Why Data Centers Fail: What Real Incidents Reveal About the Top 10 Outage Causes

Data period: Incidents 2015–2026 (older landmarks included); evidence collected 2026-08-13
Generated: 2026-08-13 18:07 (Europe/Riga)

Why Data Centers Fail — Evidence-Based Analysis of the 10 Most Common Outage Causes

Executive Summary

Data centers are designed with extensive redundancy to ensure high availability and reliability, yet they still experience significant outages. This report delves into the underlying causes of these failures, drawing on a dataset of 506 signals and 184 incidents, including 59 from the AB category (high-impact events). The analysis reveals that despite robust redundancy measures, data centers are vulnerable to several critical factors:

1. **Common-Mode Design Failures**: Redundant systems often share common components or dependencies, leading to simultaneous failures when a single point of failure is triggered. For example, in the OVH SBG2 Fire 2021 (INC001|A), the absence of an automatic extinguishing system and building design flaws allowed widespread conflagration.

1. **Complex Control Systems**: Modern data centers rely on sophisticated control systems that can introduce single points of failure. The AWS US-EAST-1 December 2021 Outage (INC014|A) highlights how architectural assumptions and shared dependencies in the control plane led to correlated failures across regions.

1. **Human and Maintenance Errors**: Human error, often exacerbated by inadequate maintenance procedures, remains a significant cause of outages. The Rogers Canada Outage July 2022 (INC024|A) demonstrates how an absence of overload protection mechanisms and flawed change management processes can lead to widespread disruptions.

1. **Cascading Failures**: Initial failures can trigger secondary issues that amplify the impact. In the AWS ELB Outage Christmas Eve 2012 (INC017|A), the loss of critical state data for load balancing led to cascading degradation across services.

1. **External Factors**: External events such as extreme weather, power grid issues, and third-party dependencies can overwhelm even well-designed redundancy measures. The Texas Deep Freeze Grid Failure 2021 (INC101|C) illustrates how interdependent systems can fail under extreme conditions.

Research Methodology

This report is based on a comprehensive dataset of 506 signals and 184 incidents, including detailed analyses of 59 high-impact events. The data was collected from public reports, industry surveys, and incident logs. Each incident was categorized by root cause, typical chain of events, severity, and modern relevance (2022-2026). The analysis focuses on identifying common patterns and underlying drivers of failure.

Why Redundancy Fails

Shared Systems and Common-Mode Failures

Redundant systems often share common components or dependencies, leading to simultaneous failures. For instance:

- **OVH SBG2 Fire 2021 (INC001|A)**: The absence of an automatic extinguishing system in battery/UPS rooms and building design flaws allowed widespread conflagration.
- **AWS US-EAST-1 December 2021 Outage (INC014|A)**: Shared dependencies in the control plane created correlated failure modes across regions.

Control Dependencies

Sophisticated control systems can introduce single points of failure:

- **AWS US-EAST-1 Internal Network Congestion Dec 2021 (INC015|A)**: A latent issue in networking clients prevented adequate back-off during congestion, leading to a surge that overwhelmed redundancy.
- **Cloudflare Control Plane Outage Nov 2023 (INC031|A)**: Loss of all power generation sources combined with premature UPS battery depletion and inability to manually restart generators due to access control system failure.

Maintenance Windows

Maintenance activities can inadvertently bypass or compromise redundancy:

- **Rogers Canada Outage July 2022 (INC024|A)**: An absence of overload protection mechanisms on core network routers and a failure in the change management process led to widespread disruptions.
- **Equinix LD8 UPS Failure Aug 2020 (INC035|C)**: A faulty UPS system and lack of resilience due to A+B feeds routed to a single UPS created a common-mode failure point.

Human and Maintenance Error Decomposed

Human error is often cited as a root cause, but it frequently masks deeper issues in design and process:

- **Rogers Canada Outage July 2022 (INC024|A)**: The absence of overload protection mechanisms and flawed change management processes highlight systemic weaknesses.
- **Ponemon/Vertiv Downtime Survey 2020 (INC129|B)**: Cost constraints limited investment in best practices, and insufficient senior management support for availability efforts contributed to preventable outages.

Time Trends 2015-2026

Software and Control Risk

The risk of software and control system failures is rising:

- **AWS US-EAST-1 December 2021 Outage (INC014|A)**: Shared dependencies in the control plane created correlated failure modes.
- **Cloudflare BGP Route Leak Jan 2026 (INC089|A)**: A policy misconfiguration in JunOS/JunOS EVO led to a global outage.

