



Collaborative Governance in AI-Enabled Resilient Cloud Networks: A Case-Based Framework from Ventilator Production

Hossein Gholami Ghadi ^a

^a Department of Mechanical, Industrial and Mechatronics Engineering, Toronto Metropolitan University, Toronto, Canada.

ARTICLE INFO

Received: 2026/03/15

Revised: 2026/04/27

Accept: 2026/05/29

Keywords:

*Collaborative governance;
Socio-technical systems;
Cloud networks; AI
governance; Resilience;
Public crisis response*

ABSTRACT

Public crises require rapid coordination across institutions, sectors, and digital infrastructures. The COVID-19 pandemic showed that emergency production cannot be treated only as a technical or industrial problem; it is also a governance challenge involving public agencies, healthcare organizations, universities, military units, manufacturers, logistics providers, regulators, and digital platforms. This study develops a case-based conceptual framework for understanding AI-enabled cloud manufacturing as a socio-technical governance mechanism for emergency production networks. Using emergency ventilator production during COVID-19 as an illustrative case, the framework explains how distributed institutional capacity can be transformed into collective crisis-response capability through public-value-oriented demand interpretation, institutional capacity mapping, AI-enabled coordination, accountability structures, traceability, ethical allocation, digital resilience, and fallback governance. The study contributes to humanities and social science by reframing service composition as institutional role allocation and by translating technical resilience concepts into governance categories such as adaptive governance, institutional redundancy, accountability, public trust, and fair allocation. The framework emphasizes that AI-enabled crisis coordination should remain human-supervised, transparent, auditable, and embedded in legitimate public authority. Future research should empirically examine emergency production networks through interviews, comparative case studies, and governance-performance indicators.

^a Corresponding author email address: Hosein.gholami72@gmail.com (Hossein Gholami Ghadi).

DOI: <https://doi.org/10.22034/ijshsc.v1i2.208>

Available online: 2026/05/31

Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0>).

3115-7297/BGSA Ltd.

1. Introduction

Public crises reveal the limitations of institutions that are organized for routine conditions but must suddenly respond to exceptional social, operational, and ethical demands. The COVID-19 pandemic showed that crisis response is not only a matter of medical capacity or industrial production; it is also a governance challenge involving public agencies, healthcare organizations, universities, military units, manufacturers, logistics providers, digital platforms, and communities. When urgent needs exceed ordinary institutional capacity, crisis response depends on the ability to coordinate heterogeneous actors, share information, allocate scarce resources, preserve accountability, and maintain public trust under uncertainty.

Emergency ventilator production during COVID-19 provides a clear example of this governance problem. Ventilators were not merely technical medical devices; they became scarce public resources linked to life, institutional responsibility, and social legitimacy. During the crisis, conventional production channels were often insufficient to meet sudden demand. As a result, alternative capacities from non-traditional actors, including universities, military organizations, industrial firms, and research groups, became relevant to emergency response. Prior research on cloud manufacturing during COVID-19 has shown that unused capacities from other supply networks can be mobilized to support ventilator production through service composition and redundancy-based coordination [11]. From a social-science perspective, this raises a broader question: how should distributed institutional capacities be governed when emergency production becomes a collective public-response function?

Cloud manufacturing offers one possible coordination logic for this problem. In technical terms, cloud manufacturing transforms distributed production resources into cloud-accessible services that can be searched, matched, composed, and coordinated. In governance terms, however, it can be interpreted as a socio-technical infrastructure for organizing institutional collaboration. It creates a digital environment in which different organizations can contribute partial capacities to a shared emergency objective. Such a system does not only allocate tasks; it also redistributes roles, responsibilities, decision rights, and accountability across a network of public and private actors.

This distinction is important because emergency production networks are not neutral technical systems. Decisions about which organization produces which component, which hospital or region receives scarce outputs, which data are shared, and which fallback procedures are activated are

deeply social and institutional. These decisions affect public value, fairness, legitimacy, and trust. A technically efficient allocation may still be socially unacceptable if it lacks transparency, excludes vulnerable regions, relies on unaccountable decision systems, or fails to document responsibility for production quality. Therefore, AI-enabled cloud manufacturing should be examined not only as an optimization mechanism but also as a governance arrangement.

Existing research on resilient cloud manufacturing and intelligent service composition has made important technical contributions. Redundancy-based models have shown how alternative capacities can support production bounce-back during disruptions [11]. Robust service composition has addressed uncertainty in cloud manufacturing networks through resilience-oriented allocation models [12]. Reinforcement learning has been used to support adaptive service composition in stochastic environments where disruptions evolve over time [13]. Generative AI, quantum-enhanced allocation, cyber-resilient migration, and fallback orchestration have further extended the technical capacity of digital production and cloud-service systems [14-18].

However, these studies have primarily framed the problem as one of manufacturing resilience, computational optimization, cyber-resilience, or service allocation. Limited attention has been given to the socio-technical governance implications of AI-enabled emergency production networks. In particular, the literature has not sufficiently examined how such systems coordinate heterogeneous institutions, distribute authority, preserve accountability, support ethical allocation, maintain traceability, and sustain trust during public crises. This gap matters because crisis-response systems cannot be evaluated only by efficiency, cost, or service availability. They must also be evaluated by their legitimacy, transparency, fairness, institutional robustness, and ability to support public value.

The COVID-19 ventilator case illustrates why this gap is significant. Activating universities, military workshops, industrial firms, and logistics providers for emergency production can expand capacity, but it also creates governance questions. Who determines whether a non-traditional provider is qualified? Who authorizes task allocation? Who is accountable if a component fails? How are scarce outputs distributed across hospitals or regions? How can AI-generated recommendations be audited? How should the system respond when the digital coordination platform fails or is compromised? These questions cannot be answered by technical models alone. They require a governance framework that integrates institutional coordination, human oversight, ethical principles, and socio-technical resilience.

Accordingly, this study develops a case-based conceptual framework for understanding AI-enabled cloud manufacturing as a socio-technical governance mechanism for public crisis response. The framework uses emergency ventilator production during COVID-19 as an illustrative case to explain how distributed institutional capacity can be transformed into a coordinated emergency production network. It links cloud manufacturing and intelligent service composition with collaborative governance, accountability structures, traceability, ethical allocation, digital resilience, and public trust.

This study addresses four research questions. First, how can AI-enabled cloud manufacturing support collaborative governance in emergency production networks during public crises? Second, what institutional roles, coordination mechanisms, and accountability structures are required to transform distributed production capacity into a collective crisis-response capability? Third, how can intelligent service composition balance efficiency, resilience, ethical allocation, traceability, and public trust in emergency production networks? Fourth, what lessons can be derived from COVID-19 ventilator production for the governance of future socio-technical crisis-response systems?

The study makes four contributions to humanities and social science. First, it reframes AI-enabled cloud manufacturing from a technical service-composition problem into a socio-technical governance mechanism for public crisis response. Second, it develops a conceptual framework linking distributed production capacity, institutional coordination, AI-enabled decision support, accountability, and public-value-oriented allocation. Third, it interprets COVID-19 ventilator production as an example of cross-sector collaborative governance involving public agencies, healthcare organizations, universities, military units, industrial firms, logistics providers, and digital platforms. Fourth, it identifies ethical and governance challenges associated with AI-enabled crisis production, including transparency, traceability, institutional trust, human oversight, and fair allocation of scarce resources.

The remainder of this paper is organized as follows. Section 2 reviews the literature on collaborative governance, socio-technical systems, emergency production networks, institutional resilience, AI-enabled decision support, and accountability in public crisis response. Section 3 presents the proposed socio-technical governance framework. Section 4 applies the framework to the illustrative case of emergency ventilator production during COVID-19. Section 5 discusses

governance, ethical, and social implications. Section 6 concludes the study and identifies directions for future research.

2. Literature Review

2.1. Public crisis response and collaborative governance

Public crises challenge the capacity of institutions to act under uncertainty, urgency, resource scarcity, and public scrutiny. In routine conditions, public agencies, hospitals, manufacturers, universities, logistics providers, and private firms usually operate according to separate mandates and organizational boundaries. During a major crisis, however, these boundaries become insufficient. Problems such as emergency medical-equipment shortages, disrupted supply chains, cyber threats, and sudden demand surges require coordinated action across sectors. This makes crisis response a problem of collaborative governance rather than a purely administrative, technical, or industrial task.

