International Journal of Engineering & Scientific Research
Vol.13 Issue 12, December 2025
ISSN: 2347-6532 Impact Factor: 8.076
Journal Homepage: http://www.ijmra.us, Email: editorijmie@gmail.com
Double-Blind Peer Reviewed Refereed Open Access International Journal - Included in the International Serial Directories Indexed & Listed at: Ulrich's Periodicals Directory ©, U.S.A., Open J-Gage as well as in Cabell's Directories of Publishing Opportunities, U.S.A

# Building a Resilience Fabric in Public Universities: Integrating Continuity, IT DR, and Emergency Response

Author: Mallesh Miryala
University of California
Email: miryalaca@gmail.com

## Abstract

Public universities behave like small cities. They run classrooms, research labs, hospitals, housing, transportation, utilities, and enterprise IT—often across multiple campuses and clinical sites. When disruption occurs, impacts rarely stay inside one department. A network outage can cancel lectures, delay clinical orders, and block payroll in the same afternoon.

Many institutions still manage continuity, IT disaster recovery (ITDR), and emergency response through disconnected spreadsheets, PDFs, and email threads. This paper proposes a Resilience Fabric for public universities: a human-centered system that connects plans, people, applications, infrastructure, and facilities into one coherent lifecycle. The concept is grounded in practical experience leading continuity and ITDR efforts across a multi-campus public university environment.

The paper introduces (1) a layered architecture that integrates continuity, ITDR, and emergency planning, (2) a continuity knowledge graph that maps operational dependencies, (3) an incident-to-learning pipeline turning real events into verified improvements, (4) persona-based workflows for planners, technical staff, facilities, and leadership, and (5) a minimal set of metrics that executive teams can understand and trust. Seven figures provide visual patterns that other public universities can adapt to their own tools and governance models.

**Keywords:** public universities, resilience, business continuity, IT disaster recovery, emergency management, facilities readiness, knowledge graph

## 1. Introduction

Public universities must deliver teaching, research, and public service under significant constraints: limited budgets, legacy systems, decentralized decision-making, and highly visible obligations to students and communities. At the same time, their risk surface keeps expanding—cyber incidents, climate-driven events, infrastructure failures, and vendor outages.

Most institutions have some continuity artifacts:

- Business continuity plans for priority units
- Disaster recovery runbooks for selected applications
- Emergency response procedures for campuses and buildings
- Facilities contingency plans for utilities and critical spaces

However, these artifacts often live in different tools, use different vocabulary, and follow different life cycles. When an incident occurs, staff spend precious time hunting for contacts, dependencies, and procedures scattered across drives, wikis, and ticketing systems.

This paper argues that public universities need something more integrated and more human-friendly: a Resilience Fabric that weaves together continuity, ITDR, emergency response, and facilities readiness. The fabric is not a single product; it is a design, governance, and measurement approach that can be implemented on top of existing platforms.

The author draws on hands-on experience designing and rolling out such a fabric across a multi-campus public university system that includes academic, research, and clinical operations.

## 2. Concept of a Resilience Fabric

A Resilience Fabric is a connected set of people, processes, and data elements that together answer three practical questions:

1. **What matters most?**
   Which services, applications, facilities, and processes are truly mission-critical?
2. **What do they depend on?**
   How do business functions rely on technology, facilities, vendors, and people?
3. **How ready are we, today?**
   Are plans current, dependencies accurate, exercises executed, and corrective actions closed?

Rather than treating business continuity, ITDR, and emergency response as separate initiatives, the fabric:

- Encodes dependencies in a shared knowledge graph.
- Uses common workflows for planning, exercises, incidents, and corrective actions.
- Presents role-specific views to planners, technical staff, facilities, and leadership.

shows a layered view of the fabric.

