International Journal of Medical and Applied Health Science (IJMAHS)

Contents lists available at IJMAHS
International Journal of Medical and Applied Health Science
journal homepage: IJMAHS
Volume 1, No. 2, 2025

A Health-System Resilience Framework for Emergency Ventilator Production Using Cloud Networks

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:
Emergency medical equipment, Ventilator supply, Cloud networks, Resilience, Service composition, Cyber-continuity.

ABSTRACT
Public health crises can rapidly expose shortages in critical medical equipment, especially when demand for life-support devices exceeds conventional production and procurement capacity. Ventilator shortages during respiratory pandemics show that emergency medical-equipment availability is not only a supply-chain issue but also a health-system resilience challenge. This study develops a case-based methodological framework for emergency ventilator production using cloud manufacturing and intelligent service composition. The framework translates distributed production capacity into a coordinated emergency response structure through five layers: health-demand sensing, cloud manufacturing services, intelligent service composition, resilience and cyber-continuity, and health-system evaluation. The COVID-19 ventilator production case is used to illustrate how hospitals, public health agencies, medical-device manufacturers, universities, military workshops, industrial firms, laboratories, component suppliers, and logistics providers can be coordinated during severe disruptions. The framework integrates redundancy, robustness, adaptive allocation, generative decision support, cyber-resilient migration, and fallback orchestration while emphasizing quality compliance, traceability, human oversight, and equitable distribution. The study contributes to applied health science by reframing emergency production as a socio-technical health-system capability rather than a purely industrial process. Future validation should use hospital demand data, certified production datasets, simulation-based disruption scenarios, and digital twins to evaluate shortage reduction, recovery time, fulfillment rate, and continuity of care.

1. Introduction

Public health crises expose structural weaknesses in the production and distribution of critical medical equipment. During large-scale emergencies, the demand for life-support devices,

[a] Corresponding author email address: Hosein.gholami72@gmail.com (Hossein Gholami Ghadi).
DOI: https://doi.org/10.22034/ijmahs.v1i2.209
Available online 2026/05/31

3115-7378/BGSA Ltd.

protective equipment, diagnostic supplies, and other essential medical technologies can rise faster than conventional manufacturing and procurement systems can respond. Ventilators represent one of the clearest examples of this vulnerability. In pandemic conditions, ventilator availability is not only a manufacturing issue but also a determinant of health-system readiness, intensive-care capacity, and continuity of critical care. When ordinary production channels are disrupted, delayed, or insufficient, health systems require alternative mechanisms for mobilizing distributed production capacity, reallocating resources, and sustaining the supply of critical devices under uncertainty.

The COVID-19 pandemic demonstrated that emergency medical-equipment shortages cannot be addressed solely through traditional centralized production models. Critical shortages of ventilators and personal protective equipment were widely recognized as major health-system risks during the early stages of the pandemic [1]. The problem was not merely the lack of individual manufacturers; it was also the limited ability to rapidly coordinate dispersed capacities across different organizations, sectors, and supply networks. Prior work on production bounce-back in cloud manufacturing examined ventilator production during the COVID-19 pandemic and showed that unused capacities from other supply networks, including military organizations and university research groups, could be considered as part of a resilience-oriented manufacturing response [2]. That perspective is relevant for applied health science because it reframes ventilator production as a coordinated health-system resilience problem rather than a narrow industrial production problem.

Cloud manufacturing provides a useful foundation for this type of emergency response. It transforms distributed manufacturing resources and capabilities into cloud-accessible services that can be searched, selected, combined, and coordinated through a service-oriented platform. Instead of relying on a single factory or a fixed supplier network, cloud manufacturing enables a central decision-making mechanism to match decomposed production tasks with available service providers. For emergency ventilator production, this means that activities such as component fabrication, assembly, quality inspection, calibration, packaging, and distribution can be assigned to different qualified entities according to capacity, time, cost, quality, and availability constraints. In this sense, service composition becomes a practical decision-support mechanism for linking urgent medical-equipment demand with distributed production resources.

However, emergency medical-equipment production requires more than ordinary service matching. Health crises are characterized by severe uncertainty, demand surges, supplier

disruptions, logistics delays, and changing operational constraints. Therefore, the relevant objective is not only efficiency but resilience. Reliability-oriented approaches are insufficient because they mainly emphasize the ability of a system to avoid failure under expected conditions. Resilience, by contrast, concerns the ability to withstand disruption, respond during disruption, and recover afterward. Prior cloud manufacturing research has emphasized this distinction and argued that resilience depends on mechanisms such as redundancy, robustness, flexibility, decentralization, and collaboration [2,3]. For emergency ventilator production, these mechanisms can be interpreted in health-system terms: redundancy reduces shortage risk, robustness protects production plans against uncertainty, flexibility enables rapid reconfiguration, decentralization reduces dependence on single production nodes, and collaboration allows non-traditional actors to support medical-device supply.

Recent studies have advanced resilient cloud manufacturing through different methodological directions. A robust service composition model introduced subentropy-based uncertainty management and a mixed-integer nonlinear programming formulation to improve the resilience of cloud manufacturing networks, with validation through a ventilator production case during COVID-19 [3]. This contribution is important because it shows that the resilience of emergency production networks can be improved by deliberately balancing service diversity, redundancy, and cost. In particular, subentropy can be interpreted as a mechanism for avoiding overly concentrated service assignments, which is relevant when medical-equipment production must remain functional despite disruption to specific suppliers or facilities.

Artificial intelligence further expands the decision-making capacity of resilient cloud manufacturing. Reinforcement learning has been proposed for stochastic cloud manufacturing environments where service availability, demand, and disruption patterns change over time. Prior work applied reinforcement learning algorithms, including Deep Deterministic Policy Gradient, Twin Delayed Deep Deterministic Policy Gradient, Proximal Policy Optimization, and Soft Actor-Critic, to resilient cloud manufacturing service composition and validated the model using a COVID-19 ventilator production case [4]. The relevance of this approach for health-system resilience is clear: emergency equipment production is a dynamic decision problem in which allocation policies must adapt as capacities fail, recover, or become newly available.

More recent developments also suggest that generative AI, quantum-enhanced learning, and fallback orchestration can support emergency production networks, although these methods should

be interpreted as complementary decision-support mechanisms rather than substitutes for regulatory oversight and medical-device quality assurance. Generative optimization can produce feasible or near-feasible task-to-resource allocations under disruption; cyber-resilient allocation can maintain quality of service during attacks on cloud-fog infrastructure; and event-driven fallback orchestration can preserve operational continuity when digital decision-support components experience provider failures, latency spikes, quota exhaustion, or data drift [5-7]. These capabilities are increasingly relevant because emergency medical production depends not only on physical capacity but also on digital coordination platforms, databases, cloud services, and communication systems.