Collaborative governance refers to arrangements in which multiple public and private actors jointly participate in decision-making, coordination, and implementation around a shared public problem [1,2]. It is especially relevant when no single actor possesses all the authority, information, resources, or legitimacy required to solve the problem alone. Public crises intensify this need because resources are dispersed across institutions, while the demand for response is concentrated and urgent. In the case of COVID-19 ventilator production, hospitals needed devices, governments needed emergency response capacity, manufacturers had partial production capabilities, universities and research laboratories had technical expertise, military organizations had logistical and operational capacities, and digital platforms could support coordination. The governance challenge was to convert these distributed capacities into collective action.

Emergency production networks therefore require more than operational efficiency. They require role definition, authority allocation, information sharing, accountability, conflict resolution, and public-value alignment [3,19]. If a university laboratory contributes technical design, an industrial firm produces components, a military workshop supports fabrication, and a hospital receives the final device, the system must clarify who validates quality, who approves task allocation, who holds responsibility for failure, and how decisions are documented. Without such governance structures, distributed production may increase capacity but also create ambiguity, mistrust, and accountability gaps.

The COVID-19 ventilator case illustrates the importance of collaborative governance. Prior cloud manufacturing research showed that unused capacities from military organizations, universities, and other supply networks could support ventilator production during the pandemic [11]. From a technical perspective, this can be described as service composition and redundancy. From a social-science perspective, it is a case of cross-sector crisis coordination. The same production task becomes a governance object: it must be assigned to qualified actors, monitored through digital systems, evaluated against public priorities, and justified to stakeholders.

A collaborative governance lens also highlights the difference between capacity and capability. A society may possess dispersed technical capacity, but this does not automatically become crisis-response capability. Capacity becomes capability only when institutions can coordinate, exchange reliable information, trust one another's roles, and act under a shared decision framework. AI-enabled cloud manufacturing can support this transformation, but only if it is embedded in governance arrangements that define authority, transparency, accountability, and ethical allocation.

2.2. Socio-technical systems and digital coordination

AI-enabled emergency production networks should be understood as socio-technical systems. A socio-technical system combines technological infrastructure with human actors, institutional rules, organizational routines, and social values. In such systems, technical design and social organization are interdependent. A digital platform may optimize task allocation, but the meaning and legitimacy of that allocation depend on institutional authority, stakeholder trust, regulatory approval, and public accountability. Therefore, cloud manufacturing cannot be analyzed only as a computational architecture; it must also be examined as a coordination infrastructure embedded in social and institutional contexts.

Cloud manufacturing provides a digital coordination environment in which distributed resources can be represented as services, matched to tasks, and combined into production plans. Technically, this supports resource sharing, service selection, and dynamic allocation. Socially, it changes how organizations participate in crisis response. It allows actors that do not normally belong to the same supply chain to become part of a temporary emergency production network. This shift creates new forms of interdependence between public agencies, private manufacturers, universities, military organizations, digital platform operators, and healthcare institutions.

Digital coordination can increase responsiveness, but it can also create new vulnerabilities. First, the accuracy of decisions depends on the quality of shared data. If capacity data, hospital demand signals, logistics information, or certification records are incomplete or outdated, the allocation system may produce recommendations that appear efficient but are institutionally or operationally inappropriate. Second, platform dependency can create concentration risk. If the digital coordination platform fails, is attacked, or becomes inaccessible, the emergency production network may lose visibility and decision continuity. Third, algorithmic coordination may obscure responsibility. If an AI system recommends assigning a critical task to a specific provider, stakeholders must still know who authorized the assignment and who remains accountable for the outcome.

Prior work on self-healing LLM-DBMS pipelines shows that digital systems operating under volatility require fallback mechanisms such as Degrade, Substitute, and Bypass to preserve continuity [14]. In a social-science interpretation, these fallback tiers are not merely technical recovery mechanisms. They represent governance rules for decision continuity under system failure. Degrade can mean using simplified decision rules when full information is unavailable. Substitute can mean activating an alternative authority, provider, database, or decision-support system. Bypass can mean relying on pre-approved emergency templates when real-time systems cannot be trusted. This interpretation connects digital resilience to institutional continuity.

Cyber-resilience is also central to socio-technical crisis governance. A cyberattack against a coordination platform can disrupt more than software performance; it can disrupt public response, institutional trust, and emergency production capacity. Research on generative AI-based service composition under DDoS and ransomware attacks demonstrates that cloud-fog systems require anomaly detection, migration, and resilience mechanisms to maintain service quality under adversarial conditions [15]. In a public crisis-response setting, this means that digital public infrastructure must be designed not only for efficiency but also for trustworthiness, continuity, auditability, and controlled degradation.

2.3. Emergency production networks and institutional resilience

Emergency production networks are temporary or semi-permanent arrangements that mobilize distributed capacities to produce critical goods during crises. Unlike conventional supply chains, they often involve actors that do not routinely collaborate. They may also operate under compressed timelines, uncertain regulations, limited information, and changing public priorities.

This makes them institutionally complex. Their performance depends not only on technical allocation but also on how well institutions coordinate roles, share authority, and maintain legitimacy.

Institutional resilience refers to the ability of organizations and governance arrangements to maintain core functions, adapt to disruption, and recover from shocks [4]. In emergency production, institutional resilience requires backup capacity, alternative organizational pathways, flexible decision rules, and the ability to learn from changing conditions. Redundancy is therefore not only a technical principle; it is also an institutional principle. A resilient crisis-response system should avoid excessive dependence on a single supplier, platform, organization, or decision authority.

The concept of anti-affinity in resilient cloud service composition provides a useful analogy. In technical systems, anti-affinity rules prevent critical tasks from being placed on resources that may fail together. Prior quantum reinforcement learning research used anti-affinity constraints to reduce single points of failure in cloud service composition [16]. In institutional terms, anti-affinity means that critical public-response functions should not be concentrated within one fragile institution, one data platform, one supplier, or one geographic cluster. This logic is important for crisis governance because institutional concentration can create systemic vulnerability.

Robust service composition also has a social-science interpretation. Prior research on robust cloud manufacturing service networks used uncertainty-aware service composition and subentropy to manage disruption scenarios [12]. Technically, this addresses uncertain parameters in service allocation. Institutionally, it reflects governance under uncertainty. In crises, decision makers rarely have complete information. Demand forecasts may be uncertain, supplier capacities may change, regulatory conditions may shift, and public expectations may evolve. Robust governance does not require perfect information; it requires decision structures that remain legitimate and functional despite incomplete information.

Reinforcement learning adds another useful governance analogy: adaptive learning. Prior work on resilient cloud networks used reinforcement learning to support service composition in stochastic environments [13]. In governance terms, crisis-response systems must learn from feedback. If an allocation decision causes delay, bottleneck, inequity, or quality risk, the system should revise its decision logic. Adaptive governance requires monitoring, feedback loops, and willingness to update procedures as crisis conditions evolve. Therefore, AI-enabled allocation can

be interpreted as a technical representation of a broader institutional need: the ability to adapt while preserving accountability.

2.4. AI-enabled decision support, ethics, and accountability

AI-enabled decision support can improve the speed and scope of crisis coordination. In emergency production networks, AI systems can identify candidate providers, generate allocation options, predict bottlenecks, detect anomalies, and recommend fallback actions. These capabilities are valuable when decision makers face urgent demand and dispersed information. However, AI-enabled crisis response also creates ethical and accountability challenges. Technical recommendations can affect access to scarce resources, institutional responsibilities, and public trust. Therefore, AI systems in crisis governance must be evaluated not only by performance but also by legitimacy [6-8].

One ethical issue is transparency. Stakeholders must understand why a system recommends one allocation over another. If a ventilator component is assigned to a particular provider, the decision should be explainable in terms of capacity, certification, urgency, risk, and public value. A black-box recommendation may increase speed but reduce trust. In crisis conditions, where decisions may be contested, transparency supports legitimacy.

A second issue is accountability. AI systems can recommend, rank, or generate options, but they should not dissolve human responsibility. Public agencies, regulators, and authorized decision makers must remain accountable for final decisions [5]. This is especially important in emergency production, where quality failures can affect safety and public confidence. The governance framework must therefore specify who approves AI-generated allocations, who validates providers, who monitors compliance, and who is responsible for correcting errors.

A third issue is traceability. Emergency production networks involve multiple actors and components. Traceability ensures that each decision, component, provider, and approval can be documented. In technical terms, traceability supports quality control. In social terms, it supports accountability and institutional learning. If a problem occurs, the system must be able to reconstruct how the decision was made, which data were used, and which actors were responsible.

Generative AI intensifies these governance concerns. Prior diffusion model-based generative optimization research shows that generative models can produce disruption-aware allocation candidates for cloud manufacturing [17]. From a governance perspective, this means AI can generate policy-relevant options during a crisis. However, option generation is not the same as

legitimate decision-making. Machine-generated options must be filtered through public values, regulatory requirements, ethical allocation rules, and human oversight. The key question is not only whether AI can generate feasible options, but who determines which options are acceptable.

