupling Task Dependencies:

- Tasks are designed with minimal coupling to prevent cascading failures.
- Fallback CRDMs at the task level isolate failures from affecting the global system.

### 3. Hierarchical Escalation Pathways:

- The escalation system prevents governance overload by resolving conflicts locally whenever possible.
- Only critical issues reach Architect Agents, ensuring scalability across multiple mandates.

### 4. Transparency and Accountability:

- The audit trail captures every decision and escalation, enabling post-incident analysis without hindsight bias.
- Transparent logs empower operators to refine CRDM logic iteratively.

### 3.4 Callbacks to Magick.AI Features

Magick.AIs tools are central to Mandates failure mitigation:

- **\*\*Memory Systems\*\***:
  - Short-term memory ensures consistent task context for all agents during execution.
  - Long-term memory provides historical data for CRDM adaptation and anomaly detection.
- **\*\*Context Bridge\*\***:
  - Real-time synchronization enables agents to adapt to external changes dynamically.
  - Continuous context updates ensure that all agents maintain a shared understanding of the task environment.

### 3.5 Example: Failure Mitigation in Action

**Scenario: Interdependent Tasks with Resource Conflict**

- **Setup**: Task A involves resource allocation, while Task B depends on Task A's completion.
- **Failure Point**: Task A's Responsible agent makes an allocation error, creating a resource shortfall for Task B.
- **Mitigation Pathway**:
  1. Task A's Responsible agent flags the issue via the Context Bridge.
  2. Task B's Orchestrator Agent invokes the Weighted Voting CRDM, incorporating SME inputs to reallocate resources dynamically.
  3. The final decision is logged in the audit trail, ensuring transparency for post-task review.

This interconnected workflow demonstrates how Mandate combines RACI roles, CRDMs, and Magick.AIs capabilities to deliver both adaptability and robustness.

## 4 Conflict Resolution Frameworks (CRDMs)

At the heart of Mandate's ability to handle complexity is its array of **Conflict Resolution and Decision-Making Frameworks (CRDMs)**. These frameworks are designed to dynamically resolve disputes, adapt to the stakes of each task, and prevent systemic failure by leveraging both agent capabilities and the Magick.AI infrastructure.

## 4.1 Overview of CRDMs

Each CRDM is tailored to specific scenarios, with predefined rules for handling disagreements among agents. CRDMs are designed with the following principles:

- **Adaptability:** CRDMs dynamically adjust to the context of the task, using data from the Context Bridge and agent memory systems.
- **Resilience:** Multi-layered decision-making ensures failures in one layer are mitigated by fallbacks.
- **Scalability:** CRDMs handle increasing complexity by distributing decision-making across agents.

The CRDMs implemented in Mandate include:

### 4.1.1 Majority Voting

- **Description:** Each agent casts one vote, and the majority decision prevails.
- **Use Case:** Suitable for low-stakes tasks where all agents have equal input weight.
- **Workflow:**
  1. Each Responsible agent submits a decision.
  2. Votes are tallied, and the majority outcome is implemented.
- **Limitations:** Inappropriate for high-stakes tasks requiring SME validation.

### 4.1.2 Weighted Voting

- **Description:** Assigns weights to agent votes based on expertise, priority, or historical accuracy.
- **Use Case:** Critical tasks where SME inputs must outweigh general votes.
- **Workflow:**
  1. Each agent's vote is assigned a weight stored in the long-term memory system.
  2. The decision is calculated using:

$$\text{Outcome} = \arg \max_{d \in D} \left( \sum_{a \in A} w_a \cdot v_{a,d} \right)$$

3. The Context Bridge synchronizes weights in real-time to account for evolving expertise.

#### 4.1.3 SME Veto

- **Description:** Grants veto authority to Subject Matter Experts (SMEs) for decisions within their domain of expertise.
- **Use Case:** High-risk tasks requiring specialized knowledge to prevent catastrophic outcomes.
- **Workflow:**
  1. Responsible agents submit their decisions.
  2. Consulted agents (SMEs) review the majority decision.
  3. If the majority decision is flawed, the SME vetoes it and proposes an alternative.
- **Strength:** Ensures domain-specific risks are mitigated by expert oversight.
- **Limitations:** Overuse of vetoes may slow progress.