Despite these advances, an important research gap remains. Existing studies on resilient cloud manufacturing, robust service composition, reinforcement learning, generative allocation, quantum-enhanced optimization, cyber-resilient orchestration, and fallback mechanisms have primarily framed the problem as a manufacturing, computing, or optimization challenge. Limited attention has been given to translating these technical mechanisms into an integrated health-system resilience framework for emergency medical-equipment production. In particular, the literature has not sufficiently explained how distributed production capacity can be organized to support ventilator availability, hospital preparedness, continuity of critical care, and public health emergency response.

To address this gap, this study develops a case-based methodological framework for emergency ventilator production using cloud manufacturing and intelligent service composition. The COVID-19 ventilator production case is used as an illustrative application because it provides a concrete example of how distributed capacities from manufacturers, universities, military units, research laboratories, logistics providers, and alternative supply networks can be coordinated to sustain the production of critical medical equipment.

This study makes three contributions. First, it reframes cloud manufacturing service composition as a health-system resilience mechanism for emergency medical-equipment production. Second, it integrates multiple resilience mechanisms - redundancy, robustness, adaptive learning, generative allocation, service migration, cyber-resilience, and fallback orchestration - into a unified framework for crisis-responsive ventilator production. Third, it derives applied health-system implications for hospitals, public health agencies, emergency planners, and medical-device

supply networks by showing how distributed manufacturing capacity can be converted into a coordinated emergency production system.

The remainder of this paper is organized as follows. Section 2 reviews the literature on emergency medical-equipment supply resilience, cloud manufacturing service composition, AI-enabled allocation, cyber-resilient orchestration, and fallback mechanisms. Section 3 presents the proposed health-system resilience framework and explains its layers, decision logic, and ventilator production case application. Section 4 discusses implications for hospitals, public health agencies, medical-device supply networks, and emergency governance. Section 5 concludes the paper and identifies directions for future validation.

2. Literature Review

2.1. Emergency medical-equipment supply resilience

Public health emergencies create abrupt and severe pressure on the supply of critical medical equipment. Unlike ordinary market demand, emergency demand is time-sensitive, uncertain, and directly connected to patient outcomes. Devices such as ventilators, oxygen systems, diagnostic equipment, and protective technologies become essential components of health-system response. When conventional procurement channels and centralized production systems cannot respond quickly, shortages can affect intensive-care capacity, emergency preparedness, and continuity of care. Therefore, the resilience of medical-equipment supply networks is not merely a logistical issue; it is an operational condition for maintaining healthcare delivery during large-scale crises.

The COVID-19 pandemic demonstrated that medical-equipment supply systems can be disrupted simultaneously by demand surges, production bottlenecks, transportation restrictions, workforce limitations, and resource scarcity. Ventilator production became a visible example of this challenge because demand increased rapidly while ordinary manufacturing capacity was limited. Ranney et al. [1] emphasized the urgency of ventilator and personal protective equipment shortages during the pandemic, while Emanuel et al. [8] showed that scarcity of critical medical resources raised ethical and policy issues regarding fair allocation. These studies show that equipment availability is tied to both operational readiness and ethical health-system governance.

The vulnerability of medical product supply chains became a major patient-safety issue during COVID-19 [9]. In parallel, empirical evidence from European healthcare supply chains showed that organizations used both buffering and bridging strategies to improve resilience in the face of medical supply scarcity [10]. OECD analysis [11] also emphasizes that reliable medical supply

chains are a cornerstone of resilient health systems and that future severe crises require better anticipation and mitigation of medicine and medical-device shortage risks (OECD, 2024).

Prior research on production bounce-back in cloud manufacturing directly addressed this issue by examining ventilator production during the COVID-19 pandemic. That study proposed a service composition model that used redundancy and subentropy to support resilient production under disruption, while also considering unused capacities from other supply networks, such as military organizations and university research groups [2]. This contribution is important for health-system resilience because it shows that emergency medical-equipment production can be treated as a distributed capacity coordination problem. Rather than depending only on existing medical-device manufacturers, a health system can potentially use a broader network of qualified organizations to support emergency production.

From an applied health science perspective, this idea extends the concept of emergency preparedness. Preparedness should not be limited to stockpiling and procurement contracts; it should also include the capacity to identify, certify, and activate distributed production resources before and during a crisis. In this setting, cloud manufacturing becomes a digital and organizational infrastructure for emergency capacity pooling.

2.2. Cloud manufacturing and service composition for crisis response

Cloud manufacturing is a service-oriented production paradigm in which distributed manufacturing resources, capabilities, and processes are virtualized and made accessible through cloud-based platforms. This allows production tasks to be allocated across a network of service providers rather than executed within a single factory or fixed supply chain. The relevance of cloud manufacturing to public health crises lies in its ability to connect emergency demand with distributed capacity. When a medical-equipment shortage occurs, the production process can be decomposed into tasks such as component fabrication, assembly, calibration, testing, packaging, and delivery. Each task can then be matched with candidate service providers based on capacity, cost, time, availability, technical compatibility, and quality requirements.

Service composition is the decision process that selects and combines multiple services to satisfy a complex demand. In ordinary manufacturing contexts, service composition is often evaluated through performance indicators such as cost, time, reliability, reputation, energy use, and availability. In emergency medical-equipment production, however, additional concerns become critical. These include continuity of supply, compliance with medical-device standards,

traceability, responsiveness to hospital demand, and the ability to recover when one or more production nodes fail.

The production bounce-back approach introduced redundancy as a central resilience strategy for cloud manufacturing networks. The model considered subentropy at the process level while controlling overall entropy at the network level, thereby supporting a balanced form of redundancy [2]. This is relevant because excessive concentration of tasks among a small number of providers can make the system fragile, while uncontrolled diversification can increase complexity and coordination cost. A balanced redundancy policy helps ensure that emergency production can continue even if some suppliers, facilities, or services are disrupted.

A later robust service composition study extended this direction by developing a mixed-integer nonlinear programming model that incorporates subentropy to handle uncertainty across different scenarios [3]. The study compared Particle Swarm Optimization, Genetic Algorithm, and Simulated Annealing, and validated the framework through a ventilator production case during COVID-19. Its contribution lies in moving from deterministic service composition toward uncertainty-aware service composition. For health-system applications, this distinction is important because emergency production plans must remain feasible under uncertain demand, uncertain capacity, disrupted logistics, and changing resource availability.

