ISSN: 2473-4901
v11 n6303

Teaching Case – Student Materials

Our Interconnected, Resilient, Modern World ... is Still Painfully and Remarkably Vulnerable

Paul Witman
witman@ieee.org

Jim Prior
jr2dg.consulting@gmail.com

California Lutheran University
Thousand Oaks, CA USA

Hook

Systems of all kinds are pervasive in everyday life. And when systems fail, often due to a lack of resilience, many aspects of daily life, including life-critical activities, may be impacted or destroyed. As future technologists and business leaders, understanding the importance and mechanisms of resilience is critical to ensuring that your organization’s systems are appropriately resilient, and that trade-offs are appropriately considered.

Abstract

In July of 2024, CrowdStrike, an IT security company, delivered updates to millions of Windows PCs around the world. A flaw in that update caused millions of Windows PCs to fail the restart process, disrupting or halting the global operations of airlines, banks, medical facilities, and much more. Many PCs were unable to reboot without direct human intervention. This resulted in delays that impacted on the recovery time for all types of businesses, as well as their employees, supply chains, and consumers. What went wrong? What was the root cause? Why didn’t we see this coming? And what can we learn from this experience that might help us help technology leaders and their business counterparts to further mitigate the risk of any similar events?

Keywords: Resilience, risks, complexity, trade-offs

ISSN: 2473-4901
v11 n6303

Our Interconnected, Resilient, Modern World ... is Still Painfully and Remarkably Vulnerable

Paul Witman and Jim Prior

1. INTRODUCTION AND OPENING STORY

The trigger for the case

On July 19, 2024, Paul (one of the authors) was traveling by plane to visit family. On arrival at the airport, he went to check the monitor near the entrance to find his flight's gate information and departure status. Rather than the usual list of flights, what he found appeared to be the Microsoft™ "blue screen of death."

He had learned of this likely outcome while en route to the airport. Fortunately, his airline's app was still able to provide status and schedule updates, and there were relatively small impacts on his flight. Other travelers were not so fortunate – other airlines cancelled thousands of flights in total, over several days, and had to scramble to recover their technology operations. They also had to communicate with delayed and disrupted passengers, track luggage, reschedule crews, and service and maintain aircraft, much as Southwest had to do during its earlier winter crisis (Witman et al., 2024). Many of those functions were dependent on Microsoft technology and were not immediately recovered.

Paul was fortunate that his travels were relatively unimpacted, the impact on the broader traveling public was highly visible turmoil, with passengers stuck in their departure or an intermediate airport, airlines unable to accurately predict "return to normal," and all of the fallout that comes from that.

Learning objectives

After working through this case study and its exercises, students should be able to:

- Define and describe system resilience and brittleness
- Identify potential causes of system brittleness and potential solutions or mitigations.
- Identify and analyze trade-offs of increasing resilience
- Identify and analyze costs and benefits of resilience testing

Case application, organization

This case has been used in introductory information systems courses for both graduate and undergraduate courses. Course applications focus on the role of system resilience in many and varied ways. Shocks to a business can come from technological failures, pricing changes, regulatory issues, and more (Rose, 2025). We have used technically-oriented case vignettes to illustrate various resiliency and brittleness issues, and have provided some broader follow-up research questions.

The remainder of this paper goes on to discuss background information on the technical concepts related to system resilience in Section 2, followed by a number of sample scenarios where resilience played a role in section 3. Each of those scenarios includes discussion questions for use by small groups or whole-class discussions or homework. Following the scenarios, Section 4 offers additional research questions for potential follow-on assignments, and a summary and conclusion follow in Section 5.

2. TECHNICAL CONCEPTS AND TERMS

There are a number of technical concepts and elements that apply to various aspects of these case scenarios, and we identify and describe them briefly here. Key concepts include system resilience, and its converse, system brittleness. Underlying those concepts are elements that affect systems, sometimes making them more or less resilient. These elements include but are not limited to system complexity, challenges that impede resilience testing, automatic updates, and pervasive client software.