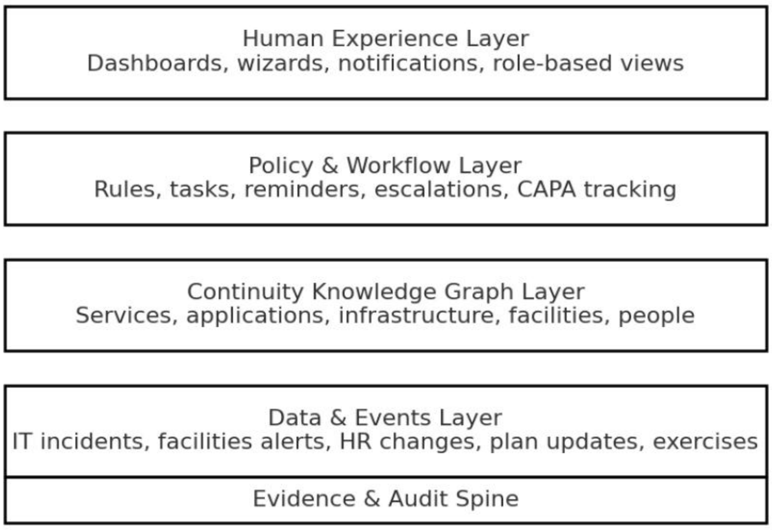

Human Experience Layer
Dashboards, wizards, notifications, role-based views

Policy & Workflow Layer
Rules, tasks, reminders, escalations, CAPA tracking

Continuity Knowledge Graph Layer
Services, applications, infrastructure, facilities, people

Data & Events Layer
IT incidents, facilities alerts, HR changes, plan updates, exercises

Evidence & Audit Spine

## 3. Layered Architecture

### 3.1 Overview

The Resilience Fabric is modeled as four vertical layers and one horizontal "spine." Each layer can be implemented using different technologies; the important part is the separation of concerns and the flow of information.

**Layered Resilience Fabric for public universities, from raw operational events at the bottom through knowledge, policy, and human interfaces at the top, with an evidence and audit spine underpinning all layers.**

### 3.2 Data & Events Layer

This layer listens to what is already happening across the institution:

- Service desk incidents and IT change records
- Building alarms and facilities work orders
- HR events such as new hires, role changes, departures
- Plan creation, review, exercise schedules, and attestations

The goal is not to centralize all data but to capture enough context so that continuity-relevant events can trigger tasks, updates, and metrics.

### 3.3 Continuity Knowledge Graph Layer

At the center of the fabric lies a continuity knowledge graph—a structured representation of how functions, applications, infrastructure, facilities, and people relate to each other (Section 4).

### 3.4 Policy & Workflow Layer

Policies define what "good" looks like:

- How often plans must be reviewed.
- What level of dependency detail is required for critical services.
- When exercises are mandatory.
- How corrective actions (CAPA) must be documented and verified.

These policies are implemented as workflows that generate tasks, reminders, escalations, and attestations—always referencing the rule that triggered them.

### 3.5 Human Experience Layer

Users never see the underlying model directly. Instead, they interact through:

- Simple wizards for creating or updating plans.
- Dashboards showing readiness and exceptions.
- Runbook views tailored to technical teams.
- Notifications through email, chat, or SMS.

The human experience layer must respect the reality that most planners and subject matter experts have very limited time for continuity work.

### 3.6 Evidence & Audit Spine

Across all layers, the system records:

- Who performed which action, when, and for what reason.
- Which policy or event generated each task.
- Evidence attachments for exercises, incidents, and CAPA closure.

This spine supports internal reviews, external audits, and accreditation visits without requiring last-minute document hunts.

## 4. Continuity Knowledge Graph

Traditional continuity efforts often rely on spreadsheets listing "system name," "owner," and "RTO." Such lists quickly become outdated and do not show how services depend on each other.

A continuity knowledge graph instead models:

- **Business Functions:** teaching, research, clinical services, student support, administration.
- **Applications:** learning management systems, student information systems, EMR, ERP, collaboration platforms.
- **Shared Services:** identity, messaging, integration platforms, data platforms.
- **Infrastructure:** compute, storage, network, cloud services, backup sites.
- **Facilities:** data centers, classrooms, labs, hospitals, utilities.
- **People & Roles:** on-call engineers, continuity coordinators, incident commanders.

Every node relates to others through edges such as "depends on," "hosts," "staffed by," or "located in."