2.3. Resilience mechanisms: redundancy, robustness, flexibility, and collaboration

The literature distinguishes reliability from resilience. Reliability refers to the probability that a system performs its intended function under specified conditions, whereas resilience emphasizes the ability to withstand disruption, respond during disruption, and recover after disruption. In emergency medical-equipment production, reliability alone is insufficient because a production network may be reliable under normal conditions but unable to adapt when demand surges, suppliers fail, or logistics routes are disrupted. Resilience therefore requires mechanisms that preserve functionality under abnormal and evolving conditions.

Redundancy is one of the most direct resilience mechanisms. In the context of ventilator production, redundancy can mean alternative suppliers for critical components, backup assembly sites, parallel logistics routes, or reserve capacity from non-traditional manufacturers. The production bounce-back model used redundancy to support disruption fulfillment and showed how unused capacities in other supply networks could support ventilator production during COVID-19

[2]. For health systems, this implies that emergency manufacturing capacity should be mapped before crises occur, rather than discovered reactively during shortages.

Robustness is another central mechanism. A robust production plan is designed to perform acceptably across uncertain scenarios. In cloud manufacturing, robustness can be implemented by considering uncertain parameters, scenario ranges, and uncertainty sets. The robust service composition model used subentropy as a scenario-related parameter and proposed a robust MINLP formulation to enhance resilience under diverse conditions [3]. In health-system terms, robustness supports continuity when demand, capacity, lead time, or supplier availability deviates from expected values.

Flexibility refers to the ability to reconfigure the system when conditions change. For emergency ventilator production, flexibility may involve switching from one supplier to another, reallocating tasks to a different production site, changing assembly sequences, or activating alternative logistics channels. Flexibility is especially important when disruptions are not static. A provider may fail temporarily, recover later, or become overloaded as demand changes. Therefore, service composition must be dynamic rather than fixed.

Collaboration is also essential because emergency production often requires capabilities distributed across sectors. During COVID-19, many organizations adapted production capabilities to meet urgent social needs. Cloud manufacturing formalizes this collaborative logic by providing a platform through which heterogeneous actors can be coordinated. In the ventilator case, the inclusion of universities, military units, and industrial firms demonstrates that resilience can involve adaptation and transformation, not only stability.

2.4. AI-enabled adaptive allocation in disrupted production networks

Classical optimization models can support service composition, but health-crisis environments often require adaptive decision-making. Disruptions may occur sequentially, and decision makers may need to update allocations as new information becomes available. Reinforcement learning provides one approach for this setting because it allows an agent to learn allocation policies through interaction with a dynamic environment. Instead of optimizing a single static plan, reinforcement learning can support repeated decision-making under stochastic disruptions.

A reinforcement-learning-based service composition study applied algorithms including Deep Deterministic Policy Gradient, Twin Delayed Deep Deterministic Policy Gradient, Proximal Policy Optimization, and Soft Actor-Critic to resilient cloud manufacturing networks [4]. The

model was applied to a real-world ventilator production case during the COVID-19 pandemic and showed the relevance of reinforcement learning for stochastic disruption scenarios. This is important for emergency medical-equipment production because the value of adaptation can exceed the cost of changing service assignments when unmet demand or delayed delivery creates serious health-system consequences.

Generative AI offers another emerging direction for disrupted allocation problems. A diffusion model-based generative optimization framework used a variational autoencoder to learn a latent manifold of feasible allocation plans and a diffusion-based model to refine disruption-aware allocations [7]. This approach is relevant because disrupted production networks often require fast generation of feasible candidate allocations. Instead of solving each disruption from scratch, generative methods can learn the structure of good allocations and produce high-quality warm starts. For health-system resilience, this capability can reduce response time when emergency demand changes rapidly.

Quantum-enhanced methods provide a further possible direction for complex allocation problems, especially when the decision space becomes large. Prior work on quantum machine learning proposed a service migration framework that maximizes quality of service while minimizing migration cost under capacity, energy, and jitter constraints [12]. Another study proposed quantum reinforcement learning for resilient cloud service composition and integrated anti-affinity constraints to reduce single points of failure [13]. In the health-system context, these methods are best interpreted as future-facing decision-support tools. They may eventually support rapid exploration of large service allocation spaces, but their practical deployment must consider hardware limitations, explainability, and governance requirements.

2.5. Cyber-resilience and fallback orchestration

Emergency medical-equipment production increasingly depends on digital infrastructure. Cloud manufacturing platforms, databases, dashboards, allocation engines, supplier registries, logistics systems, and communication tools all contribute to production coordination. Therefore, resilience cannot be limited to physical production capacity. Digital and cyber-resilience are also required. If the digital infrastructure coordinating emergency production fails, suffers latency spikes, experiences data drift, or is compromised by cyberattacks, the production network may lose visibility and responsiveness.

Cyberattacks such as distributed denial-of-service and ransomware can degrade quality of service in cloud-fog systems by overloading links and compromising nodes. A generative AI-based cyber-resilient service composition study proposed a conditional generative adversarial network for task-to-server allocation and combined it with anomaly detection and reactive migration [6]. The model evaluated DDoS and ransomware scenarios and showed that generative allocation and migration can help maintain quality of service under attack. For health-system applications, this line of research is relevant because hospital-linked production networks, digital procurement systems, and medical-equipment coordination platforms may become targets during crises.

Fallback orchestration provides another operational resilience mechanism. Event-driven fallback systems can preserve service continuity when primary digital components fail. A self-healing LLM-DBMS pipeline study proposed three fallback tiers: Degrade, Substitute, and Bypass [5]. In that architecture, Degrade reduces service complexity, Substitute switches to alternative models or methods, and Bypass uses cached or deterministic outputs when higher-tier intelligence is unavailable. For emergency ventilator production, fallback logic can be translated into operational terms. A Degrade strategy may use simplified allocation rules when a full optimization model is unavailable. A Substitute strategy may switch to an alternative decision engine, supplier pool, or logistics provider. A Bypass strategy may activate pre-approved emergency allocation templates when real-time systems fail.