#### 4.1.4 Consensus Decision-Making

- **Description:** Requires unanimous agreement among agents before proceeding.
- **Use Case:** Tasks with extreme consequences, where dissent must be fully resolved.
- **Workflow:**
  1. All Responsible agents propose a decision.
  2. Iterative discussions resolve disagreements, leveraging the Context Bridge for shared understanding.
  3. Task execution begins only when unanimous agreement is achieved.
- **Limitations:** Time-intensive for large agent casts.

#### 4.1.5 Arbitration by Orchestrator

- **Description:** The Orchestrator Agent acts as the final arbitrator when other CRDMs fail to resolve disputes.
- **Use Case:** Tasks requiring fast decisions or hierarchical oversight.
- **Workflow:**
  1. Responsible agents escalate unresolved conflicts to the Orchestrator.
  2. The Orchestrator evaluates the inputs and historical context.
  3. A final decision is issued and logged for audit purposes.



#### 4.1.6 Hierarchical Escalation

- **Description:** Escalates conflicts through a hierarchy of agents, up to the Architect if necessary.
- **Use Case:** Critical system-level tasks where localized resolution is insufficient.
- **Workflow:**
  1. Task conflicts unresolved by CRDMs are escalated to the Orchestrator.
  2. If the Orchestrator fails to resolve the issue, the Architect Agent intervenes.
  3. The Architect uses audit logs and Context Bridge data to finalize the decision.

#### 4.1.7 Hybrid CRDMs

- **Description:** Combines multiple CRDMs for complex tasks requiring flexibility.
- **Use Case:** Tasks involving diverse agent casts or evolving conditions.
- **Workflow:**
  1. Start with Majority or Weighted Voting.
  2. Trigger SME Veto or Consensus mechanisms if disagreements persist.
  3. Escalate to Arbitration or Hierarchical Escalation as a fallback.
- **Strength:** Balances efficiency and resilience.

### 4.2 Callbacks to RACI and Context Bridge

CRDMs operate seamlessly within Mandate's RACI framework:

- **Responsible agents** initiate decision-making.
- **Consulted agents** apply SME Veto or validate outputs.
- **Orchestrators (Accountable agents)** arbitrate disputes and manage hybrid workflows.

The **Context Bridge** enables real-time CRDM adaptation:

- Dynamically synchronizes external data for Weighted Voting.
- Provides historical insights from memory systems for Arbitration decisions.
- Flags unresolved disagreements, triggering fallback CRDM layers.

## 4.3 CRDM Failure Mitigation in Action

### Scenario: Task Chain Deadlock with High Stakes

- **Setup:** Task A involves resource allocation, but a dispute arises between Responsible agents.
- **Failure Point:** SME Veto is invoked but fails to reach consensus.
- **CRDM Pathway:**
  1. Weighted Voting adjusts priorities to resolve partial disagreements.
  2. Orchestrator arbitration resolves the remaining ambiguity using historical task outcomes.
  3. Final escalation to the Architect Agent ensures alignment with the Mandate objective.
- **Outcome:** The conflict is resolved with minimal escalation, and the process is logged for future optimization.

## 5 Agent Prompt Examples and Task Execution

Mandate leverages Magick.AIs dynamic memory systems, Context Bridge, and IoT integrations to generate tailored prompts for agents. These prompts are dynamically assembled based on the state of the mandate, task checklists, and agent roles (RACI framework).

### 5.1 Dynamic Prompts Based on Agent Roles

The following illustrates how agent prompts differ across roles (Responsible, Accountable, Consulted, Informed) during a task’s lifecycle.

#### 5.1.1 Scenario: Energy Grid Optimization

**Objective:** Balance power distribution across a smart grid during peak demand, integrating IoT sensor data and historical outage reports.

**Mandate State and Checklist:**

- Task: Redistribute power from over-supplied regions to high-demand areas.
- Checklist:
  1. Analyze IoT sensor data for real-time grid status.
  2. Retrieve historical outage patterns from long-term memory.
  3. Prioritize critical zones (e.g., hospitals, data centers).
  4. Validate the proposed redistribution plan with SME input.
  5. Execute the redistribution and log all decisions.

**Dynamic Prompts:**

1. **\*\*Responsible Agent Prompt\*\***:

Task: Redistribute power in the smart grid during peak demand.

IoT Data: [Region A: 20% overload, Region B: 30% underutilized, ...]

Historical Context: "Region A has experienced recurring outages under similar

Checklist Progress:

1. IoT data analyzed: Complete.
2. Historical patterns retrieved: Complete.
3. Critical zones identified: Pending.

Instructions:

1. Identify and prioritize critical zones (e.g., hospitals, data centers).
2. Draft a redistribution plan to mitigate overloads and ensure continuity.