System Resilience was defined by Google AI as the "ability to maintain its essential functions and capabilities when facing disturbances, disruptions, or adverse events." This definition clearly leaves a lot of room for flexibility, including a need to define which functions are "essential", and whether doing those functions in a degraded way (more slowly, at higher cost, etc.) might also be treated as a failure. Fundamentally, we believe that the concept of system resilience is one that must be defined for each system, and agreed to by each stakeholder in a system's functions.

System Brittleness (or Fragility) is to some extent the polar opposite of resilience. It speaks to a state of being vulnerable to one or more kinds of disruptions or adversity, both alone and simultaneously. Often, adversity and disruptions come in bursts - weather problems in the travel industry will force schedule changes, which will affect staffing, and will also cause bursts of customer service activity to re-plan travels (Witman et al., 2024).

Complexity

Baccarini (1996) defines complexity in projects as "consisting of many varied interrelated parts" and distinguishes between two critical types that directly impact system resilience. Organizational complexity relates to interconnections among organizational units, stakeholders, and procedures — the human and procedural elements that must coordinate for system success. Technological complexity concerns the number and diversity of technologies and their interdependencies. This dual framework remains highly relevant for understanding modern information systems, where both organizational factors and technical architectural decisions contribute to system vulnerability and resilience challenges.

The complexity challenge extends beyond Baccarini's framework, as modern systems operate within intricate webs of external dependencies that multiply complexity exponentially. Baskerville et al (2018) note that complexity arises from "the number and diversity of components and the nature of interconnections among them, leading to emergent and often unpredictable behaviors." This complexity is amplified by dependencies on external providers outside the system owner's direct control: power grid operators, telecommunications providers, cloud computing platforms, and facilities maintenance contractors, among others. These effects are further exacerbated by process complexity, the intricate nature of workflows, handoffs, poor process management, and the like (Shi et al., 2024).

Each external dependency introduces its own complexity and creates additional interdependencies and potential failure points. This multi-layered complexity — combining internal organizational and technological factors with external provider dependencies — creates emergent behaviors that can be difficult to predict or control, fundamentally challenging traditional approaches to system resilience planning.

Challenges to testing resilience

Active, "invasive" testing on running systems presents significant operational and strategic challenges for organizations seeking to validate their resilience capabilities. Fault injection testing involves "the deliberate introduction of errors and faults to a system to validate and harden its stability and reliability" with the goal of improving system design for resiliency under intermittent failure conditions. However, such invasive testing on production systems carries inherent risks — testing that introduces various disruptions may induce behaviors that degrade or destroy system functionality, necessitating repair and recovery operations that can be costly and disruptive. Production fault injection should be considered one of many approaches used to gain confidence in the safety and resiliency of a system, similar to unit testing and code review, though it is limited in which surprising events it can prevent. To mitigate these risks, system owners often schedule "planned outages" to allow for controlled testing scenarios, but tolerance for even planned outages decreases significantly as system criticality increases.

Organizations like the University of Wisconsin Hospital spend weeks preparing and testing before planned instances of downtime to help things run as quickly and smoothly as possible. However, they recognize that the balance between resilience validation and operational continuity becomes increasingly delicate for mission-critical systems where downtime can have severe consequences for users, revenue, public safety, and more (Otkhozoria et al., 2025).

Cloud Computing and other shared services

Cloud Computing generally refers to contracting with a vendor to operate some portion of an organization's technological needs in a remote data center, sharing that infrastructure with other clients. The cloud provider handles infrastructure operations including power, cooling, monitoring, and security. Such providers are often selected based on their ability to operate at a higher level of quality than clients could otherwise achieve. While this upgrade in operational capability is usually real, it does not eliminate the risk inherent in running complex technological systems, and cloud providers can still experience failures in both primary and backup operations (Alozie et al., 2024).

Pervasive client software

"Pervasive client software" refers to software that must be installed across many or all computing devices in an organization. A healthcare industry CIO used this term to describe software deployed on everything from desktop computers to embedded systems like airport display controllers or ATM operating systems. The critical risk with pervasive software is that a single software failure can simultaneously disable multiple systems across the organization. This creates a "single point of failure" scenario, as demonstrated during the July 2024 CrowdStrike incident when a faulty update disrupted operations globally.