2.6. Synthesis and research gap

The reviewed literature shows that resilient cloud manufacturing has developed through several complementary directions: redundancy-based service composition, robust optimization, reinforcement learning, generative allocation, quantum-enhanced migration, cyber-resilient orchestration, and fallback-based self-healing. Each stream contributes a useful mechanism for emergency production. Redundancy enables backup capacity, robustness protects against uncertainty, reinforcement learning supports adaptive reallocation, generative AI accelerates allocation generation, quantum-enhanced methods expand future optimization possibilities, cyber-resilience protects digital coordination layers, and fallback orchestration preserves continuity when intelligent systems degrade.

However, these contributions remain fragmented from a health-system perspective. Most prior work frames the problem as manufacturing optimization, cloud-resource allocation, cyber-resilience, or AI-enabled service composition. Limited attention has been given to integrating these

mechanisms into a health-oriented framework that addresses emergency medical-equipment availability, hospital demand fulfillment, continuity of care, public health preparedness, and medical-device governance. This creates a clear research gap: existing technical models need to be translated into an applied health-system resilience framework.

Table 1. Synthesis of technical mechanisms and health-system relevance.

Technical mechanism	Source stream	Health-system relevance	Emergency ventilator example
Redundancy	Cloud manufacturing resilience	Reduces shortage risk	Backup services for key components
Subentropy	Resilient service composition	Balances redundancy and complexity	Avoids overdependence on one supplier
Robust optimization	Uncertainty-aware allocation	Maintains feasible production under disruption	Plans remain feasible under demand surge
Reinforcement learning	Adaptive allocation	Supports repeated reallocation	Tasks reassigned when providers fail
Generative allocation	Fast solution generation	Reduces response time	Candidate allocations generated rapidly
Cyber-resilient migration	Cloud-fog resilience	Maintains digital coordination	Tasks migrated after cyber disruption
Fallback orchestration	Self-healing systems	Prevents decision paralysis	Pre-approved templates activated
Quantum-enhanced search	Future decision support	Explores large allocation spaces	Future support for complex networks

3. Proposed Framework and Case-Based Methodology

3.1. Study design

This study adopts a case-based methodological framework design to translate AI-enabled cloud manufacturing mechanisms into a health-system resilience model for emergency medical-equipment production. The purpose is not to introduce a new numerical optimization experiment, but to synthesize and reposition prior technical mechanisms - cloud manufacturing service composition, redundancy, robust optimization, reinforcement learning, generative allocation, cyber-resilient migration, quantum-enhanced decision support, and fallback orchestration - within an applied health-science context.

The proposed framework uses emergency ventilator production during the COVID-19 pandemic as an illustrative case. Ventilators are selected because they represent a critical medical device whose shortage can directly affect intensive-care capacity and continuity of care during respiratory pandemics. The case also provides a suitable bridge between medical-equipment supply resilience and cloud manufacturing because ventilator production involves multiple tasks, heterogeneous resources, strict quality requirements, uncertain demand, and the possibility of activating non-traditional production capacities.

The framework is developed around five analytical layers: the health-demand layer, the cloud manufacturing service layer, the intelligent service composition layer, the resilience and cyber-continuity layer, and the health-system evaluation layer. These layers connect hospital demand signals to distributed manufacturing capacities and then to AI-enabled decision support mechanisms.

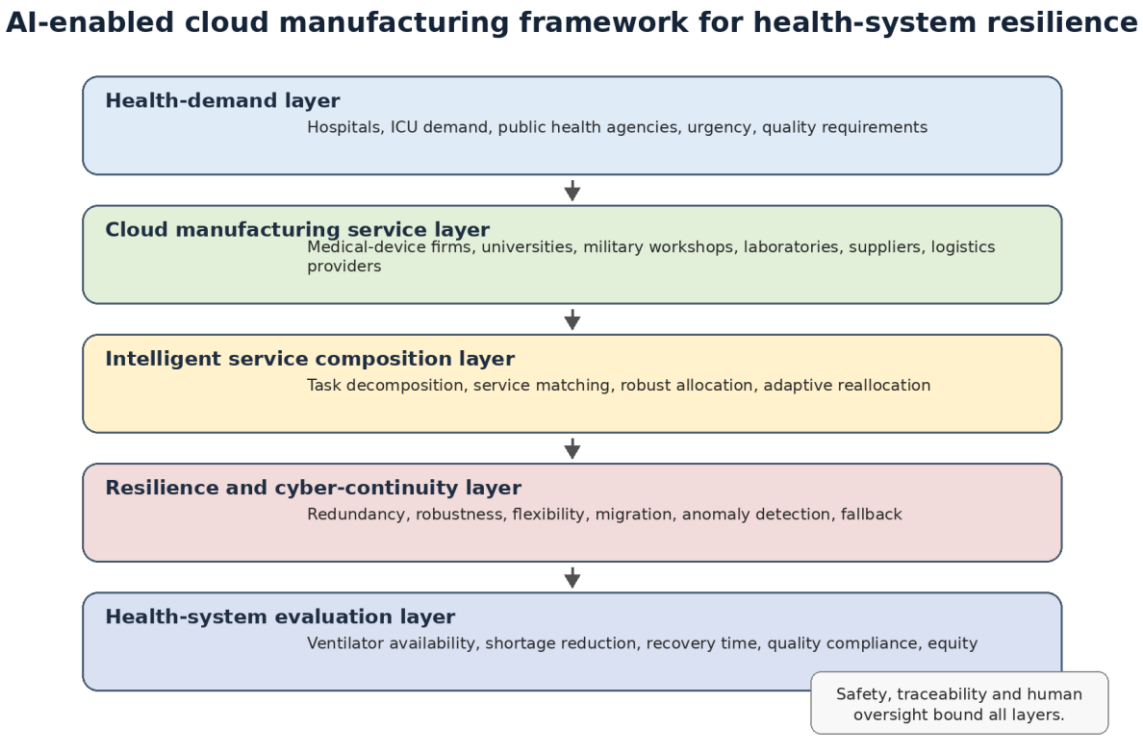

Fig. 1. AI-enabled cloud manufacturing framework for health-system resilience.

3.2. Conceptual logic of the framework

The central premise of the framework is that emergency medical-equipment production can be represented as a dynamic service composition problem. In this representation, hospitals and public health agencies generate urgent demand for medical devices. Each device is decomposed into

production and delivery tasks. These tasks are then matched with qualified services distributed across a cloud manufacturing network. The network may include medical-device manufacturers, industrial firms, university laboratories, military workshops, component suppliers, logistics providers, and quality-control units.