2. **\*\*Accountable Agent (Orchestrator) Prompt\*\***:

Task: Oversee power redistribution in the smart grid.

Current Status:

- Responsible agent is drafting a redistribution plan.
- Consulted agents are validating historical context and SME input.

Critical Dependencies:

- Redistribution must prioritize hospitals and critical zones.
- All decisions must be logged for audit review.

Instructions:

- Coordinate between Responsible and Consulted agents.
- Finalize the redistribution plan and ensure execution readiness.

3. **\*\*Consulted Agent (SME) Prompt\*\***:

Task: Validate the power redistribution plan for Region A.

Current Status:

- Proposed redistribution: Transfer 15% capacity from Region B to Region A.
- Context: Region A has experienced recurring failures due to transformer issues.

SME-Specific Data:

- Transformer ratings for Region A: 80% max capacity.
- Real-time IoT readings: 20% overload.

Instructions:

- Assess whether the redistribution plan addresses overload risks effectively.
- Suggest adjustments if required.

4. **\*\*Informed Agent Prompt\*\***:

Task: Power redistribution in progress for Region A and Region B.

Updates:

- Redistribution plan has been drafted and is under SME review.
- Critical zones identified: Hospitals in Region A are top priority.

Instructions:

- Monitor updates for final redistribution decisions.
- No direct action required.

### 5.1.2 Output Analysis and RACI Feedback

- **Responsible Agent Output:** Drafts a plan to transfer 15
- **Consulted Agent (SME) Feedback:** Suggests lowering the transfer to 10
- **Accountable Agent Action:** Incorporates SME feedback and authorizes the revised plan for execution.
- **Informed Agents:** Updated on final decisions, ensuring alignment across dependent tasks.

## 5.2 Complex Scenario: Disaster Response Coordination

**Objective:** Coordinate rescue efforts after a natural disaster using IoT sensors, drone feeds, and historical data.

### Setup

- Task: Prioritize rescue zones based on severity and accessibility.
- IoT Data: Sensor readings from affected areas (e.g., population density, environmental hazards).
- Context: Drone imagery shows blocked roads and collapsed buildings in key zones.
- Checklist:
  1. Analyze IoT and drone data for rescue zone prioritization.
  2. Retrieve historical data on disaster response times and bottlenecks.
  3. Allocate resources (e.g., rescue teams, medical supplies) to top-priority zones.
  4. Coordinate logistics to ensure resource delivery.

### Dynamic Execution with Code Integration

Below is Python code to demonstrate how Magick.AI dynamically handles task completion:

```
from magick_ai import ContextBridge, AgentMemory, IoTIntegration

# Initialize agents and systems
context = ContextBridge()
memory = AgentMemory()
iot_data = IoTIntegration()

# Fetch IoT and historical data
```

```

iot_readings = iot_data.get_sensor_data(["Zone A", "Zone B", "Zone C"])
historical_data = memory.fetch("disaster_response_bottlenecks")

# Responsible agent generates rescue priorities
responsible_agent_prompt = f"""
IoT Data: {iot_readings}
Historical Data: {historical_data}
Checklist: Analyze data and draft rescue priorities.
"""
responsible_output = context.query_agent("ResponsibleAgent", responsible_agent_prompt)

# SME reviews priorities and provides feedback
sme_prompt = f"""
Proposed Priorities: {responsible_output}
Historical Failures: {historical_data}
"""
sme_feedback = context.query_agent("SMEAgent", sme_prompt)

# Orchestrator finalizes the plan
orchestrator_prompt = f"""
Rescue Priorities: {responsible_output}
SME Feedback: {sme_feedback}
Checklist: Validate and execute the rescue plan.
"""
orchestrator_decision = context.query_agent("OrchestratorAgent", orchestrator_prompt)

# Execute and log
context.log("Final Rescue Plan", orchestrator_decision)

```

## Outputs and Impact

- **\*\*Responsible Agent Output\*\***: Drafts initial rescue priorities based on IoT and historical data.
- **\*\*SME Feedback\*\***: Refines priorities by identifying logistical bottlenecks.
- **\*\*Orchestrator Decision\*\***: Combines SME feedback to authorize final resource allocation.
- **\*\*Outcome\*\***: Resources are deployed efficiently, minimizing delays and ensuring maximum impact.