Remote and automated updates

Enterprise IT teams managing thousands of devices face a complex challenge balancing two competing risks: the risk of deploying a faulty update that could crash multiple systems simultaneously, versus the risk of delaying critical security patches that leave the organization vulnerable to cyberattacks. This creates a fundamental trade-off in system resilience (Usman & Asplund, 2025).

For example, Carol, the CIO of a large healthcare organization, worked with her Chief Information Security Officer to develop an update deployment policy. Most routine updates underwent thorough testing before organization-wide deployment. However, for a select group of highly trusted vendors, updates marked as "critical and urgent" were deployed immediately to the most at-risk systems. This policy aimed to minimize vulnerability windows but increased exposure to potential update failures.

Root Cause Analysis

Root cause analysis refers to the process of attempting to discern the most fundamental, or root, cause of a problem. It is often conducted by repeatedly asking "why" after identifying a partial cause. The goal is to focus on solving the fundamental cause rather than addressing symptoms (Rooney & Heuvel, 2004).

For example, analyzing system failures: Why did the system experience widespread service disruptions? Because an external shock overwhelmed the system's ability to manage operational complexity. Why? Because the system was not designed to deal with such disruption levels. Why? … A related Why? might be a question of why the organization allowed a system in production without sufficient resilience testing. Often, the answer to that is some combination of pressures to reduce costs and/or time to market (TBD Citation).

This analysis reveals the real cause: lack of investment in resilience-building measures left the system with dangerous brittleness and vulnerability to cascading failures.

Risk management

Risk refers to the probability that a certain type of event will occur, while impact is the magnitude of damage that might result. There are several fundamental approaches to managing risk (Ahmed, 2017).

- **Avoidance**: Avoiding activities that create unacceptable risk
- **Transfer**: Moving risk to another party (insurance, contracts)
- **Reduction**: Taking steps to reduce risk to acceptable levels
- **Acceptance**: Accepting remaining "residual risk" after mitigation

Sometimes, organizations take an approach that externalizes some part of their risk. This means that they (deliberately or inadvertently) transfer some of their risk to external parties - customers, vendors, employees, supply chain partners, etc. In brittle systems, organizational choices for resilience investment essentially transfer some of the impact of these risks to various other parties, also known as creating a trade-off.

3. SCENARIOS: SYSTEM RESILIENCE FAILURES

3.1 CrowdStrike scenario completion

Getting through a US airport is always an engaging experience, and one guided by information made available to the travelers and to security and other staff. While going through the airport was relatively easy for Paul (his airline's systems were generally OK, and the airline’s app on his phone provided correct gate and flight information), other travelers were not as fortunate.

Security protocols limit the secure area of the airport to those with tickets for active flights. With many flights being canceled, it was difficult for security staff to be sure who they could properly admit. Customer service staff from all airlines, both at the airport and in call centers, were working hard to help passengers re-arrange their travel plans. However, the scope of the problems was so broad as to create significant queues of people that were impacted by the technology issues. These customers added greatly to the delay faced by customers dealing with “normal” travel challenges, forcing airlines to rapidly ramp up on-site and call-center staffing to help rearrange travel plans.

Overall, the airline industry alone suffered many billions of dollars of damages due to the CrowdStrike issue. One airline, Delta, reported its own costs at over USD500 million, with impacts on revenues, flight crews, manual rescheduling, and compensation to affected customers (Cerullo, 2024).

In brief, the underlying technical problem that caused the outage to occur was a faulty software update released by CrowdStrike for its cybersecurity software that caused 8.5 million Windows computers to crash. CrowdStrike is a cybersecurity company whose software runs on millions of computers worldwide to protect them from cyber threats. The faulty configuration update caused an error in the computer's memory, resulting in the infamous "Blue Screen of Death" that made affected computers unusable.

The impact was massive because CrowdStrike's software operates at a very deep level within computer systems, and when it crashed, it took the entire computer down with it. Airlines, healthcare systems, media companies, and organizations all over the world were unable to operate properly. The computers affected were unable to recover without direct manual intervention, meaning IT teams had to physically access each machine to fix the problem. While CrowdStrike's CEO confirmed this was not a cyberattack and deployed a fix, the incident highlighted how a single software error from one company could bring down critical infrastructure worldwide (Kerner, 2024).