Quantum-enhanced and quantum reinforcement learning methods raise similar future-oriented questions. Studies on quantum machine learning and quantum reinforcement learning for resilient cloud networks suggest that advanced computational methods may improve exploration of large allocation spaces and support service migration under disruption [16,18]. For social science, the significance of these methods lies less in the technical details and more in the governance questions they raise: How explainable are advanced decision-support systems? How should public institutions evaluate recommendations produced by complex models? What standards of evidence are required before such systems can influence crisis decisions?

2.5. Public value, trust, and fair allocation

Emergency production networks operate in public-value contexts. Their purpose is not only to produce outputs efficiently but also to support collective welfare, protect vulnerable populations, and preserve trust in public institutions [10]. Public value includes effectiveness, fairness, transparency, safety, accountability, and responsiveness. A crisis-response system that maximizes output but distributes resources unfairly may fail socially even if it succeeds technically.

Fair allocation is particularly important when emergency outputs are scarce. During COVID-19, ventilators became ethically significant because scarcity could affect life-and-death decisions. Although production networks do not directly make bedside allocation decisions, they influence which hospitals, regions, or populations receive scarce resources first. Therefore, emergency production governance should include explicit allocation principles. These may consider urgency, shortage severity, population vulnerability, regional equity, and public health priorities.

Trust is also essential. Trust affects whether institutions share data, whether providers accept task assignments, whether hospitals rely on emergency-produced devices, and whether the public accepts crisis decisions. AI-enabled systems can support trust when they are transparent, auditable, and human-supervised. They can undermine trust when they appear opaque, biased, unsafe, or unaccountable. Therefore, the governance of AI-enabled emergency production must include mechanisms for explanation, appeal, audit, and oversight.

Traceability contributes to trust by making actions visible and reviewable. In a distributed production network, each component and decision should be linked to a responsible actor. This is

especially important when non-traditional providers are mobilized. Universities, military organizations, and industrial firms may contribute valuable capacity, but their roles must be documented and bounded by qualification criteria. Public trust depends on the perception that emergency expansion of capacity does not mean abandonment of safety or accountability.

2.6. Research gap and synthesis

The reviewed literature suggests that AI-enabled cloud manufacturing and resilient service composition have substantial relevance for crisis response, but their social-science implications remain underdeveloped. Technical studies have shown how service composition, robust optimization, reinforcement learning, generative allocation, quantum-enhanced search, cyber-resilient migration, and fallback orchestration can improve resilience in disrupted cloud and manufacturing systems. However, these mechanisms have not been sufficiently translated into a governance framework for public crisis response.

This gap has three dimensions. First, prior technical research has focused mainly on allocation performance, service quality, cost, resilience metrics, and computational efficiency. Less attention has been paid to institutional coordination, public legitimacy, and accountability. Second, crisis governance research has emphasized coordination, collaboration, and public management, but has not fully engaged with AI-enabled cloud manufacturing as a digital infrastructure for emergency production. Third, AI ethics and accountability literature has addressed transparency, fairness, and human oversight, but often without connecting these principles to concrete emergency production networks.

The present study addresses this gap by developing a socio-technical governance framework for AI-enabled emergency production networks. The framework interprets technical concepts such as service composition, redundancy, robustness, reinforcement learning, generative allocation, anti-affinity, cyber-resilience, and fallback orchestration through social-science categories such as institutional coordination, distributed authority, adaptive governance, accountability, trust, and fair allocation. Using emergency ventilator production during COVID-19 as an illustrative case, the study explains how distributed capacity can become a collective crisis-response capability when it is governed through transparent, accountable, and public-value-oriented mechanisms.

3. Proposed Socio-Technical Governance Framework

3.1. Framework overview

This study proposes a socio-technical governance framework for AI-enabled emergency production networks. The framework explains how distributed institutional capacities can be

transformed into a coordinated crisis-response capability through digital coordination, intelligent service composition, accountability structures, and public-value-oriented decision rules. It is designed for crisis contexts in which no single organization has sufficient capacity, authority, information, or legitimacy to act alone.

The framework is built on the premise that emergency production is simultaneously a technical, institutional, and ethical process. Technically, emergency production requires decomposition of urgent needs into tasks, identification of available services, and allocation of tasks to qualified actors. Institutionally, it requires coordination across public agencies, hospitals, universities, military organizations, private manufacturers, logistics providers, and digital platform operators. Ethically, it requires transparent decision-making, fair allocation, traceability, and human oversight when outputs are scarce or socially consequential.

The framework consists of six interrelated layers: the public-value and crisis-demand layer, the distributed institutional capacity layer, the AI-enabled coordination layer, the accountability and traceability layer, the ethical allocation and public trust layer, and the resilience and fallback governance layer. These layers are not sequential in a narrow technical sense; rather, they represent governance functions that must operate together during a crisis.

Socio-technical governance framework for AI-enabled emergency production

A public crisis-response view of cloud manufacturing, service composition, accountability, and fallback governance

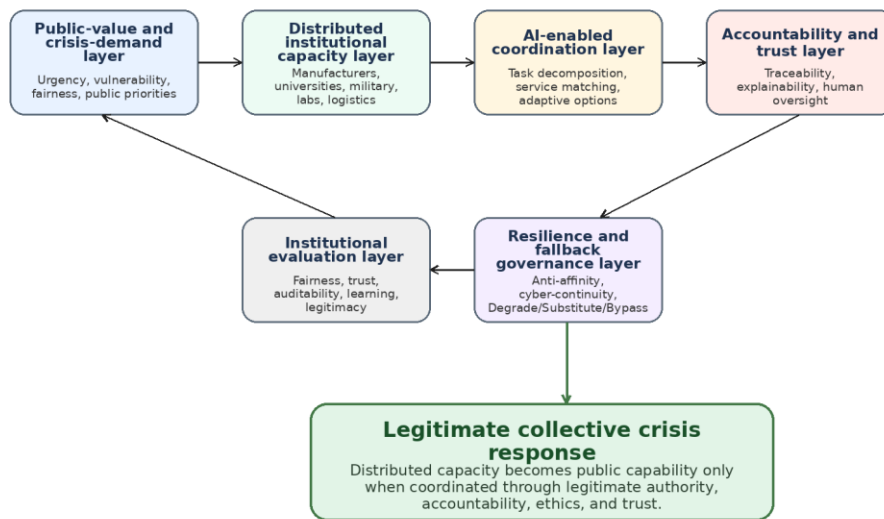


Figure 1. Socio-technical governance framework for AI-enabled emergency production networks.

3.2. Public-value and crisis-demand layer

The public-value and crisis-demand layer defines the social purpose of the emergency production network. In a public crisis, production is not guided only by market demand or organizational self-

interest. It is guided by public value: protection of life, continuity of essential services, fairness, legitimacy, and collective welfare. Therefore, the first governance task is to translate crisis needs into public-value-oriented production priorities.

In the COVID-19 ventilator case, the demand for ventilators was not merely a product requirement; it represented an urgent public need tied to intensive-care capacity, institutional responsibility, and social trust. Hospitals, public health agencies, and emergency planners generated demand signals, but those signals required interpretation. A high number of requested ventilators in one region may indicate clinical urgency, poor inventory, population vulnerability, or data-reporting differences. Therefore, the governance system must distinguish between raw demand and justified priority.

This layer asks several questions. Which public need is the production network trying to address? Who has authority to define urgency? Which indicators should determine priority? How should competing claims from different hospitals, regions, or populations be compared? How can the system avoid privileging actors with stronger political, economic, or digital visibility over actors with greater vulnerability?

The output of this layer is a crisis-demand profile. Unlike a conventional demand forecast, this profile includes public-value criteria. It specifies not only what should be produced, but why, for whom, under what urgency, and according to which allocation principles. For emergency production networks, this is essential because socially legitimate production requires more than output maximization.

3.3. Distributed institutional capacity layer

The distributed institutional capacity layer identifies and organizes the actors that can contribute to crisis production. In ordinary conditions, these actors may belong to separate institutional fields. Universities may focus on research, military units on defense operations, manufacturers on commercial production, hospitals on care delivery, and logistics firms on distribution. During crises, however, these boundaries can become more permeable. Cross-sector coordination becomes necessary when public needs exceed ordinary institutional capacity.

Cloud manufacturing provides a technical logic for representing these capacities as services. However, the social-science significance is broader: it makes latent institutional capacity visible and actionable. Prior work on production bounce-back in cloud manufacturing showed that unused capacities from other supply networks, including military organizations and university research

groups, could support ventilator production during the COVID-19 pandemic [11]. In governance terms, this means that emergency production requires institutional mapping before allocation can occur.