The framework assumes that a health crisis creates three simultaneous pressures. First, demand becomes time-sensitive and volatile. Second, ordinary production capacity may be insufficient or disrupted. Third, the digital coordination infrastructure required for allocation, monitoring, and communication may itself be exposed to failures, cyberattacks, latency spikes, or data-quality problems. Therefore, an emergency production system should not only optimize cost and production time; it should also maintain resilience, traceability, quality compliance, and continuity under disruption.

3.3. Health-demand layer

The health-demand layer represents the demand side of emergency medical-equipment production. Its purpose is to translate public health needs into operational production requirements. During a crisis, demand for ventilators may arise from hospitals, intensive-care units, public health agencies, procurement authorities, or regional emergency-management bodies. These demand signals should be expressed not only as quantities but also as time-critical requirements.

For emergency ventilator production, the health-demand layer includes hospital-level demand, acceptable delivery windows, clinical criticality, quality requirements, and distribution priority. Hospital-level demand indicates the number of ventilators required by specific hospitals or regions. Time urgency indicates the acceptable delivery window. Clinical criticality indicates the severity of shortage and the potential effect on care continuity. Quality requirements define the minimum acceptable medical-device standards, certification requirements, and testing protocols. Distribution priority determines whether available units should be allocated based on ICU pressure, regional shortage severity, population risk, or emergency policy criteria.

The output of this layer is an emergency demand profile. This profile defines what must be produced, where it is needed, how urgently it is required, and what quality and regulatory constraints must be satisfied. In an operational implementation, this profile could be updated periodically as hospital admissions, ICU occupancy, inventory levels, and supply forecasts change.

3.4. Cloud manufacturing service layer

The cloud manufacturing service layer represents the supply side of the framework. It transforms distributed production capacities into selectable services. Each service corresponds to a resource

or capability that can contribute to emergency ventilator production. Examples include component manufacturing, electronics assembly, mechanical part fabrication, sensor integration, sterilization, calibration, testing, packaging, warehousing, and delivery.

In this layer, service providers are not limited to conventional medical-device manufacturers. A crisis may require the activation of alternative capacities from other sectors. Universities may provide prototyping, design validation, laboratory testing, or engineering expertise. Military workshops may provide rapid fabrication, logistics, and disciplined operational coordination. Automotive or industrial firms may provide machining, assembly, and production-scale capability. Logistics providers may support distribution to hospitals. However, all such services must be filtered through quality, certification, and safety requirements before being included in the emergency production network.

Each service can be described through capacity attributes, operational attributes, quality attributes, and resilience attributes. Capacity attributes indicate how many units or components can be produced within a given time. Operational attributes describe technical compatibility, production capability, and required inputs. Quality attributes describe certification status, testing capability, reliability, and traceability. Resilience attributes describe availability under disruption, backup capacity, cyber-dependence, geographic risk, and ability to substitute for other services.

The output of this layer is a qualified service pool. This pool defines which services are available, what they can produce, what constraints they face, and how they can be combined to satisfy emergency ventilator demand.

3.5. Intelligent service composition layer

The intelligent service composition layer links the health-demand profile to the qualified service pool. Its purpose is to select and combine services so that emergency ventilator production can continue under constraints and disruptions. In this layer, ventilator production is decomposed into tasks, and each task is assigned to one or more candidate services.

The service composition problem can be described as follows. Let the set of emergency production tasks be represented by T, and let the set of qualified services be represented by S. Each task t in T requires capacity, technical compatibility, quality compliance, and completion within an acceptable time window. Each service s in S has limited capacity, cost, availability, production time, and quality attributes. The objective is to construct a feasible task-service allocation that

supports ventilator production continuity while satisfying health-system, operational, and resilience constraints.

Unlike ordinary service composition, emergency medical-equipment production requires health-sensitive constraints. First, quality compliance must not be relaxed below medical-device safety thresholds. Second, critical components should not depend on a single fragile service provider. Third, production lead time should be evaluated relative to the urgency of hospital demand. Fourth, allocation decisions should consider shortage consequences, not only manufacturing cost. Fifth, traceability should be maintained so that each produced unit can be linked to approved services, components, and testing records.

The framework allows different decision-support modules to support the service composition process. A robust optimization module can generate allocations that remain feasible under uncertainty. A reinforcement learning module can support repeated reallocation as disruptions occur over time. A generative allocation module can produce rapid candidate allocations when the network state changes. A quantum-enhanced module can be considered for future high-dimensional allocation settings. A fallback module can preserve decision continuity when the primary decision-support system is unavailable.

3.6. Resilience and cyber-continuity layer

The resilience and cyber-continuity layer protect the emergency production system against physical, operational, and digital disruptions. This layer is necessary because health-crisis production depends on both material capacity and digital coordination. Even if production resources exist, the system may fail if allocation databases, cloud platforms, dashboards, supplier registries, or communication tools become unavailable or compromised.

The layer includes five resilience mechanisms. Redundancy ensures that alternative services are available for critical production tasks. In ventilator production, redundancy may involve multiple suppliers for valves, sensors, tubes, control boards, or assembly operations. Robustness ensures that allocations remain acceptable under uncertain demand, capacity, lead time, and disruption scenarios. Flexibility enables task reallocation when a service becomes unavailable, overloaded, delayed, or newly available. Cyber-continuity protects the digital orchestration layer against attacks, provider outages, data drift, and latency spikes. Fallback orchestration preserves continuity when advanced decision support becomes unavailable.

The framework adapts three fallback tiers: Degrade, Substitute, and Bypass. In the Degrade tier, the system uses simplified allocation rules or fewer optimization criteria to produce a timely decision. In the Substitute tier, an alternative decision engine, model, provider, or service pool is activated. In the Bypass tier, pre-approved emergency allocation templates or cached feasible plans are used to preserve continuity when real-time optimization is not possible. Each fallback tier must be constrained by safety, traceability, and quality-control rules.

3.7. Case illustration: emergency ventilator production during COVID-19

The ventilator production case illustrates how the proposed framework can operate during a public health crisis. The process begins when hospitals or public health agencies report an urgent shortage of ventilators. This demand is translated into an emergency demand profile specifying quantity, location, urgency, and minimum quality requirements.

The production requirement is then decomposed into tasks. These tasks may include mechanical component fabrication, electronic control unit assembly, airflow system production, sensor integration, casing production, software configuration, calibration, quality testing, sterilization, packaging, and transportation to hospitals. Each task requires specific capabilities and may have different constraints.

The cloud manufacturing service layer identifies qualified service providers for each task. Conventional medical-device manufacturers may cover core certified production activities. Industrial firms may support component production. Universities may support prototyping, testing, or engineering validation. Military workshops may support rapid fabrication or emergency logistics. Logistics providers may distribute finished ventilators or components across regions.