Questions:

- Of the technical factors around the topic of system resilience, documented in Section 2, which seems most applicable to this situation?
- Given the role of strong competitors to become a dominant presence within many technology operations (e.g., iOS, Windows, CrowdStrike, etc.), what can an organization do to protect themselves from this sort of incident?
- Even if a company does protect itself from such risks, what can they do to protect their supply chain against risks among their vendors?
- What other questions or observations can you make about this scenario as it relates to resilience, risk management, information, and information systems?

3.2 Ninety seconds of terror

Air traffic controllers have very high stress jobs. With little room for error, they simultaneously coordinate the movement of multiple aircraft carrying hundreds of people. They do so for extended periods of time, making critical decisions with often incomplete and/or rapidly changing information. They communicate via radio directly with pilots and they monitor aircraft via radar.

Picture an air traffic controller in a radio conversation with a pilot bringing a plane in for landing at one moment. And then, in the next moment, the controller loses radio contact with the pilot, and their radar screen goes blank.

This happened to the air traffic controllers at the Philadelphia Pennsylvania Terminal Radar Approach Control (TRACON) center, which is responsible for separating and sequencing planes in and out of Newark Liberty International Airport in New Jersey. The outage – no radio, no radar – lasted for ninety seconds.

Fortunately, there were no serious incidents during the blackout. Pilots have procedures they follow under circumstances like these, and once radio communications were restored, the air traffic controllers and pilots quickly recovered, despite the stress involved.

But the follow-on impact was profound. A number of the controllers on duty during the outage took time off due to the stress and trauma of the event. The Newark airport was shut down for two hours while the Federal Aviation Administration (FAA) worked on the issue. Over 65 flights were diverted to other airports, 160 flights were canceled and 424 delayed. Those problems continued into the next day with more than 370 flights delayed and more than 500 cancelled (Martínez & Dumas, 2025). Additional outages at Newark prompted the FAA to reduce the number of departures and arrivals per hour by 50% as compared to peak times in order to maintain safety and reduce flight delays.

Two issues were at the root of the problem:

- A telecommunication line that transmits data and audio to air traffic controllers failed, and
- A radar feed that transmits data from an FAA facility to the Philadelphia TRACON and finally to the Newark airport also failed.

These issues were very likely tied to a contentious move of Newark air traffic controllers from Westbury, N.Y. to Philadelphia in 2024. Many controllers quit rather than relocate. But with the move there were no redundant data connections built between the radar processing center in New York and the Philadelphia facility where the controllers were moved. Nor was a backup system put in place (Sherman, 2025).

On top of all that, the Newark air traffic controllers faced long-standing, FAA-wide challenges:

- Outdated technology:
    - Telecommunication links built upon older copper wire.
    - Air traffic controllers using paper strips to track flights, manage traffic flow, and communicate between themselves.
    - Floppy disks used to update software and transfer data.
    - Computers running Windows 95.
- Understaffing
    - The FAA says that nationwide they are around 3,500 air traffic controllers short of staffing targets (Federal Aviation Administration, 2025).
    - The demanding nature of the air traffic control profession leads to early retirements and additional turnover. Filling those empty seats requires candidates with unique skills. According to the air traffic controllers' union, only 50% of trainees complete the training and achieve full certification.

In response to the outages, the FAA quickly outlined several measures to strengthen air traffic control at Newark Liberty International Airport. These included:

- Installing three additional high-capacity telecommunications links to the Philadelphia TRACON to increase reliability, speed, and system redundancy.
- Upgrading existing copper connections to fiber-optic lines, which support faster data transfer, greater bandwidth, longer transmission distances, and improved security.
- Introducing a temporary backup system at the Philadelphia TRACON to ensure operational continuity (Sherman, 2025).

Furthermore, the FAA continues their decades-long effort to implement solutions for these problems across the U.S., but it remains slow-going. For instance, technology upgrades are challenging in circumstances where existing information systems must be up and running 24/7, which is paramount given the importance of the air traffic control system.