This layer therefore classifies actors according to their crisis-response roles. Public agencies define emergency mandates, legal authority, and public priorities. Hospitals and healthcare organizations provide demand signals and user requirements. Manufacturers provide production capacity. Universities and laboratories provide technical expertise, prototyping, testing, and validation support. Military organizations may provide disciplined logistics, fabrication capacity, and emergency coordination. Logistics providers connect production outputs to points of need. Digital platform operators provide the information infrastructure for coordination.

However, capacity is not equivalent to qualification. A provider may have technical ability but lack certification, traceability, cyber readiness, or accountability structures. Therefore, this layer must distinguish between available capacity, qualified capacity, and governable capacity. Available capacity refers to what an organization can physically do. Qualified capacity refers to what it can do according to safety, quality, and technical standards. Governable capacity refers to what can be integrated into a transparent, accountable, and auditable crisis-response network.

The output of this layer is an institutional capacity registry. This registry should include organizational role, technical capability, certification status, geographic location, capacity limits, data-sharing readiness, cyber-dependence, quality-control capability, and accountability contact points. Without such a registry, emergency coordination becomes improvised and may reproduce inequality, confusion, and delay.

3.4. AI-enabled coordination layer

The AI-enabled coordination layer explains how intelligent systems can support task allocation, coordination, and adaptation across distributed institutions. In technical terms, this layer corresponds to service composition: decomposing a crisis need into tasks and matching those tasks with available services. In governance terms, it is a mechanism for distributing responsibility across actors.

For emergency ventilator production, tasks may include component fabrication, electronics assembly, quality testing, calibration, packaging, transportation, and inventory allocation. An AI-enabled coordination system can identify which actors are candidates for each task, compare constraints, and recommend allocation options. Robust service composition is relevant here

because crisis conditions involve uncertainty. Prior research developed a robust service composition model incorporating subentropy and validated it through a ventilator-production case during COVID-19 [12]. In the present framework, such technical methods are interpreted as tools for governance under uncertainty.

The coordination layer should not be understood as an autonomous decision authority. AI can generate options, rank alternatives, detect bottlenecks, and support reallocation, but it should not replace institutional judgment. Public agencies, regulators, and authorized crisis managers must define the rules under which AI recommendations are considered acceptable. This includes constraints related to fairness, safety, certification, regional priority, and accountability.

The layer should also support adaptive governance. In a crisis, the system state changes continuously. A supplier may become unavailable, a hospital may report urgent shortage, a logistics route may fail, or a digital platform may experience disruption. Adaptive decision support can help revise allocations in response to feedback. Reinforcement-learning-based service composition is technically relevant to such dynamic settings [13], but the governance interpretation is broader: crisis-response systems must learn, update, and correct decisions without losing accountability.

The output of this layer is not a single optimal allocation. It is a set of explainable, auditable, and governable coordination options. Each option should show which actor performs which task, what data supported the recommendation, what constraints were considered, and what public-value criteria were applied. This shifts the role of AI from hidden optimizer to accountable decision-support mechanism.

3.5. Accountability and traceability layer

The accountability and traceability layer addresses one of the most important social risks of distributed crisis production: responsibility can become fragmented. When multiple institutions jointly produce a scarce and critical good, responsibility for quality, delay, failure, and distribution can become unclear. This is especially problematic when AI-enabled systems influence task allocation.

Accountability requires that decision rights and responsibilities be explicit. The framework distinguishes five accountability roles. The mandate holder defines the public objective and has authority to activate the emergency production network. The data steward is responsible for the quality, timeliness, and security of shared data. The allocation authority approves or rejects AI-

generated coordination options. The service provider performs assigned tasks and is responsible for task-level compliance. The audit authority reviews decisions, traces failures, and supports post-crisis learning.

Traceability supports accountability by making actions reconstructable. Each task assignment should be linked to the data used, the decision rule applied, the approving authority, the responsible provider, and the quality-control record. In technical systems, provenance mechanisms and transaction guards can support safe reuse and prevent unsafe actions. The event-driven fallback orchestration study, for example, used provenance-aware caching and transaction guards to block unsafe database actions while maintaining service continuity [14]. In governance terms, these mechanisms correspond to auditability and procedural safeguards.

This layer also defines explainability requirements. AI recommendations should be interpretable enough for public decision makers to understand why an allocation was suggested. Explainability does not require exposing every computational detail; it requires providing decision-relevant reasons. For example, an allocation recommendation should identify whether a provider was selected because of capacity, certification, proximity, resilience, prior performance, or equity priority.

The output of this layer is an accountability map. This map specifies who is responsible for demand validation, data provision, AI recommendation review, task approval, production compliance, quality verification, distribution decisions, and post-crisis audit. Without such a map, AI-enabled emergency production risks becoming technically efficient but institutionally unaccountable.

3.6. Ethical allocation and public trust layer

The ethical allocation and public trust layer ensures that emergency production networks serve public legitimacy, not merely operational performance. In crises, production outputs may be scarce, time-sensitive, and politically visible. Decisions about where outputs go, which providers are activated, and which risks are tolerated can affect trust in public institutions.

This layer introduces fairness, transparency, and human oversight as governance requirements. Fairness means that scarce outputs should be allocated according to justified public criteria rather than institutional power, visibility, or market access. Transparency means that stakeholders can understand the principles behind allocation and production decisions. Human

oversight means that AI-generated recommendations remain subject to review by authorized and accountable decision makers.

Trust is produced through repeated evidence that the system is competent, fair, and accountable. Competence requires reliable coordination and quality outputs. Fairness requires justified prioritization. Accountability requires traceability and correction mechanisms. If any of these conditions fails, the system may lose legitimacy even if it produces outputs efficiently.

This layer is particularly important for non-traditional providers. Universities, military organizations, and industrial firms can contribute valuable capacity, but their inclusion may raise questions about safety, mandate, authority, and oversight. The governance system must therefore make clear why these actors are included, what they are allowed to do, what standards they must meet, and who supervises their contribution.

The output of this layer is a public-trust protocol. This protocol defines ethical allocation principles, transparency practices, human oversight mechanisms, complaint or appeal pathways, and communication responsibilities. It helps ensure that the emergency production network remains socially legitimate.

3.7. Resilience and fallback governance layer

The resilience and fallback governance layer defines how the emergency production network continues functioning when institutions, platforms, providers, or data flows fail. In technical systems, resilience is often measured through continuity, availability, latency, or recovery. In socio-technical governance, resilience also includes continuity of authority, accountability, and legitimacy.

This layer adapts three fallback concepts: Degrade, Substitute, and Bypass. In the technical literature, these tiers are used to preserve service continuity when a primary system becomes unavailable or unreliable [14]. In governance terms, Degrade means using simpler but approved decision rules when full optimization or complete data are unavailable. Substitute means activating an alternative provider, authority, model, or coordination pathway. Bypass means using pre-approved emergency templates when real-time decision support cannot be trusted.

Fallback governance is essential because crises rarely unfold under ideal information conditions. A platform may fail, a provider may withdraw, a data source may become unreliable, or an AI model may produce recommendations that violate policy constraints. The framework therefore requires fallback rules to be defined before the crisis. These rules should identify which

functions can be simplified, which actors can substitute for others, and which pre-approved templates can be activated under severe disruption.

The concept of anti-affinity can also be translated into governance design. Technically, anti-affinity prevents critical tasks from being co-located within the same failure domain; this reduces the risk of simultaneous failure [16]. Institutionally, anti-affinity means avoiding concentration of critical crisis-response functions in one organization, one region, one supplier, one digital platform, or one authority. This reduces dependence on a single point of institutional failure.

The output of this layer is a fallback governance plan. This plan specifies degraded decision rules, substitute actors, emergency templates, authority transfer conditions, communication protocols, and audit requirements during failure. It ensures that the system can continue operating without abandoning accountability.

3.8. Institutional evaluation layer

The institutional evaluation layer defines how the governance framework should be assessed. Since the framework is designed for humanities and social science, evaluation should not be limited to technical performance. It should include social, institutional, ethical, and governance outcomes.

The first evaluation criterion is coordination effectiveness. This measures whether the system successfully connects crisis demand with distributed institutional capacity. The second criterion is accountability clarity. This evaluates whether decision rights and responsibilities are explicit. The third criterion is traceability. This assesses whether decisions, data, providers, and outputs can be audited. The fourth criterion is fairness. This evaluates whether allocation principles are transparent and justified. The fifth criterion is institutional resilience. This assesses whether the system can continue operating when actors or platforms fail.