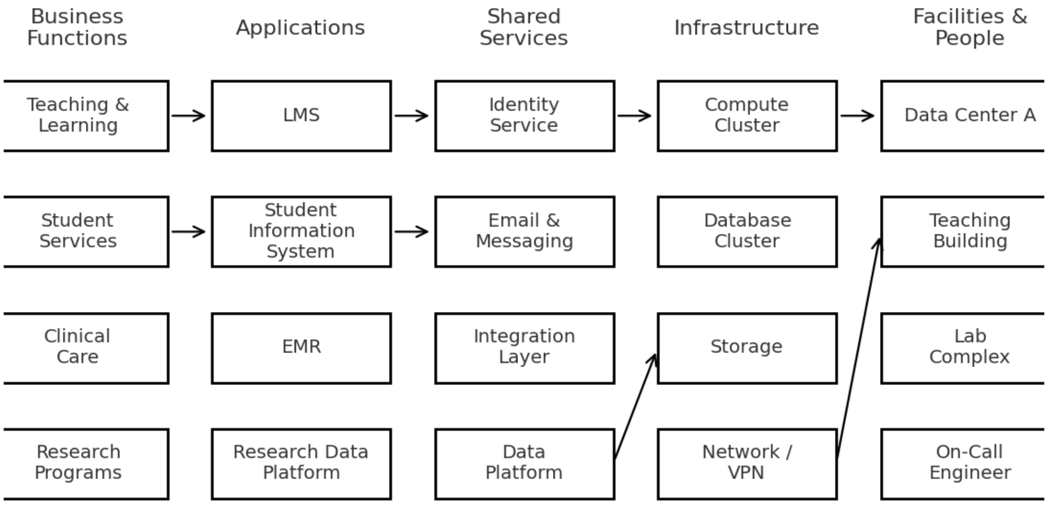

**Continuity knowledge graph for public universities, showing how business functions depend on applications, shared services, infrastructure, facilities, and people.**

## 5. Incident-to-Learning Pipeline

Many universities conduct exercises and endure real incidents, yet learning remains local and temporary. Email threads capture "lessons learned," but corrective actions are not tracked to closure.

The Resilience Fabric defines an incident-to-learning pipeline:

1. **Event or Exercise** – something happens: a network outage, building flood, tabletop, or failover test.
2. **Observation Capture** – participants record what worked, what failed, and what surprised them.
3. **Corrective Action (CAPA)** – specific actions are created with owners and due dates.
4. **Verification** – evidence is attached and reviewed; CAPA is accepted or re-opened.
5. **Metrics Update** – dashboards update automatically as actions close or remain overdue.

Each observation and CAPA is linked back to the continuity knowledge graph, so improvements remain tied to the services and facilities they affect.

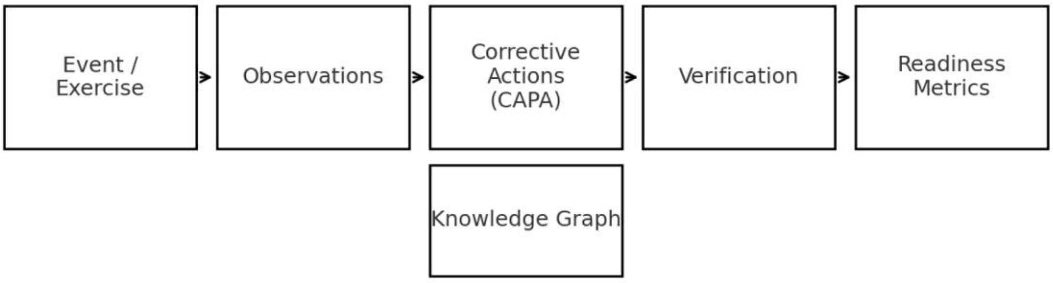

**Incident-to-learning pipeline converting events and exercises into observations, corrective actions, verification, and updated readiness metrics.**

---

## 6. Personas and Human Workflows

Real-world success depends on how well the system fits people's jobs. Experience in a multi-campus public university setting suggests four core personas:

1. **Department Planner** – manages continuity for an academic or administrative unit.
2. **Technical Owner** – responsible for one or more applications or shared services.
3. **Facilities & Safety Coordinator** – manages building information, hazards, and emergency response details.
4. **Executive & Risk Leadership** – needs systemwide visibility and assurance.