Questions:

- What does this scenario reveal about the resilience (or lack thereof) of the U.S. air traffic control system?
- What are the risks of relying on outdated technology in critical infrastructure like air traffic control?
- Why do you think only 50% of trainees successfully complete air traffic control certification? What might be done to improve that rate?
- If you were tasked with reimagining the air traffic control system for the next 5 years, what would be your top three priorities—and why?
- What other questions or observations can you make about this scenario as it relates to resilience, risk management, information, and information systems?

3.3 Ransomware attack on Baltimore

One Tuesday morning in May 2019, Maria Santos arrived at her real estate office in Baltimore with excitement bubbling inside her. After months of searching, her clients, the Johnsons, had finally found their dream home — a charming rowhouse in Federal Hill. The closing was scheduled for that afternoon, and all that remained was the routine verification of city liens and water bills. But when Maria tried to access the city's online system to pull the necessary documents, she was met with an ominous message. The city's servers were down. What she did not know yet was that Baltimore had become the latest victim of a devastating ransomware attack—one where hackers had demanded $76,000 in ransom and would hold the city's computer systems hostage for over five weeks.

The attack had brought the real estate market to a grinding halt as property transfers could not be completed digitally, and the city was unable to issue lien certificates or generate water bills. Maria spent hours on the phone with increasingly frustrated city employees who could only tell her that their systems were compromised, and that they had no timeline for recovery. The Johnsons' closing was postponed indefinitely, along with hundreds of other real estate transactions across the city. For weeks, Baltimore's housing market essentially froze as buyers, sellers, and real estate professionals found themselves caught in digital limbo.

Meanwhile, across town, David Kim was dealing with his own crisis. As a city employee working in the health department, he arrived at work to find that not only could he not access his email, but the entire network was locked down. By day 36 of the attack, only 70 percent of city employees had regained access to their email accounts, with recovery efforts aiming to restore 95 percent of employee access by week's end. David's team was responsible for sending out health alerts and coordinating with local hospitals, but with their systems compromised, they were forced to resort to personal phones and handwritten notes. Critical public health communications that should have taken minutes to distribute now took hours, creating potential risks for the city's most vulnerable residents (Mathews, 2025).

The attack's effects rippled through Baltimore's daily life for months. As Deputy Chief of Staff Sheryl Goldstein warned residents, "I do not expect June bills to go out," water bills for June were delayed indefinitely, creating a backlog of charges that would eventually hit residents all at once. Families who budgeted for monthly water payments suddenly faced the uncertainty of not knowing their usage or owing amounts. The city's inability to process online payments meant residents had to find alternative ways to pay fines and taxes, though parking tickets could eventually be looked up in person. The city ultimately invested more than $18 million in recovery efforts — far exceeding the hackers' original $76,000 ransom demand that Mayor Young had refused to pay. But the true cost — measured in disrupted lives, delayed dreams, and shaken confidence in digital infrastructure — was immeasurable. The attack served as a stark reminder of how deeply intertwined modern urban life had become with digital systems that could vanish in an instant (Gallagher, 2019).

Questions:

- What are the pros and cons of paying a ransom to a ransomware attacker?
- How can systems be protected from ransomware attacks?
- Baltimore operated most of the systems in this scenario in its own data centers. Could they have reduced their risk by running their systems in the cloud instead? Does moving to the cloud have risk trade-offs?
- What other questions or observations can you make about this scenario as it relates to resilience, risk management, information, and information systems?

3.4 Change Healthcare Cyber Attack

A series of incidents related to patients, providers, and hospitals impacted stakeholders in Change Healthcare, a health insurance company based in Nashville, TN.

Alan regularly went to his Naperville, Illinois, pharmacy to pick up medications for his congestive heart failure and diabetes, which are covered by Medicaid. At his most recent visit, his pharmacist told him that he would have to pay out of pocket, that Medicaid would not cover the cost. Unable to do so, Alan ended up in a hospital (De Mar, 2024).

For a month now, Margaret, a dermatologist in California, has been unable to submit insurance claims to get paid for the services her private practice has provided. She’s been considering borrowing money to pay rent and her staff (Johnson & Ibarra, 2024).