Additional criteria include trust, public legitimacy, data governance quality, cyber-continuity, human oversight, and learning capacity. Learning capacity is especially important because crisis-response systems should improve over time. Post-crisis audits should identify which decisions worked, which coordination failures occurred, and how future governance rules should change.

The output of this layer is a governance evaluation dashboard. This dashboard should not be understood as a purely technical dashboard. It should combine operational indicators with governance indicators. For example, it may include production completion rate, provider diversity,

decision approval time, number of fallback activations, audit completeness, equity of distribution, stakeholder complaints, and correction actions.

3.9. Framework synthesis

The proposed framework reframes AI-enabled cloud manufacturing as a governance mechanism rather than only a production technology. It shows that emergency production networks require more than digital platforms and allocation algorithms. They require public-value criteria, institutional capacity mapping, accountable AI-supported coordination, traceability, ethical allocation, fallback governance, and evaluation of trust and legitimacy.

The framework also clarifies the relationship between technical and social concepts. Service composition becomes institutional role allocation. Redundancy becomes backup institutional capacity. Robustness becomes governance under uncertainty. Reinforcement learning becomes adaptive governance. Generative allocation becomes AI-generated policy options. Anti-affinity becomes avoidance of single points of institutional failure. Cyber-resilience becomes trustworthy digital public infrastructure. Fallback orchestration becomes continuity of public decision-making. Traceability becomes accountability. Equity becomes legitimacy.

This translation is the main conceptual contribution of the framework. It allows technical advances in cloud manufacturing and AI-enabled allocation to be interpreted through social-science categories relevant to crisis governance.

Table 1. Translation of technical concepts into governance concepts.

Technical concept	Social-science reinterpretation	Governance question
Service composition	Institutional role allocation	Which organization should perform which crisis-response task?
Redundancy	Backup institutional capacity	How can dependency on one actor be avoided?
Robustness	Governance under uncertainty	How should decisions remain legitimate when information is incomplete?

Reinforcement learning	Adaptive governance	How should crisis response learn from feedback?
Generative allocation	AI-generated policy options	Who validates machine-generated options?
Anti-affinity	Distributed institutional risk	How can single points of institutional failure be avoided?
Cyber-resilience	Trustworthy digital public infrastructure	How can coordination continue under digital disruption?
Fallback orchestration	Continuity of public decision-making	What happens when the main decision system fails?
Traceability	Accountability	Who is responsible for each production decision?
Equity	Public legitimacy	How are scarce outputs allocated fairly?

Table 2. Governance layers and their functions.

Governance layer	Main function	Key actors	Expected outcome
Public-value and crisis-demand layer	Defines public priorities	Public agencies, hospitals, communities	Crisis-demand profile
	Maps qualified capacity		

Distributed institutional capacity layer		Manufacturers, universities, military, logistics	Institutional capacity registry
AI-enabled coordination layer	Supports task allocation	Digital platforms, public agencies, providers	Explainable coordination options
Accountability and traceability layer	Assigns responsibility	Mandate holders, auditors, providers	Accountability map
Ethical allocation and trust layer	Preserves legitimacy	Public agencies, regulators, affected institutions	Public-trust protocol
Resilience and fallback governance layer	Maintains continuity under failure	Backup providers, platform operators, public authorities	Fallback governance plan

4. Case Illustration: COVID-19 Ventilator Production as a Socio-Technical Governance Problem

4.1. Why ventilator production is a governance case

Emergency ventilator production during COVID-19 provides a useful illustrative case for examining AI-enabled emergency production networks as socio-technical governance systems. Ventilators became highly visible public resources during the pandemic because they were directly connected to intensive-care capacity, public anxiety, institutional preparedness, and the legitimacy of emergency response. Their scarcity was not only a technical production problem; it also raised questions about public authority, cross-sector coordination, ethical allocation, institutional trust, and accountability.

The ventilator case is particularly relevant because it involved multiple types of actors. Hospitals and healthcare systems generated urgent demand signals. Public agencies interpreted these signals and defined emergency priorities. Manufacturers and suppliers provided production and component capacity. Universities and research laboratories could contribute design, testing,

prototyping, and technical expertise. Military organizations could contribute disciplined logistics, fabrication resources, and operational coordination. Digital platforms could support service matching, monitoring, information sharing, and task allocation. The challenge was to transform these distributed capacities into a coordinated public-response network.

From a cloud manufacturing perspective, this can be described as service composition: decomposing ventilator production into tasks and matching those tasks with qualified services. From a social-science perspective, however, the same process is better understood as a form of collaborative governance. Each task allocation also assigns responsibility, redistributes authority, and creates accountability relationships among institutions. Therefore, the ventilator case helps explain why AI-enabled emergency production networks must be governed, not merely optimized.

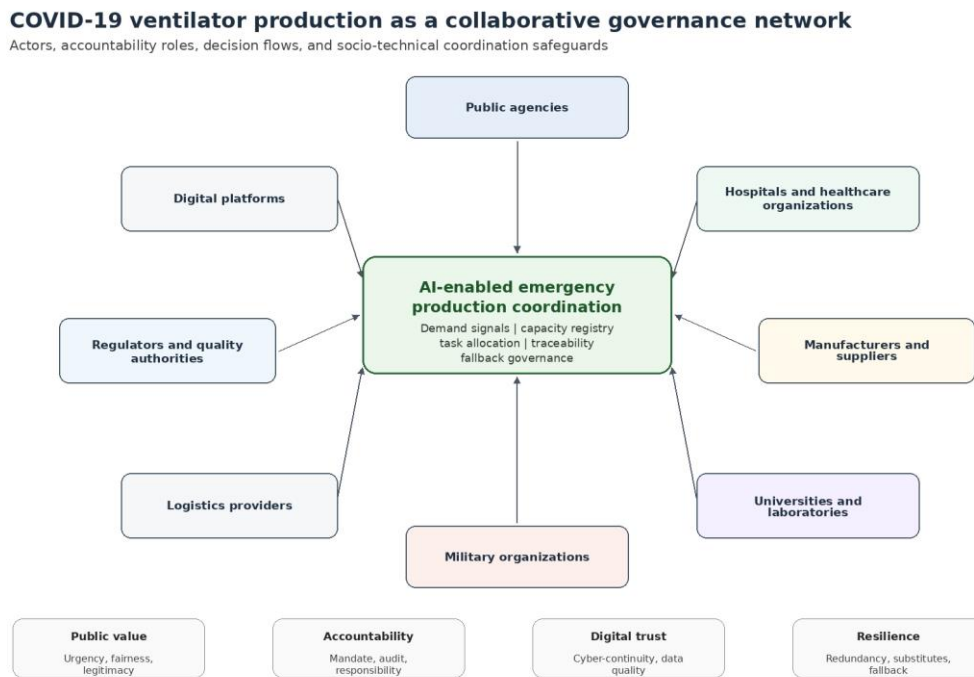


Figure 2. COVID-19 ventilator production as a collaborative governance network.

4.2. Institutional actors and crisis roles

The proposed framework identifies several institutional roles in the ventilator production case. Public agencies act as mandate holders. They define the emergency objective, authorize the activation of non-traditional production capacity, set public priorities, and coordinate among healthcare, industry, and regulatory actors. Their role is not simply to request more ventilators; it is to create a legitimate governance structure for emergency production.

Hospitals and healthcare organizations act as demand signal providers and end users. They identify shortages, communicate urgency, define clinical requirements, and evaluate usability. Their role is essential because production decisions should be guided by real crisis demand rather than abstract production targets. Manufacturers act as production service providers. Some may be conventional medical-device manufacturers; others may be industrial firms with convertible capacity. Their role depends on technical capability, certification status, quality-control systems, and ability to operate under emergency timelines.

Universities and research laboratories act as knowledge and validation actors. They may support prototyping, engineering assessment, design adaptation, testing, and technical problem solving. Their role is particularly important when ordinary supply channels cannot respond quickly enough. Military organizations act as operational support actors. They may provide logistics, emergency coordination, fabrication capacity, disciplined execution, and crisis-management infrastructure. Their inclusion can increase response capacity, but it also requires clear authority boundaries and accountability rules.

Digital platform operators act as coordination infrastructure providers. They maintain the systems through which demand, capacity, allocation, quality, and logistics information are shared. Their role is critical because the emergency production network depends on digital visibility and continuity. Regulators and quality authorities act as safety and legitimacy guardians. They define minimum standards, approve emergency pathways, verify provider qualifications, and ensure that accelerated production does not undermine public safety.

Together, these actors form an emergency production network. The network is not only a supply chain; it is a temporary governance arrangement in which public and private actors share responsibility for a socially consequential output.