The intelligent service composition layer then matches tasks to services. The allocation should satisfy capacity, time, cost, compatibility, and quality constraints. It should also avoid fragile configurations. For example, if a critical sensor component depends on a single provider that is geographically exposed or capacity constrained, the framework should identify backup capacity or split the task across multiple qualified services if quality requirements allow.

During the crisis, disruptions may occur. A supplier may lose capacity, a logistics route may be delayed, a digital coordination platform may suffer latency, or a cyberattack may compromise a node. The resilience layer responds by reallocating tasks, activating alternative services, or applying fallback rules. If the primary optimization model cannot be used quickly enough, a degraded but safe allocation procedure may be applied. If a service provider becomes unavailable,

a substitute provider may be activated. If digital coordination is severely disrupted, pre-approved emergency production templates may be used until full functionality is restored.

3.8. Health-system evaluation layer

The health-system evaluation layer defines how the framework should be assessed. Since the framework is designed for applied health science, evaluation should not be limited to manufacturing cost or computational performance. Instead, performance should be measured through indicators that connect production decisions to health-system resilience.

The first indicator is ventilator production continuity. This measures whether production can continue despite disruption to specific services, suppliers, or digital components. The second indicator is shortage reduction, which measures the extent to which emergency demand is fulfilled. The third indicator is recovery time, which measures how quickly the system restores production after a disruption. The fourth indicator is fulfillment rate, which measures the proportion of demand satisfied within the required time window. The fifth indicator is lead-time responsiveness, which measures whether production and delivery occur quickly enough to support hospital needs.

Additional indicators include service availability, cost of emergency reallocation, quality compliance, traceability, cyber-disruption tolerance, and distribution equity. Quality compliance is essential because faster production must not compromise patient safety. Traceability is required to ensure that components and finished devices can be audited. Cyber-disruption tolerance measures whether the digital coordination layer can continue operating under attacks or platform failures. Distribution equity measures whether equipment allocation across hospitals or regions follows justified public health priorities.

3.9. Governance, safety, and implementation assumptions

The proposed framework assumes that emergency production is conducted under appropriate regulatory and quality-control oversight. Cloud manufacturing and AI-enabled allocation can support emergency production, but they cannot replace medical-device certification, safety testing, or clinical governance. WHO technical guidance [14] on ventilators for COVID-19 emphasized minimum requirements for safety, quality, and effectiveness. Therefore, each service included in the qualified service pool should satisfy predefined technical, regulatory, and traceability requirements.

The framework also assumes that decision-support systems are used to assist human decision makers rather than autonomously authorize unsafe production. Public health agencies, medical-device regulators, hospital systems, and production partners should define acceptable quality

thresholds, emergency authorization rules, and accountability structures before crisis conditions occur.

Finally, the framework assumes that preparedness is built before a crisis. A cloud manufacturing response cannot be fully improvised during an emergency. The most effective use of the framework would require pre-crisis mapping of potential service providers, pre-qualification of emergency production capacities, simulation-based stress testing, and agreement on fallback rules.

4. Discussion: Health-System and Emergency Preparedness Implications

4.1. Reframing emergency production as a health-system resilience capability

The proposed framework reframes emergency medical-equipment production as a health-system resilience capability rather than a purely industrial or procurement function. In conventional planning, medical-device availability is often treated as a supply-chain or purchasing issue. However, public health crises show that the availability of critical devices such as ventilators depends on the ability of the broader system to sense demand, mobilize capacity, coordinate heterogeneous actors, preserve digital continuity, and recover from disruptions. Therefore, emergency production should be understood as a socio-technical health-system function that links hospitals, public health agencies, manufacturers, laboratories, logistics providers, and digital coordination platforms.

This reframing has practical importance. If ventilator shortages are viewed only as market failures or procurement delays, the response is likely to focus on emergency purchasing, imports, or stockpiling. These strategies remain important, but they may be insufficient when global demand rises simultaneously and international supply channels become constrained. A resilience-oriented perspective adds another option: the ability to activate distributed and pre-qualified production capacity through a cloud manufacturing network.

The proposed framework therefore expands the concept of health-system preparedness. Preparedness should include not only inventory reserves and emergency contracts, but also a live map of convertible production capacities, pre-approved service providers, digital task-allocation mechanisms, alternative logistics routes, and fallback decision procedures. Such an approach can reduce dependence on single suppliers and improve the ability of the system to respond when demand exceeds ordinary production capacity.

4.2. Implications for hospitals and healthcare providers

Hospitals are the primary users of emergency medical equipment, but they are often downstream actors in production decisions. The proposed framework suggests that hospitals should be more directly connected to emergency production networks through structured demand signals. During a crisis, hospitals can provide information on equipment inventory, ICU occupancy, projected demand, device failure rates, and urgency levels. These data can be translated into an emergency demand profile that guides production and distribution priorities.

This connection can improve responsiveness. Instead of waiting for delayed procurement cycles, hospital demand data can trigger service composition decisions across a distributed production network. For example, if a region reports a rapid increase in ventilator demand, the framework can identify available production services, allocate component tasks, activate backup providers, and update distribution plans. This does not remove the need for public health coordination; rather, it provides a more operational link between clinical demand and production capacity.

However, hospitals should not be expected to manage manufacturing complexity directly. Their role should be to provide accurate demand data, define clinical urgency, participate in prioritization rules, and verify that delivered equipment satisfies safety and usability requirements. Manufacturing coordination should remain the responsibility of public health agencies, certified production partners, and regulated emergency-response structures.

4.3. Implications for public health agencies and emergency planners

Public health agencies can use the proposed framework as a preparedness-planning tool. Before a crisis occurs, agencies can identify and pre-qualify potential production services across sectors. These may include medical-device manufacturers, automotive firms, machining workshops, electronics producers, universities, military facilities, testing laboratories, and logistics providers. Each provider can be assessed according to capacity, technical capability, quality compliance, geographic location, cyber-dependence, and disruption exposure.

This preparation can support the creation of an emergency production registry. Such a registry would not be a simple contact list; it would be a structured database of capabilities, constraints, certification status, available equipment, production lead times, and potential task compatibility. When a crisis occurs, the registry can feed the cloud manufacturing service layer and allow rapid task-service matching. Without this preparation, emergency production networks may be built reactively, which increases delay, uncertainty, and quality risk.