A cancer clinic in Oregon has been unable to bill for its $500,000 to $1 million daily chemotherapy costs. Mel, the company’s Chief Financial Officer, has managed so far out of cash flow, but with reserves running low, she worries they may not have the money to pay for labor and keep their doors open (Steenhuysen, 2024).

These three unfortunate circumstances, along with countless others, were due to the largest healthcare data breach in U.S. history.

Change Healthcare, a subsidiary of UnitedHealth Group (UHG), based in Nashville, Tennessee, provides billing and data services for the healthcare industry, processing 15 billion transactions a year. On February 21st, 2024, the company discovered that their information systems were under cyber attack. In order to protect their partners and patients, they quickly disconnected their systems which were relied upon by hospitals, pharmacies, and doctors’ offices.

The cyber attack actually began on February 12th, when a hacker gained access to Change Healthcare’s information system network by logging into a remote access service that lacked multi factor authentication using compromised credentials (see vulnerabilities and their remediations in Table 1, below). Undetected over a ten-day period, the hacker was able to create privileged administrative accounts, remove vast amounts of sensitive data, and encrypt files on Change Healthcare’s systems via ransomware - malicious software designed to block access to a computer system until a ransom payment is made. The attack was only detected when system files were encrypted, preventing access.

Five days after Change Healthcare discovered the attack, a ransomware group named ALPHV/Blackcat claimed responsibility for the attack, said that 6TB of data was stolen, and demanded a $22 million ransom to prevent the publication of the stolen data. Change Healthcare ultimately paid the ransom, but their payment did not secure the stolen data due to a scam by the ALPHV/Blackcat group. As of June 2025, there have been partial leaks of the stolen data - screenshots and documents - but there have been no indications of complete data sets being disclosed.

Change Healthcare restored operations of their electronic payments platform and reinstated 99% of its pharmacy network services in mid-March (Hyperproof Team, 2024). And while many of their services recovered within a few months, Change Healthcare didn’t declare that their clearinghouse services had been fully restored until November, nine months after the attack. Furthermore, in October, Change Healthcare estimated that the breach involved the data of approximately 100 million people, and amid ongoing analysis revised that number to 190 million in January of 2025 - eleven months after the attack began. The stolen data included medical records, insurance records, dental records, payments/claims information, and patients’ Personally Identifiable Information (PII) - phone numbers, addresses, social security numbers, driver’s license numbers, and email addresses (Gatlan, 2024).

The financial impacts of the attack were staggering. UnitedHealth Group estimated the overall cost of the Change Healthcare ransomware attack at $2.87 billion in 2024. According to the American Medical Association, the attack resulted in a $100 million impact to the healthcare industry per day, as claims were delayed and physicians were unable to get paid for their services (Hatton, 2024). UnitedHealth Group made loans totaling close to $9 billion to healthcare providers to help ease the financial strain caused by the extended outage.

Vulnerability	**Remediation**
Remote access service that lacked multi factor authentication.	All external facing systems have been or will be enabled with multi-factor authentication.
Hacker was able to disable both Change Healthcare's primary information systems as well as their backup because the two weren’t isolated.	Company is rebuilding legacy systems, shifting on-premise systems to cloud-based systems designed with built-in security controls, including segmentation, isolation and enhanced backup strategies (Jones, 2024).

Table 1: Vulnerabilities and Remediations

UnitedHealth Group and Change Healthcare learned some hard lessons from the attack, providing opportunities to improve their cybersecurity resilience. The two primary vulnerabilities and their remediations are captured in Table 1.

<u>Questions:</u>

- How does the use of cloud architecture enhance organizational security and resilience?
- How can organizations discover and evaluate third-party risks in their information systems strategy?
- How does this incident raise the importance of cybersecurity to a board-of-directors-level concern?
- What Key Performance Indicators (for instance the number of security incidents detected over a period of time) would you track to assess the effectiveness of security in a healthcare IT system?
- What other questions or observations can you make about this scenario as it relates to resilience, risk management, information, and information systems?

4. ADDITIONAL RESEARCH OPPORTUNITIES

If one pays attention, it is possible to see new examples of both brittle and resilient systems in many settings and situations. Please feel free to use these prompts as you and your instructor see fit, to conduct additional research and analysis of this broadly visible topic.