4.3. From distributed capacity to coordinated response

The first governance challenge in the ventilator case is converting distributed capacity into coordinated response capability. A society may have many actors with useful resources, but these resources do not automatically become an emergency production system. Coordination requires institutional mapping, qualification, information sharing, decision rules, and trusted authority.

The process begins with demand interpretation. Hospitals may report shortages, but those reports must be translated into public-value-oriented priorities. A shortage in one hospital may require urgent action if it is connected to ICU pressure, regional vulnerability, or lack of substitute

capacity. Therefore, the crisis-demand profile should include not only quantity but also urgency, justification, location, equity considerations, and public health priority.

The next step is capacity mapping. Potential providers must be identified and classified. This includes what they can produce, how quickly they can produce it, what standards they can meet, what data they can share, and what risks they face. Capacity mapping is a governance activity because it determines which actors become visible and eligible in the emergency response system.

The third step is qualification. Not all available capacity is appropriate for emergency ventilator production. Providers must be assessed for technical capability, safety compliance, traceability, cyber readiness, and accountability. This prevents the emergency network from becoming an uncontrolled collection of improvised actors.

The fourth step is task allocation. AI-enabled service composition can recommend which provider should perform which task, but final authority should remain human-supervised. Allocation decisions should be explainable and auditable. A recommendation should show why a provider was selected, which constraints were considered, and which public-value criteria were applied.

The fifth step is monitoring and adaptation. Crisis conditions change. Providers may fail, demand may rise, logistics may be disrupted, or digital systems may degrade. The governance framework must therefore include feedback loops and reallocation procedures. Adaptive coordination is not only a technical need; it is an institutional requirement for maintaining legitimacy under uncertainty.

4.4. Accountability in the ventilator production network

Accountability is central in distributed emergency production because responsibility can easily become fragmented. If a ventilator component is produced by an industrial firm, tested by a laboratory, assembled by another provider, and distributed through a public agency, accountability must be clearly structured.

The framework assigns accountability across several roles. The mandate holder is responsible for activating the emergency production network and defining public priorities. The data steward is responsible for the quality and security of shared information. The allocation authority is responsible for approving AI-supported task assignments. The service provider is responsible for performing assigned tasks according to quality and safety requirements. The audit authority is responsible for reviewing decisions and tracing failures.

In the ventilator case, accountability should cover at least five questions. Who authorized the inclusion of a non-traditional provider? Who approved the allocation of a production task? Who verified that the output met safety requirements? Who decided where completed ventilators should be distributed? Who is responsible if a failure occurs?

These questions are not secondary administrative details. They are central to public legitimacy. A crisis-response system that cannot answer them may produce outputs but still fail as a governance system. Therefore, AI-enabled cloud manufacturing should include accountability mapping as part of its design.

4.5. Ethical allocation and public legitimacy

Ventilator production during COVID-19 also illustrates the ethical dimension of emergency production networks. When outputs are scarce, production and distribution decisions become socially sensitive. The system must decide not only how to produce more ventilators, but also how to prioritize limited outputs across hospitals, regions, or populations.

The proposed framework requires ethical allocation principles to be defined before allocation decisions are made. These principles may include urgency, shortage severity, ICU pressure, regional vulnerability, population risk, and fairness across jurisdictions. Without explicit principles, allocation may be influenced by institutional power, visibility, political pressure, or data availability rather than justified public need.

AI-enabled coordination can support ethical allocation by making trade-offs visible. For example, an allocation dashboard could show how different distribution options affect regions with different levels of shortage. However, AI cannot itself determine what is ethically legitimate. Public authorities and accountable decision makers must define the values embedded in allocation rules.

Public legitimacy depends on the perception that decisions are fair, transparent, and accountable. If universities, military units, or industrial firms are activated for emergency production, the public should be able to understand why they were selected, what standards they must meet, and how their contributions are supervised. Trust is strengthened when emergency expansion of capacity is accompanied by clear oversight.

4.6. Digital coordination and cyber-governance

The ventilator case also shows that emergency production depends on digital coordination. A cloud manufacturing platform would need to integrate demand data, provider capacity, task

decomposition, allocation logic, certification records, logistics information, and monitoring signals. This digital infrastructure creates visibility, but it also creates dependency.

If the platform fails, suffers cyberattack, or produces unreliable recommendations, the production network may lose coordination capacity. Therefore, cyber-governance is part of crisis governance. The system must define who maintains the digital platform, who validates data, who can access sensitive information, who monitors anomalies, and what happens when the platform cannot be trusted.

Fallback governance is especially important. A Degrade rule may allow decision makers to use simplified allocation criteria if full optimization is unavailable. A Substitute rule may activate an alternative platform, provider registry, or coordination authority. A Bypass rule may use pre-approved emergency templates if real-time digital coordination fails. These fallback rules should be defined before the crisis, because improvising them under pressure can undermine trust and accountability.

In this interpretation, technical fallback mechanisms become public governance safeguards. They ensure that the system can continue operating without abandoning oversight, traceability, or fairness.

4.7. Institutional resilience lessons from the ventilator case

The COVID-19 ventilator case provides several lessons for the governance of future emergency production networks. First, distributed capacity must be mapped before a crisis. Emergency response is slower and less accountable when capacity is discovered reactively. Public agencies should maintain registries of qualified or potentially convertible capacities across sectors.

Second, cross-sector collaboration requires pre-defined roles. Universities, military units, private firms, hospitals, and public agencies can all contribute, but their responsibilities must be explicit. Collaboration without role clarity can produce confusion and accountability gaps. Third, AI-enabled coordination should be explainable and human-supervised. AI can support service matching and reallocation, but public decision makers must remain responsible for final decisions.

Fourth, ethical allocation must be embedded into production governance. Emergency production networks should not only maximize output; they should support fair and publicly justified distribution. Fifth, digital resilience is institutional resilience. If the digital coordination layer fails, the institutional response can fail. Therefore, cyber-continuity, fallback rules, and auditability are governance requirements.

Sixth, traceability should be treated as a public trust mechanism. Each component, task, provider, and decision should be documented so that the system can learn, correct errors, and maintain legitimacy.

4.8. Case synthesis

The ventilator production case demonstrates that emergency production networks are not merely technical systems for producing scarce goods. They are socio-technical governance arrangements that coordinate distributed institutions under public pressure. Cloud manufacturing and intelligent service composition can support such coordination, but they must be embedded in governance structures that define public purpose, institutional roles, accountability, ethical allocation, traceability, and fallback procedures.

The main lesson is that AI-enabled emergency production should not be evaluated only by speed, cost, or output. It should also be evaluated by whether it strengthens collaborative governance, preserves public trust, avoids single points of institutional failure, and supports fair crisis response. In this sense, the COVID-19 ventilator case provides a foundation for designing future socio-technical crisis-response systems.

5. Discussion: Governance, Ethics, and Social Implications

5.1. Emergency production as collaborative governance rather than technical mobilization

The proposed framework suggests that emergency production should not be understood merely as the rapid mobilization of technical capacity. It should be understood as a collaborative governance process through which multiple institutions coordinate under uncertainty to address a shared public problem. This distinction is important because crisis production involves public authority, institutional trust, ethical allocation, and accountability, not only production efficiency.

The COVID-19 ventilator case demonstrates this point. The need for ventilators was generated by a public health emergency, but the response required coordination beyond the healthcare sector. Universities, military organizations, industrial firms, logistics providers, manufacturers, regulators, hospitals, and public agencies could all contribute different forms of capacity. However, their participation required rules for inclusion, task allocation, quality assurance, responsibility, and public justification. A purely technical view may describe this as resource allocation. A governance view recognizes that each allocation also defines institutional roles and public responsibilities.

This means that AI-enabled cloud manufacturing should be treated as a governance infrastructure. It can help identify who can do what, where capacity exists, and how tasks can be

distributed. Yet the system also creates new questions: who controls the platform, who validates the data, who approves AI-generated recommendations, who monitors performance, and who explains decisions to stakeholders. Therefore, emergency production networks require institutional design, not only technological design.

5.2. From service composition to institutional role allocation

One of the central arguments of this study is that service composition can be reinterpreted as institutional role allocation. In technical literature, service composition refers to assigning tasks to services based on criteria such as capacity, cost, time, availability, and quality. In a socio-technical crisis-response setting, the same process determines which institution is responsible for which part of a public emergency response.

This reinterpretation changes the meaning of optimization. The best allocation is not necessarily the one with the lowest cost or fastest completion time. It is the allocation that is technically feasible, institutionally legitimate, ethically defensible, and publicly accountable. For example, assigning a ventilator component to a provider may be efficient, but if the provider lacks certification, traceability, or accountability mechanisms, the allocation may be unacceptable from a governance perspective.