Each persona interacts with the fabric differently yet shares the same underlying data.

| | | | | | |
|---|---|---|---|---|---|
| Department Planner | Plan Creation | Dependency Modeling | Exercise / Test | After-Action Review | Attestation & Reporting |
| Technical Owner | Plan Creation | Dependency Modeling | Exercise / Test | After-Action Review | Attestation & Reporting |
| Facilities & Safety Coordinator | Plan Creation | Dependency Modeling | Exercise / Test | After-Action Review | Attestation & Reporting |
| Executive & Risk Leadership | Plan Creation | Dependency Modeling | Exercise / Test | After-Action Review | Attestation & Reporting |

**Persona-centered workflow map showing how planners, technical owners, facilities and safety staff, and leadership participate in the planning, exercise, and review cycle.**

## 7. Planning and ITDR Lifecycle

The Resilience Fabric introduces a consistent lifecycle for units and systems, regardless of campus:

1. **Scope & Ownership** – identify which functions, systems, and facilities are in scope and who owns them.

2. **Modeling** – capture impacts, recovery objectives, and dependencies in the knowledge graph.
3. **Strategy Alignment** – align ITDR strategies, facility contingencies, and manual workarounds with business priorities.
4. **Exercise & Scenario Testing** – run technical and tabletop tests against realistic scenarios.
5. **Corrective Action & Attestation** – document CAPA, verify closure, and collect leadership attestations.

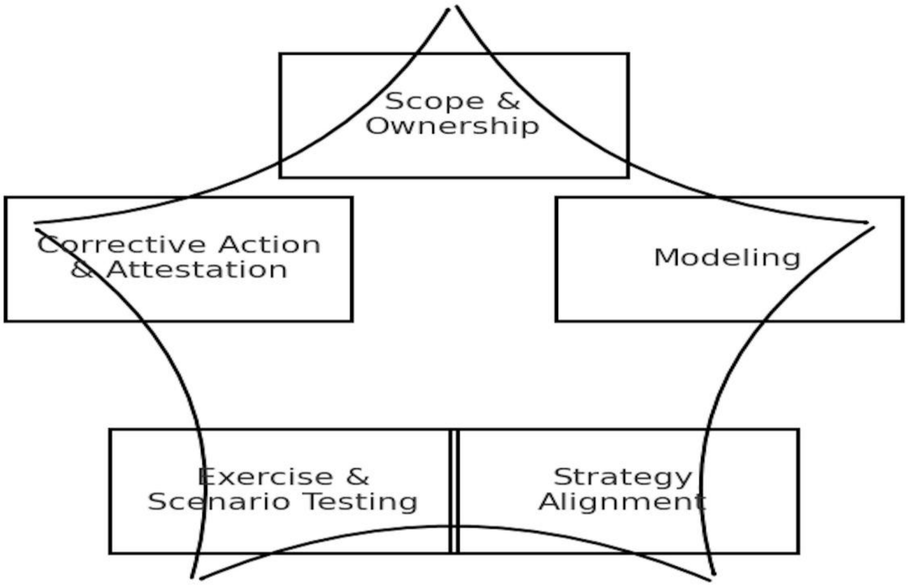

**Planning and IT disaster recovery lifecycle for public universities, forming a continuous loop from scope and ownership through modeling, strategy alignment, exercises, and corrective action with attestation.**

## 8. Measurement and Dashboards

University leaders do not need dozens of metrics; they need a small, stable set that can be monitored over time. Experience suggests the following four are both practical and powerful:

1. **Plan Currency (%)**
   Percentage of in-scope plans that are updated and attested within their defined review window and have no overdue CAPA.
2. **Dependency Coverage (%)**
   Percentage of critical services whose continuity chain (business function → app → shared services → infrastructure → facilities → people) is fully documented.
3. **CAPA Closure Time (days)**
   Median time from CAPA creation to verified closure.

4. **Attestations On-Time (%)**
   Percentage of required attestations completed before their due date.

Together, these metrics provide a concise view of readiness and improvement pace.