Public health agencies can also use the framework to design simulation-based stress tests. For example, agencies can simulate a sudden ventilator demand surge, the failure of a major component supplier, a logistics disruption, or a cyberattack on the digital coordination platform. These exercises can identify fragile dependencies before a real crisis occurs. The outputs can guide investment in backup capacity, supplier diversification, digital redundancy, and emergency allocation templates.

4.4. Implications for medical-device supply networks

Medical-device supply networks are often optimized for efficiency, cost control, and lean inventory management. While these objectives are valuable under normal conditions, they can increase fragility during large-scale disruptions. The proposed framework suggests that medical-device supply networks should balance efficiency with resilience. This requires maintaining alternative capacity, mapping critical dependencies, and preserving the ability to reconfigure production under stress.

For ventilator production, some components may be more critical than others. Sensors, valves, control boards, oxygen-flow components, and calibration processes may create bottlenecks if only one qualified supplier is available. A resilience-oriented service composition framework can identify such bottlenecks and recommend backup services or alternative production routes. This is particularly important when different suppliers face different disruption risks.

However, distributed production creates coordination and quality challenges. Medical devices cannot be produced through informal task allocation without strict safety controls. Every service included in the production network must satisfy technical, regulatory, and traceability requirements. Therefore, resilience should not be interpreted as uncontrolled diversification. A resilient medical-device network is not one that uses any available capacity; it is one that uses pre-qualified, traceable, and governable capacity.

4.5. Digital and cyber-resilience implications

Emergency production networks increasingly depend on digital infrastructure. Cloud manufacturing platforms, supplier registries, hospital demand dashboards, allocation engines, quality records, and logistics systems are all part of the operational response. If these systems fail, production coordination may be delayed even when physical capacity exists. Therefore, cyber-resilience is a necessary component of health-system production resilience.

The proposed framework integrates cyber-continuity because digital disruption can affect both information flow and production decisions. A distributed denial-of-service attack may slow access

to allocation systems. Ransomware may compromise supplier data or production records. Cloud provider failures may prevent the execution of optimization models. Data drift may reduce the reliability of AI-based allocation recommendations. In each case, the production network needs the ability to continue operating safely.

Fallback orchestration is particularly important in this context. The Degrade, Substitute, and Bypass logic can be adapted to emergency medical production. A Degrade strategy may use simplified allocation rules when full optimization is unavailable. A Substitute strategy may switch to a backup model, backup database, or alternative service provider. A Bypass strategy may activate pre-approved emergency production templates when real-time digital systems cannot be trusted. These fallback strategies can prevent operational paralysis, but they must be bounded by quality and safety rules.

4.6. Governance, ethics, and regulatory considerations

The use of cloud manufacturing and AI-enabled service composition in emergency medical-equipment production raises important governance issues. The first issue is safety. Ventilators are life-support devices, and their production must comply with medical-device standards, testing protocols, and regulatory requirements. Emergency conditions may justify accelerated procedures, but they do not justify unsafe production. Any AI-enabled allocation framework must therefore operate within defined safety constraints.

The second issue is accountability. When production tasks are distributed across multiple providers, responsibility for quality, delay, failure, or defect can become unclear. Governance structures should define who approves service providers, who authorizes production assignments, who validates quality, and who is accountable for final device release. Without these rules, distributed production may increase legal and operational ambiguity.

The third issue is equity. During a public health crisis, equipment allocation can become ethically sensitive. If production is insufficient to satisfy all demand immediately, allocation rules should be transparent and justified. Criteria may include ICU pressure, shortage severity, regional vulnerability, patient volume, or public health risk. The proposed framework can support equitable distribution only if equity criteria are explicitly embedded in the health-system evaluation layer.

The fourth issue is data governance. Hospital demand data, supplier capacity data, logistics data, and production-quality data may be sensitive or commercially confidential. Secure data-sharing protocols are necessary to prevent misuse, leakage, or manipulation. In addition, AI-based

decision support should be explainable enough for emergency planners and regulators to understand why certain providers or allocation plans were selected.

The fifth issue is human oversight. AI-enabled allocation can improve speed and coordination, but emergency medical-device production should not be fully delegated to autonomous systems. Human decision makers should remain responsible for approving high-risk allocations, resolving conflicts between efficiency and equity, and ensuring compliance with medical-device regulations.

4.7. Practical implementation roadmap

A practical implementation of the proposed framework would require staged development. The first stage is capacity mapping. Public health agencies and industry partners should identify potential production services and classify them according to capability, certification status, capacity, lead time, and disruption risk. This creates the foundation for an emergency manufacturing registry.

The second stage is pre-qualification. Candidate services should be evaluated before a crisis. This includes technical assessment, quality-system verification, cybersecurity review, data-sharing readiness, and compatibility with medical-device production requirements. Pre-qualification reduces the delay associated with activating new providers during emergencies.

The third stage is digital platform development. A secure cloud manufacturing platform should connect demand data, service provider data, task decomposition, allocation logic, quality records, and logistics coordination. The platform should support interoperability with hospital systems and public health dashboards while preserving privacy and security.

The fourth stage is decision-support integration. Robust optimization, adaptive learning, generative allocation, and fallback rules can be integrated gradually. A practical system does not need to deploy the most advanced AI module at the beginning. It can start with rule-based allocation and robust optimization, then add adaptive and generative modules as data and governance maturity increase.

The fifth stage is simulation and stress testing. Emergency scenarios should be tested before deployment. These scenarios may include sudden demand surges, supplier failures, regional logistics disruption, cyberattacks, data delays, or simultaneous multi-region shortages. Stress testing can reveal hidden dependencies and help refine fallback procedures.

The sixth stage is crisis activation and continuous monitoring. During a real emergency, the system should update demand, capacity, allocation, and quality data continuously. Monitoring

should detect production bottlenecks, cyber anomalies, delivery delays, and quality deviations. The framework should then support reallocation, provider substitution, and fallback activation as needed.

4.8. Limitations and future research

This study is conceptual and case-based. It synthesizes prior technical mechanisms and translates them into an applied health-system resilience framework, but it does not provide a new empirical validation using real hospital demand data or certified medical-device production datasets. Therefore, the framework should be interpreted as a structured methodological contribution rather than a fully validated operational system.

A second limitation concerns generalizability. Ventilator production is a strong illustrative case because it combines urgent demand, technical complexity, strict quality requirements, and distributed production potential. However, other medical devices may have different production requirements, regulatory constraints, and supply-network structures. The framework may need adaptation for diagnostic equipment, oxygen systems, personal protective equipment, or pharmaceuticals.