Resilience of emerging technologies

Emerging technologies significantly impact system resilience, offering both opportunities and challenges. While they can enhance a system's ability to withstand and recover from disruptions, they also introduce new vulnerabilities and complexities that require careful management. Key areas of impact include cybersecurity, infrastructure resilience, and the ability to adapt to evolving threats. Examples might include autonomous vehicles, drone deliveries, augmented and virtual realities, artificial intelligence in all its forms, and many more.

In consultation with your instructor, pick one or more related emerging technologies, and perhaps a particular industry to apply them to. Research the resilience of that technology and then analyze the potential resilience impacts of that technology on your particular industry choice.

Heathrow airport power outage

On March 21, 2025, London Heathrow Airport was forced to shut down for nearly a full day due to a fire at the North Hyde electrical substation in

Hayes, west London (Mathews, 2025). The fire began at 23:23 on March 20, 2025, and the airport remained closed until 23:59 on March 21.

Marhea, a 74-year-old passenger booked on a Brussels Airlines flight to Liberia, described arriving to "darkness and confusion" with no staff available to explain what had happened. "They didn't let us into terminal five. People were standing at the door everywhere," she said. Her airline eventually rebooked her on a Saturday flight with a different carrier. Another passenger, Ellen, had her surprise 30th birthday trip to Venice cancelled, telling Al Jazeera: "We were supposed to fly to Venice this morning from Heathrow for a day trip for my 30th birthday present, it was a surprise booked by my cousin for the two of us". Perhaps even worse, "transit" passengers who were passing through London en route to another country, often did not have a valid UK entry visa, and thus were restricted to a small part of the airport - without power, heating, cooling, or food.

The subsequent Kelly Review found that while Heathrow made appropriate decisions during the crisis, it highlighted the need for better systemic resilience planning and recommended prioritizing investment in backup power systems for critical operations The Kelly Review: Lessons from Heathrow's power outage. The review also noted communication gaps between technical teams and other airport staff regarding potential vulnerabilities (Mathews, 2025).

In consultation with your instructor, identify some aspects of the Heathrow power outage that you can analyze to better understand what went wrong, why such a heavily redundant system can fail, and what were the root causes of the failure.

5. CONCLUSIONS

In this modern world, where so much of daily life is dependent on systems of one form or another, it is important to have the systems designed, built, and tested for resilience. Doing so requires all stakeholders, including business and technology teams, to understand the concepts of system resilience, as well as the factors that influence it. We believe that these scenarios will illustrate a significant range of resilience failure modes, and substantial fodder for discussion of system resilience issues in the classroom, as well as writing and research assignments outside the classroom.

6. ACKNOWLEDGEMENTS

The authors appreciate the support and insights of various industry experts who shared their time and expertise to discuss this case. The authors also appreciate the productive feedback provided by their students, their peers, and by the reviewers and conference chairs.

7. REFERENCES

Ahmed, R. (2017). Risk Mitigation Strategies in Innovative Projects. In *Key Issues for Management of Innovative Projects*. IntechOpen. https://doi.org/10.5772/intechopen.69004

Alozie, C. E., Akerele, J. I., Kamau, E., & Myllynen, T. (2024). Disaster recovery in cloud computing: Site reliability engineering strategies for resilience and business continuity. *International Journal of Management and Organizational Research*, *3*(1), 36-48.

Cerullo, M. (2024). Delta cancels hundreds more flights as fallout from CrowdStrike outage persists. *CBS News*. Retrieved 7/1/2025, from https://www.cbsnews.com/news/delta-crowdstrike-outage-flight-status/

De Mar, C. (2024). Hacking of health care company leaves Chicago area man stuck in hospital. https://www.cbsnews.com/chicago/news/hacking-health-care-company-chicago-area-man-stuck-in-hospital/

Federal Aviation Administration. (2025). *Update: Newark Liberty International Airport* (FAA General Statements, Issue. https://www.faa.gov/newsroom/statements/general-statements

Gallagher, S. (2019). Baltimore ransomware nightmare could last weeks more, with big consequences. *Harm City*. https://arstechnica.com/information-technology/2019/05/baltimore-ransomware-nightmare-could-last-weeks-more-with-big-consequences

Gatlan, S. (2024). Ransomware gang claims they stole 6TB of Change Healthcare data. *Bleeping Computer*. https://www.bleepingcomputer.com/news/security/ransomware-gang-claims-they-stole-6tb-of-change-healthcare-data/

Hatton, R. (2024). How has the Change Healthcare cyberattack affected physicians?