## 6 Failure Mitigation and System Resilience

While the Mandate framework is designed to be adaptive and resilient, its critical to recognize that no system is entirely free of failure. Complex systems like Mandate must embrace the concept of **\*\*failures as part of the design\*\***, managing these failures through redundancy, transparency, and proactive strategies. The following sections summarize key failure mitigation techniques.

## 6.1 Proactive Monitoring and Error Detection

The Context Bridge and Agent Memory systems are integrated to continuously monitor task status and detect potential failures before they propagate through the system:

- **Anomaly Detection:**
  - Continuous scanning of IoT sensor data, agent feedback, and task progress to identify anomalies.
  - Early warning signals, such as unexpected task delays or inconsistent agent outputs, trigger proactive reviews.
- **Task Redundancy:**
  - Tasks are designed with built-in redundancy. If one agent fails or produces inconsistent results, other agents can step in with alternative approaches.
  - This redundancy is controlled through RACI roles, ensuring that critical functions are not entirely dependent on any single agent.

## 6.2 Escalation Paths and Dynamic Adjustments

The Mandate system builds multiple layers of **escalation** to ensure that disputes or task failures are handled at the appropriate level, preventing unnecessary delays or errors from cascading:

- **Hierarchical Escalation:**
  - If lower-level agents (e.g., Responsible or Consulted) cannot resolve a conflict, the issue is escalated to the Orchestrator or Architect Agent.
  - This hierarchical structure ensures that complex disputes receive attention from higher-level agents who have the broader context to resolve them.
- **Fallback CRDM Layers:**
  - In cases where CRDMs fail to reach consensus, fallback mechanisms like SME Veto, Arbitration by Orchestrator, and Hybrid CRDMs activate, ensuring that resolution continues without stagnating.

## 6.3 Continuous Improvement Through Audit and Feedback

Mandate incorporates a comprehensive **Audit Trail** that logs all agent decisions, escalations, and resolutions, making post-task analysis and failure identification more effective:

- **Audit Logs:**
  - Every decision made by agents, from the Responsible to the Architect, is logged with contextual details.
  - Logs are used for post-incident analysis, where failures are reviewed, and corrective actions are proposed for future workflows.
- **Feedback Loops:**

- Agent feedback, especially from Consulted and Informed roles, provides continuous learning for the system, improving the accuracy and speed of future decision-making.
- Over time, these logs enable agents to dynamically adapt to new challenges and optimize workflows.

## 6.4 Real-Time Decision Adjustments and Memory Synchronization

A core feature of Mandate’s failure mitigation is the synchronization of task-relevant data across agents using **Magick.AI’s Memory Systems**. These systems ensure that all agents have access to up-to-date information, preventing decisions based on outdated or incomplete data:

- **Memory Systems Integration:**

- Tasks that evolve over time or involve fluctuating data (e.g., environmental conditions, resource availability) are continuously updated in the system’s memory.
- Each agent can dynamically access this data through the Context Bridge, ensuring that decisions are based on the most recent inputs.

- **Contextual Decision-Making:**

- Agents receive **contextual data** (e.g., updated sensor information, historical performance) tailored to their specific role in the task. This ensures that each agent makes decisions based on their relevant expertise and real-time task conditions.

## 6.5 Leveraging IoT Data for Proactive Adaptation

In addition to memory and context, Mandate integrates **IoT data** for real-time decision-making and task management. This ensures that the system is not only reactive but also anticipatory in handling failures:

- **IoT Sensor Data Integration:**

- Data from IoT devices (e.g., environmental sensors, resource tracking) is fed into the Context Bridge, dynamically adjusting tasks based on real-world conditions.
- Tasks like disaster response or energy grid management rely heavily on IoT data to adjust resource allocation and priorities as new information arrives.

- **Predictive Failure Detection:**

- By analyzing historical and real-time IoT data, Mandate can predict potential system failures before they occur, triggering preventative actions like resource reallocation or task redefinition.

## 6.6 Mitigating Complex Failure Scenarios with Hybrid CRDMs

In some scenarios, multiple CRDMs are combined to handle tasks that involve conflicting data sources or highly complex dependencies:

- **Hybrid CRDMs** combine:
  - Majority Voting for general decisions.
  - SME Veto for specialized expertise.
  - Hierarchical Escalation if the previous layers fail.
- This approach provides **flexibility** while ensuring that decisions are made by the most appropriate agents and mechanisms available, preventing bottlenecks.