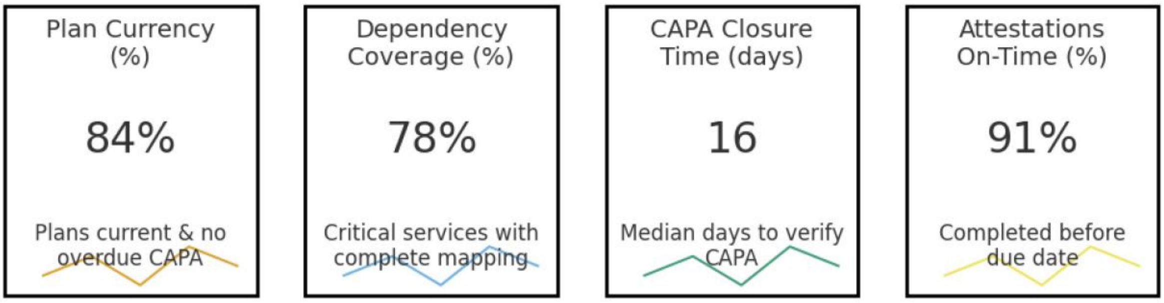

**Example leadership dashboard for the Resilience Fabric, highlighting plan currency, dependency coverage, CAPA closure time, and on-time attestations with simple trends.**

## 9. Governance and Release Management

Resilience work competes with many other institutional priorities. A light but disciplined governance model helps maintain progress without overloading participants.

A practical structure includes:

- **Resilience Steering Committee**
  Representatives from academic affairs, IT, facilities, emergency management, risk, and key campuses. Sets priorities, approves standards, and reviews metrics.
- **Configuration & Change Forum**
  Technical and process owners who design changes to workflows, forms, and dashboards. They plan releases and coordinate testing.
- **Local Champions Network**
  Continuity coordinators embedded in colleges, departments, or hospitals who help colleagues adopt the fabric and provide feedback.

Release management follows a predictable rhythm:

- **Quarterly Releases** with:
  - One-page change notes written in plain language.
  - Two-minute explainer videos.
  - Updates to FAQs and training decks.
- Releases deliberately avoid peak academic periods and known clinical "blackout dates."

## 10. Phased Rollout in Public Universities

Implementing a Resilience Fabric across a decentralized public university system is best done in stages. A realistic roadmap looks like this:

1. **Phase 1 – Baseline Discovery (Months 0–2)**
   - Inventory existing plans, ITDR documentation, and emergency procedures.
   - Identify critical services and facilities across campuses.
   - Capture "as-is" metrics, even if incomplete.
2. **Phase 2 – Design and Pilot (Months 3–4)**
   - Select a small set of diverse pilot units (for example, a central IT service, a research lab, and a student services office).
   - Build initial continuity knowledge graph for pilot scope.
   - Configure workflows for planning, exercises, incidents, and CAPA.
   - Run at least one tabletop and one technical failover through the new pipeline.
3. **Phase 3 – Evaluate and Refine (Month 5)**
   - Collect structured feedback from all personas.
   - Simplify forms, reduce redundant fields, and adjust notifications.
   - Fix modeling patterns that proved confusing.
   - Finalize metric definitions.
4. **Phase 4 – Systemwide Expansion (Months 6–12)**
   - Onboard units by category (e.g., all student-facing services, then research, then administrative).
   - Reuse proven modeling patterns and runbook templates.
   - Publish comparative dashboards by campus or unit type.
5. **Phase 5 – Optimization and Integration (Months 12+)**
   - Integrate the fabric with enterprise systems (HR, ITSM, building management, monitoring).
   - Expand to third-party dependencies and inter-institutional collaborations.
   - Use KPI trends to guide investments in infrastructure, training, and staffing.