Becker's ASC Review. Retrieved June 28, 2025, from https://www.beckersasc.com/asc-news/how-has-the-change-healthcare-cyberattack-affected-physicians/

Hyperproof Team. (2024). Understanding the Change Healthcare Breach and Its Impact on Security Compliance. https://hyperproof.io/resource/understanding-the-change-healthcare-breach

Johnson, K., & Ibarra, A. B. (2024). California doctors struggle to make payroll one month after ransomware attack. *News from the States*. Retrieved July 5, 2025, from https://www.newsfromthestates.com/article/california-doctors-struggle-make-payroll-one-month-after-ransomware-attack

Jones, S. (2024, July 3). What the Change Healthcare Cyber Attack Means for the US Healthcare Industry. *Hornet Security*. https://www.hornetsecurity.com/en/blog/change-healthcare-cyber-attack/

Kerner, S. M. (2024). CrowdStrike outage explained: What caused it and what's next. *Tech Accelerator*. https://www.techtarget.com/whatis/feature/Explaining-the-largest-IT-outage-in-history-and-whats-next

Martínez, A., & Dumas, N. (2025). How archaic tech, staff shortages and construction made a meltdown at Newark airport. *NPR National*. Retrieved 06/23/2025, from https://www.npr.org/2025/05/07/nx-s1-5388438/airlines-trade-group-vp-discusses-newark-airport-delays

Mathews, R. (2025). *The Kelly Review: Lessons from Heathrow's power outage*. https://www.thebci.org/news/the-kelly-review-lessons-from-heathrow-s-power-outage.html

Otkhozoria, N., Petriashvili, L., Zhvania, T., & Imerlishvili, A. (2025). Advancing information system testing: challenges, methods, and practical recommendations. *International Science Journal of Engineering & Agriculture*, *4*(2), 203-214.

Rooney, J. J., & Heuvel, L. N. V. (2004). Root Cause Analysis For Beginners. *Quality Progress*. https://servicelink.pinnacol.com/pinnacol_docs/lp/cdrom_web/safety/management/accident_investigation/Root_Cause.pdf

Rose, J. (2025). Disbelief, then fury: A Newark air traffic controller says they saw a crisis coming. *NPR*. Retrieved 05/29/2025, from https://www.npr.org/2025/05/22/g-s1-68333/newark-air-traffic-controller-atc

Sherman, T. (2025). Newark's air traffic nightmare continues as controllers lose contact with planes a 4th time. *NJ News*. Retrieved May 20, from https://www.nj.com/news/2025/05/newarks-air-traffic-nightmare-continues-as-controllers-lose-contact-with-planes-a-4th-time.html

Shi, X., Liu, W., & Lim, M. K. (2024). Supply chain resilience: new challenges and opportunities. *International Journal of Logistics Research and Applications*, *27*(12), 2485-2512.

Steenhuysen, J. (2024, Mar 7). Patients or Payroll? US Healthcare Hack Creates Hard Choices. *Reuters*, https://www.reuters.com/world/us/patients-or-payroll-us-healthcare-hack-creates-hard-choices-2024-03-06/

Usman, A. B., & Asplund, M. (2025, May). Update at Your Own Risk: Analysis and Recommendations for Update-Related Vulnerabilities. In *IFIP International Conference on ICT Systems Security and Privacy Protection,* Maribor, Slovenia (pp. 97-110). Cham: Springer Nature Switzerland.

Witman, P. D., Prior, J., Nickl, T., & Mackelprang, S. (2024). The Southwest Airlines Winter Meltdown Case Studies on Risk, Technical Debt, Operations, Passengers, Regulators, Revenue, and Brand. *Information Systems Education Journal*, *22*(5), 59-71. https://doi.org/https://doi.org/10.62273/EFWA2093