The implication is that AI-enabled coordination systems should include institutional constraints alongside technical constraints. These may include certification status, public mandate, accountability contact, auditability, cybersecurity readiness, history of compliance, regional equity, and ability to participate in public reporting. This makes the service composition problem more complex, but also more socially realistic.

5.3. Accountability as a condition for trustworthy AI-enabled crisis response

Accountability is a core condition for trustworthy emergency production networks. AI-enabled systems can generate recommendations, detect bottlenecks, rank providers, and support reallocation, but they cannot remove responsibility from human institutions. If the system recommends an allocation that later produces delay, quality failure, or unfair distribution, responsibility must remain traceable.

The framework therefore distinguishes between AI-supported decision-making and AI-authorized decision-making. AI-supported decision-making means that intelligent systems assist human decision makers by generating options and explaining trade-offs. AI-authorized decision-making would mean that the system itself has final authority. In public crisis response, the first

model is more appropriate. Emergency production involves safety, scarcity, public legitimacy, and legal responsibility; these cannot be delegated fully to algorithmic systems.

Accountability requires institutional clarity before the crisis occurs. Public agencies should define who has the mandate to activate an emergency production network, who can approve non-traditional providers, who validates quality, who controls access to sensitive data, and who can override AI-generated recommendations. Without these rules, AI-enabled coordination can produce accountability gaps in which decisions appear technically rational but institutionally unowned.

5.4. Transparency, explainability, and public legitimacy

Transparency is essential because crisis decisions are often contested. During emergencies, stakeholders may disagree about priorities, resource distribution, and risk tolerance. If allocation decisions are opaque, public trust can decline even when the decisions are technically defensible. Therefore, AI-enabled emergency production systems should provide decision explanations that are understandable to public officials, providers, auditors, and affected institutions.

Explainability in this context does not require exposing every mathematical detail of an algorithm. It requires explaining the practical reasons behind a recommendation. A decision-support system should be able to indicate whether a provider was selected because of capacity, proximity, certification, resilience, cost, delivery speed, regional need, or equity priority. It should also identify which constraints prevented other providers from being selected.

Transparency also includes communication with the public. When non-traditional actors such as universities, military units, or industrial firms are included in emergency production, the public may ask whether safety standards are being maintained. Public legitimacy depends on demonstrating that emergency expansion of capacity does not mean abandonment of oversight. Clear communication about qualification, testing, traceability, and accountability can help preserve trust.

5.5. Fair allocation and public value

Emergency production networks operate under public-value conditions. Their purpose is not simply to produce more units, but to support collective welfare under scarcity. This requires explicit attention to fair allocation. In the ventilator case, production decisions may influence which hospitals or regions receive scarce devices first. Even if production networks do not make bedside clinical allocation decisions, they shape the availability of resources across institutions and populations.

Fair allocation should therefore be included in the governance framework from the beginning. Allocation criteria may include urgency, shortage severity, ICU pressure, regional vulnerability, population risk, and equity across jurisdictions. These criteria should be publicly justifiable and auditable. If the system allocates outputs primarily to institutions with better data infrastructure, stronger political influence, or greater purchasing power, it may reproduce social inequality.

AI can support fair allocation by making trade-offs visible. For example, it can show how different allocation strategies affect regions with different levels of shortage. However, AI cannot determine fairness by itself. Fairness is a normative and political concept that must be defined through public governance. The role of AI should be to operationalize agreed principles, not to replace democratic or institutional judgment.

5.6. Trust and the inclusion of non-traditional actors

The inclusion of non-traditional actors is both an opportunity and a risk. Universities, military organizations, industrial firms, and laboratories can expand emergency production capacity. However, their involvement may raise questions about authority, quality, safety, and responsibility. Trust depends on whether stakeholders believe that these actors are included through clear, justified, and accountable procedures.

The framework suggests that non-traditional actors should be included through pre-defined qualification pathways. These pathways should specify what each actor can do, what standards must be satisfied, what data must be shared, and who supervises their contribution. This prevents emergency collaboration from becoming informal improvisation.

Trust also depends on role boundaries. A university laboratory may be suitable for prototyping or testing but not for final certified production. A military workshop may support logistics or fabrication but still require civilian regulatory oversight. An industrial firm may produce components but need quality verification by certified authorities. Defining these boundaries protects both public safety and institutional legitimacy.

5.7. Digital public infrastructure and cyber-governance

AI-enabled emergency production depends on digital public infrastructure. The platform that connects demand, capacity, service matching, quality records, and logistics becomes a critical governance object. If it fails, is compromised, or becomes inaccessible, the crisis-response network may lose coordination capacity.

This creates a need for cyber-governance. Cybersecurity should not be treated as a purely technical add-on. It is part of public crisis governance because digital failure can disrupt social

response. A ransomware attack, data manipulation, platform outage, or denial-of-service event may affect not only system performance but also public trust, institutional coordination, and resource allocation.

The proposed framework therefore supports fallback governance. Degrade, Substitute, and Bypass should be interpreted as public decision-continuity mechanisms. Degrade allows the system to use simplified but approved decision rules. Substitute allows an alternative provider, platform, authority, or data source to take over. Bypass allows the use of pre-approved emergency templates when real-time systems cannot be trusted. These fallback modes must be documented, authorized, and auditable.

5.8. Institutional resilience and avoidance of single points of failure

Resilience in emergency production networks is not only about production capacity; it is also about institutional architecture. A system may fail if it depends too heavily on one supplier, one platform, one public agency, one logistics channel, or one decision authority. The technical idea of anti-affinity can be translated into a governance principle: critical public-response functions should not be concentrated in a single failure domain.

Institutional anti-affinity can be implemented through distributed authority, backup providers, alternative data channels, regional redundancy, and multiple qualified service pools. For example, a ventilator production network should avoid relying on one component supplier, one certification pathway, or one digital platform. If one actor fails, the system should be able to continue through substitutes without losing accountability.

However, redundancy must be governed carefully. Excessive fragmentation can create confusion, duplication, and accountability problems. The objective is not unlimited decentralization but structured redundancy. Institutions should know who substitutes for whom, under what conditions, and with what authority.

5.9. Human oversight and responsible automation

The framework supports human-supervised AI rather than autonomous crisis governance. This is especially important because emergency production decisions involve ethical trade-offs, safety risks, public legitimacy, and institutional responsibility. AI systems may process information faster than human decision makers, but they do not possess public mandate or moral responsibility.

Human oversight should be meaningful, not symbolic. Decision makers should have the ability to understand recommendations, question assumptions, override outputs, request additional

information, and document reasons for approval or rejection. If human oversight is reduced to automatic approval of algorithmic recommendations, accountability becomes weak.

Responsible automation also requires clear escalation rules. Routine low-risk allocations may be handled with minimal review if providers are pre-qualified and constraints are satisfied. High-risk decisions, such as use of non-traditional providers, allocation to severely affected regions, or activation of emergency fallback templates, should require higher-level approval. This risk-based oversight model can preserve speed without eliminating responsibility.

5.10. Implications for public administration and crisis management

The proposed framework has several implications for public administration. First, governments should treat emergency production capacity as part of crisis preparedness. This means developing registries of convertible capacities, defining legal pathways for emergency activation, and establishing coordination protocols across sectors before crises occur.

Second, public agencies should develop governance standards for AI-enabled crisis coordination. These standards should address data quality, transparency, explainability, accountability, cyber-resilience, human oversight, and audit requirements. Without such standards, digital coordination may expand faster than governance capacity.

Third, crisis management should include institutional stress testing. Just as technical systems are tested under simulated failures, governance systems should be tested under scenarios involving provider failure, data gaps, platform outage, conflicting priorities, and ethical allocation disputes. These exercises can reveal weaknesses in authority, communication, trust, and fallback procedures.

Fourth, public administration should recognize that collaboration cannot be improvised entirely during a crisis. Effective collaboration requires pre-existing relationships, shared terminology, legal agreements, data-sharing protocols, and mutual trust. AI-enabled platforms can support collaboration, but they cannot compensate for the absence of institutional preparedness.

5.11. Limitations and future research directions

This study is conceptual and case-based. It does not empirically test the proposed framework using interviews, surveys, archival data, or comparative case analysis. Therefore, its contribution is theoretical and interpretive rather than empirically validated. Future studies should examine whether the proposed governance layers reflect actual crisis-response practices in different national, regional, or institutional contexts.

A second limitation is that the framework is developed around emergency ventilator production. Ventilators are a useful case because they are technically complex, publicly visible, and ethically significant. However, other crisis goods may require different governance arrangements. Personal protective equipment, vaccines, diagnostic kits, food supplies, and energy resources may involve different actors, risks, and allocation principles.