A third limitation is data availability. Effective implementation requires reliable data on hospital demand, provider capacity, production lead times, component availability, quality status, and logistics constraints. In many health systems, these data are fragmented across organizations and may not be updated in real time. Without accurate data, AI-enabled allocation may produce recommendations that are technically feasible but operationally unrealistic.

Future research should validate the proposed framework using real or realistic emergency demand datasets. Hospital-level demand data, ICU occupancy data, ventilator inventory records, and production capacity data could be used to evaluate whether cloud manufacturing service composition improves shortage reduction, fulfillment rate, and recovery time during simulated crises. Another direction is the development of digital twins for emergency medical-equipment production. A digital twin could represent hospitals, production services, suppliers, logistics routes, quality-control nodes, and disruption scenarios. This would allow researchers and public health agencies to test allocation policies before crises occur.

5. Conclusion

Public health crises reveal that the availability of critical medical equipment is not only a procurement or manufacturing problem, but also a health-system resilience challenge. During

severe disruptions, conventional production and supply systems may be unable to respond quickly enough to demand surges for life-support devices such as ventilators. This study developed a case-based methodological framework that positions cloud manufacturing and intelligent service composition as mechanisms for improving emergency medical-equipment production continuity during public health crises.

The proposed framework connects hospital and public health demand signals to distributed manufacturing capacities through five layers: the health-demand layer, the cloud manufacturing service layer, the intelligent service composition layer, the resilience and cyber-continuity layer, and the health-system evaluation layer. By organizing emergency ventilator production as a service composition problem, the framework explains how production tasks can be decomposed, matched with qualified service providers, and reallocated under disruption. This structure allows non-traditional but qualified actors, including universities, military workshops, industrial manufacturers, laboratories, component suppliers, and logistics providers, to contribute to emergency production while remaining subject to quality, safety, and traceability requirements.

The framework also emphasizes that resilience cannot be reduced to efficiency or reliability. Emergency medical-equipment production requires redundancy to provide backup capacity, robustness to protect against uncertainty, flexibility to enable reconfiguration, collaboration to activate cross-sector resources, cyber-continuity to preserve digital coordination, and fallback orchestration to maintain decision-making when advanced systems fail. These mechanisms are particularly important in crisis conditions, where demand, capacity, logistics, and digital infrastructure may change rapidly.

The COVID-19 ventilator production case illustrates the practical relevance of this framework. Ventilator shortages during respiratory pandemics can directly affect intensive-care capacity and continuity of care. A cloud manufacturing approach can support emergency response by identifying qualified distributed capacities, assigning production tasks to suitable services, maintaining alternative production routes, and updating allocation decisions as disruptions evolve. In this way, emergency production becomes an organized health-system capability rather than an improvised response.

This study contributes to applied health science by translating technical advances in resilient cloud manufacturing, robust service composition, adaptive allocation, generative optimization, cyber-resilient migration, quantum-enhanced decision support, and fallback orchestration into a

health-oriented emergency production framework. The contribution is conceptual and methodological rather than empirical. It provides a structured basis for future operational models, simulation studies, digital twins, and public health preparedness tools.

Several limitations should be acknowledged. First, the study is case-based and conceptual; it does not validate the proposed framework using real hospital demand data, certified ventilator production datasets, or a new simulation experiment. Second, the framework is illustrated through ventilator production, and adaptation may be required for other medical devices, pharmaceuticals, or diagnostic technologies. Third, implementation depends on the availability of accurate data on hospital demand, supplier capacity, production lead times, regulatory status, logistics constraints, and cyber risks.

Future research should empirically test the framework using realistic emergency scenarios and operational datasets. Simulation-based studies can compare robust optimization, reinforcement learning, generative allocation, and fallback-based decision rules under demand surges, supplier failures, logistics disruptions, cyberattacks, and combined crises. Digital twins of emergency medical-equipment production networks can further support preparedness planning by allowing public health agencies to stress-test production capacity before crises occur. Future studies should also integrate regulatory approval pathways, quality assurance, traceability, and equity-based allocation rules directly into service composition models.

Overall, the proposed framework shows that cloud manufacturing and intelligent service composition can support health-system resilience when they are governed as safety-critical, traceable, and human-supervised decision-support mechanisms. Emergency medical-equipment production should not depend solely on centralized capacity or reactive procurement. A prepared health system should be able to sense demand, activate distributed qualified capacity, coordinate production tasks, preserve digital continuity, and recover from disruption while maintaining patient-safety standards.

6. References

[1] Ranney, M. L., Griffeth, V., & Jha, A. K. (2020). Critical supply shortages - The need for ventilators and personal protective equipment during the Covid-19 pandemic. New England Journal of Medicine, 382(18), e41. https://doi.org/10.1056/NEJMp2006141

[2] 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

[3] Shahab, E., Rabiee, M., Mobasseri, 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

[4] Shahab, E., Taleb, M., Gholian-Jouybari, F., & Hajiaghaei-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

[5] 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

[6] 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

[7] 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

[8] Emanuel, E. J., Persad, G., Upshur, R., Thome, B., Parker, M., Glickman, A., Zhang, C., Boyle, C., Smith, M., & Phillips, J. P. (2020). Fair allocation of scarce medical resources in the time of Covid-19. New England Journal of Medicine, 382(21), 2049-2055. https://doi.org/10.1056/NEJMsb2005114

[9] Miller, F. A., Young, S. B., Dobrow, M., & Shojania, K. G. (2021). Vulnerability of the medical product supply chain: The wake-up call of COVID-19. BMJ Quality & Safety, 30(4), 331-335. https://doi.org/10.1136/bmjqs-2020-012133

[10] Spieske, A., Gebhardt, M., Kopyto, M., & Birkel, H. (2022). Improving resilience of the healthcare supply chain in a pandemic: Evidence from Europe during the COVID-19 crisis. Journal of Purchasing and Supply Management, 28(5), 100748. https://doi.org/10.1016/j.pursup.2022.100748

[11] Organisation for Economic Co-operation and Development. (2024). Securing medical supply chains in a post-pandemic world. OECD Health Policy Studies. OECD Publishing. https://doi.org/10.1787/119c59d9-en

[12] 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

[13] Shahab, E., & Taghipour, S. (2026). Quantum reinforcement learning for resilient cloud service composition. Soft Computing. https://doi.org/10.1007/s00500-025-11052-6

[14] World Health Organization. (2020). Technical specifications for invasive and non-invasive ventilators for COVID-19: Interim guidance, 15 April 2020. World Health Organization.