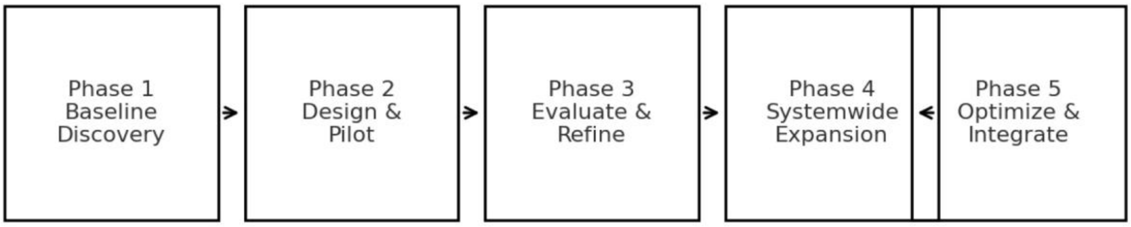

**Phased rollout roadmap for the Resilience Fabric, from discovery and pilot design through systemwide expansion and ongoing optimization.**

## 11. Discussion: Benefits and Challenges

### 11.1 Benefits

Implementing a Resilience Fabric in public universities offers several tangible benefits:

- Faster, better decisions in crises – leaders can see what is affected, which dependencies matter, and what recovery sequence makes sense.
- Reduced duplication of effort – continuity, ITDR, and emergency management share the same data and workflows.
- Stronger evidence for accreditation and audits – the evidence spine records exercises, incidents, and CAPA closure automatically.
- Greater ownership – clear personas, tasks, and metrics help departments and services understand their responsibilities.

### 11.2 Challenges

Challenges include:

- Data quality and backlog – legacy plans and system lists may be incomplete or inconsistent.
- Change fatigue – staff may be tired of new tools and forms.
- Customization pressure – campuses and units may want different fields or workflows, risking fragmentation.

Mitigation strategies include focusing on critical services first, using small pilots, and protecting a governed core schema while allowing limited, carefully reviewed extensions.

## 12. Conclusion

Public universities will continue to face complex risks that cut across campuses, disciplines, and technologies. Spreadsheets and isolated documents cannot keep pace with evolving threats, staff turnover, and the growing interdependence of digital and physical infrastructure.

The Resilience Fabric proposed in this paper offers a practical, human-centered way forward. By combining a layered architecture, a continuity knowledge graph, an incident-to-learning pipeline, persona-based workflows, clear metrics, and phased rollout, public universities can transform continuity from a compliance project into an operational capability.

Although the specific tooling can vary, the underlying patterns are reusable. Any public institution that recognizes itself in the challenges described here can adapt these ideas to build its own fabric—and in doing so, make its teaching, research, and public service more resilient for the communities it serves.

## References

1. Rasiah, R., Kaur, H., & Guptan, V. (2020). *Business continuity plan in the higher education industry: University students' perceptions of the effectiveness of academic continuity plans during COVID-19 pandemic*. Applied System Innovation, 3(4), 51. https://doi.org/10.3390/asi3040051 MDPI
2. EDUCAUSE. (2024). *Business continuity and disaster recovery toolkit: What is business continuity and disaster recovery?* EDUCAUSE Working Group Paper. EDUCAUSE
3. Yang, C.-L., Huang, C.-Y., Kao, Y.-S., & Tasi, Y.-L. (2017). *Disaster recovery site evaluations and selections for information systems of academic big data*. Eurasia Journal of Mathematics, Science and Technology Education, 13(8), 5151–5173. Eurasia J. Math. Sci. Tech. Educ.
4. EDUCAUSE Center for Applied Research. (2007). *Shelter from the storm: IT and business continuity in higher education – key findings*. EDUCAUSE. EDUCAUSE Library
5. **University of California, Office of the President. (n.d.). *UC Ready: Continuity planning – Business/Mission Continuity program*.** Retrieved from the Enterprise Risk and Resilience site: UC Ready is a systemwide, web-based continuity planning tool supporting departmental, application, and enterprise-level continuity and IT disaster recovery across UC campuses and medical centers. University of California+1
6. University of North Carolina at Chapel Hill. (n.d.). *Mission continuity*. Emergency Management and Planning Program. Emergency Management and Planning
7. Yale University. (n.d.). *Develop a plan – Business continuity planning*. Yale Emergency Management. emergency.yale.edu
8. Temple University. (n.d.). *Continuity planning*. Office of Emergency Management, Campus Operations. Campus Operations
9. Cornell University. (n.d.). *Continuity and recovery*. Cornell Emergency Management. Office of Emergency Management
10. Info-Tech Research Group. (2023). *Develop a business continuity plan for higher education*. Info-Tech Research. Info-Tech Research Group