A third limitation is that the framework assumes the possibility of digital coordination. In some regions, digital infrastructure, data standards, or institutional trust may be insufficient to support AI-enabled emergency production. The framework may therefore require adaptation for low-resource settings or contexts with weak digital governance.

Future research can extend this study in several directions. Empirical studies should examine how emergency production networks actually operated during COVID-19. Interviews with public officials, hospital administrators, manufacturers, university laboratories, military logistics actors, and platform operators could identify real coordination challenges and accountability gaps. Comparative case studies could examine different emergency goods and different governance contexts. Future research should also develop measurable indicators of socio-technical governance performance, including accountability clarity, traceability completeness, fairness of allocation, stakeholder trust, fallback activation quality, digital continuity, and institutional learning.

6. Conclusion

Public crises expose not only material shortages but also weaknesses in institutional coordination, accountability, and public decision-making. The COVID-19 ventilator case demonstrates that emergency production cannot be understood only as a technical or industrial problem. It is also a socio-technical governance challenge in which public agencies, hospitals, universities, military organizations, manufacturers, logistics providers, regulators, and digital platforms must coordinate under urgency, uncertainty, and public scrutiny.

This study developed a case-based conceptual framework for interpreting AI-enabled cloud manufacturing as a socio-technical governance mechanism for emergency production networks. The framework reframes service composition as institutional role allocation, redundancy as backup institutional capacity, robustness as governance under uncertainty, reinforcement learning as adaptive governance, generative allocation as AI-generated crisis-response options, anti-affinity as avoidance of single points of institutional failure, cyber-resilience as trustworthy digital public

infrastructure, fallback orchestration as continuity of public decision-making, traceability as accountability, and equity as a condition of public legitimacy.

The proposed framework consists of six interrelated governance layers: the public-value and crisis-demand layer, the distributed institutional capacity layer, the AI-enabled coordination layer, the accountability and traceability layer, the ethical allocation and public trust layer, and the resilience and fallback governance layer. Together, these layers explain how distributed capacity can be transformed into collective crisis-response capability. The framework emphasizes that technical capacity becomes socially meaningful only when it is embedded in legitimate authority, transparent decision rules, accountable institutions, auditable data practices, and human-supervised AI.

The COVID-19 ventilator production case illustrates the importance of this perspective. During a crisis, universities, military units, industrial firms, laboratories, logistics providers, and healthcare organizations may all contribute useful capacity. However, their participation must be governed through clear role definitions, qualification pathways, traceability requirements, oversight mechanisms, and fair allocation principles. Without such governance, emergency production networks may increase output while creating ambiguity, mistrust, or accountability gaps.

The study contributes to humanities and social science by moving AI-enabled cloud manufacturing beyond its usual technical framing. It shows that emergency production networks should be evaluated not only by speed, efficiency, cost, or resilience metrics, but also by institutional legitimacy, public trust, accountability clarity, ethical allocation, transparency, and continuity of public decision-making. This perspective is important because digital coordination platforms and AI-supported allocation systems increasingly shape how societies respond to crises.

Several limitations should be acknowledged. The study is conceptual and case-based; it does not empirically test the framework through interviews, surveys, archival data, or comparative case analysis. The framework is also developed around ventilator production, which is a highly visible and technically complex emergency good. Other crisis goods, such as personal protective equipment, vaccines, diagnostic kits, food, or energy resources, may require different governance arrangements. In addition, the framework assumes the availability of digital coordination capacity, which may not exist equally across regions or institutional contexts.

Future research should empirically examine how emergency production networks operated during COVID-19 and other crises. Comparative case studies can investigate how different governance systems coordinated distributed capacity, handled accountability, and used digital platforms. Interview-based research can explore the experiences of public officials, hospital administrators, manufacturers, university laboratories, military logistics units, regulators, and platform operators. Future studies should also develop measurable indicators for socio-technical governance performance, including accountability clarity, traceability completeness, fairness of allocation, stakeholder trust, fallback readiness, cyber-continuity, and institutional learning.

Overall, the proposed framework suggests that AI-enabled emergency production should be governed as a public socio-technical system, not merely deployed as an optimization tool. A crisis-response system is socially effective only when it can coordinate distributed capacity while preserving legitimacy, responsibility, fairness, and trust. Designing such systems requires integration of digital technology, institutional governance, ethical reasoning, and public accountability.

References

- [1] Ansell, C., & Gash, A. (2008). Collaborative governance in theory and practice. *Journal of Public Administration Research and Theory*, 18(4), 543-571. <https://doi.org/10.1093/jopart/mum032>
- [2] Emerson, K., Nabatchi, T., & Balogh, S. (2012). An integrative framework for collaborative governance. *Journal of Public Administration Research and Theory*, 22(1), 1-29. <https://doi.org/10.1093/jopart/mur011>
- [3] Comfort, L. K. (2007). Crisis management in hindsight: Cognition, communication, coordination, and control. *Public Administration Review*, 67(s1), 189-197. <https://doi.org/10.1111/j.1540-6210.2007.00827.x>
- [4] Boin, A., & Lodge, M. (2016). Designing resilient institutions for transboundary crisis management: A time for public administration. *Public Administration*, 94(2), 289-298. <https://doi.org/10.1111/padm.12264>
- [5] Bovens, M. (2007). Analysing and assessing accountability: A conceptual framework. *European Law Journal*, 13(4), 447-468. <https://doi.org/10.1111/j.1468-0386.2007.00378.x>
- [6] Floridi, L., & Cowls, J. (2019). A unified framework of five principles for AI in society. *Harvard Data Science Review*, 1(1). <https://doi.org/10.1162/99608f92.8cd550d1>
- [7] Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389-399. <https://doi.org/10.1038/s42256-019-0088-2>
- [8] Mittelstadt, B. D., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). <https://doi.org/10.1177/2053951716679679>
- [9] Mergel, I., Edelmann, N., & Haug, N. (2019). Defining digital transformation: Results from expert interviews. *Government Information Quarterly*, 36(4), 101385. <https://doi.org/10.1016/j.giq.2019.06.002>
- [10] Moore, M. H. (1995). *Creating public value: Strategic management in government*. Harvard University Press.
- [11] Shahab, E., Kazemisaboor, A., Khaleghparast, S., & Fatahi Valilai, O. (2023). A production bounce-back approach in the Cloud manufacturing network: Case study of COVID-19 pandemic. *International Journal of Management Science and Engineering Management*, 18(4), 305-317. <https://doi.org/10.1080/17509653.2022.2112781>
- [12] Shahab, E., Rabiee, M., Mobasser, N., & Fatahi Valilai, O. (2025). A robust service composition for a resilient cloud manufacturing service network. *International Journal of Computer Integrated Manufacturing*. <https://doi.org/10.1080/0951192X.2025.2504088>

- [13] Shahab, E., Taleb, M., Gholian-Jouybari, F., & Hajiaghahi-Keshteli, M. (2024). Designing a resilient cloud network fulfilled by reinforcement learning. *Expert Systems with Applications*, 255, 124606. <https://doi.org/10.1016/j.eswa.2024.124606>
- [14] Shahab, E., & Taghipour, S. (2025). Resilient LLM-DBMS pipelines via event-driven fallback orchestration. In 2025 IEEE International Conference on Collaborative Advances in Software and Computing (CASCON). <https://doi.org/10.1109/CASCON66301.2025.00113>
- [15] Shahab, E., Taghipour, S., & Moghadam, P. (2026). A generative AI-based approach for resilient service composition under cybersecurity attacks in cloud-fog networks. *Future Generation Computer Systems*, 177, 108273. <https://doi.org/10.1016/j.future.2025.108273>
- [16] Shahab, E., & Taghipour, S. (2026). Quantum reinforcement learning for resilient cloud service composition. *Soft Computing*. <https://doi.org/10.1007/s00500-025-11052-6>
- [17] Shahab, E., Moghadam, P., & Taghipour, S. (2026). Diffusion model-based generative optimisation for resilient service composition in cloud manufacturing. *International Journal of Production Research*. <https://doi.org/10.1080/00207543.2026.2625247>
- [18] Shahab, E., & Taghipour, S. (2025). Designing a resilient cloud network fulfilled by quantum machine learning. *International Journal of Management Science and Engineering Management*. <https://doi.org/10.1080/17509653.2025.2544566>
- [19] Kapucu, N. (2006). Interagency communication networks during emergencies: Boundary spanners in multiagency coordination. *The American Review of Public Administration*, 36(2), 207-225. <https://doi.org/10.1177/0275074005280605>