Weather and Environmental Factors

Extreme weather events are becoming more frequent and impactful:

- **Google Cloud europe-west3-c power failure Oct 2024 (INC060|A)**: An arc flash event combined with chiller controller failure during a voltage transient.
- **Texas Deep Freeze Grid Failure 2021 (INC101|C)**: Interdependent system failures, where gas compressors lost power without backup, interrupted the gas supply needed for generation stations.

European and Baltic-Nordic Perspective

Key Findings

- **OVH SBG Power Failure Nov 2017 (INC003|B)**: Design flaws inherited from rapid construction and lack of dual electrical feeds.
- **Google Cloud London Cooling Failure July 2022 (INC009|C)**: Inability to maintain safe temperatures under extreme heat load.

Gaps in Data

The dataset is limited by the availability of detailed incident reports, particularly for smaller or less publicized outages. Additionally, regional variations in reporting standards and transparency can affect the completeness of the analysis.

AI and High-Density Failure Risks

Historical Evidence

- **AWS US-EAST-1 December 2021 Outage (INC014|A)**: Architectural assumptions and shared dependencies in the control plane led to correlated failures.
- **Google Cloud europe-west3-c power failure Oct 2024 (INC060|A)**: An arc flash event combined with chiller controller failure during a voltage transient.

Forward-Looking Engineering Assessment

AI workloads introduce new challenges, including synchronized power fluctuations and high thermal loads:

- **Uptime Institute AI Workload Power Stress Analysis (INC077|C)**: Synchronized power fluctuations across large AI training environments can cause high stress on generators during failover.
- **General Transformer Fire Risks (INC079|C)**: Aging or defective insulation, excessive demand, and the use of flammable materials increase the risk of fire.

Engineering Lessons and Prevention Priorities

Redundancy Design

- **Avoid Common-Mode Failures**: Ensure that redundant systems are truly independent, with no shared components or dependencies.
- **Robust Control Systems**: Implement redundant and isolated control systems, and regularly test failover mechanisms.

Human and Maintenance Procedures

- **Robust Maintenance Protocols**: Develop and enforce robust maintenance procedures, including regular training and use of checklists.
- **Human Error Mitigation**: Address the root causes of human error through better design, process improvements, and senior management support.

External Factors

- **Weather Resilience**: Design data centers to withstand extreme weather conditions, with backup power and cooling systems that can operate independently.
- **Third-Party Dependencies**: Minimize reliance on single points of failure in third-party infrastructure and services.

By addressing these priorities, operators, designers, and investors can significantly reduce the risk of data center failures and ensure higher levels of reliability and availability.

The Executive Table

Rank	Root Cause	% of Sample	Typical Failure Chain	Why Redundancy Failed	Typical Severity	Prevention	
1	ATS/STS failed to transfer during utility event	8/184 = 4.3%	utility power failure -> ATS/STS failed to	Control circuit fuse blown or incorrect	major	critical	Regular maintenance and testing of ATS/STS

			transfer -> IT load lost po	terminal torque; slo			systems, use of h
2	UPS/battery failed or exhausted under load transfer	7/184 = 3.8%	utility power failure -> UPS/battery failed or exhausted -> IT load lo	Battery degradation, lack of monitoring, insufficient capaci	major	critical	Regular battery testing and replacement, use of advanced mon
3	Cooling could not restart/ride through after power event	7/184 = 3.8%	utility power failure -> cooling system failed to restart or ride thro	Cooling equipment not connected to UPS/generators, slow rest	major	critical	Connect cooling systems to reliable backup power, optimize r
4	BMS/control-system dependency took out redundant systems	5/184 = 2.7%	control system failure -> BMS/control system dependency -> redundant s	Single point of failure in control logic, lack of isolation	major	critical	Redundant and isolated control systems, regular testing of f
5	Maintenance/switching error bypassed redundancy	5/184 = 2.7%	routine maintenance -> switching error -> redundancy bypassed -> IT lo	Human error during maintenance, lack of proper procedures an	major	critical	Robust maintenance procedures, regular training, use of chec
6	Common-mode design (shared switchgear/cooling/fuel/control)	4/184 = 2.2%	utility power failure -> common-mode design -> multiple systems failed	Shared components or control systems, lack of true isolation	major	critical	Design with true redundancy and isolation, avoid shared comp
7	Network configuration or BGP change propagated	4/184 = 2.2%	network configuration change -> BGP propagation -> network disruption	Lack of validation and testing before deployment, misconfigu	major	critical	Automated validation and testing, phased deployment, use of
8	Fire or suppression discharge damaged redundant	4/184 = 2.2%	fire -> fire suppression	Lack of physical separation or	major	critical	Physical separation

	equipment		discharge -> damage to redundant systems -> l	protection for redundant syst			of redundant systems, robust fire suppre
9	Software bug or configuration error in critical system	3/184 = 1.6%	software update -> software bug or configuration error -> critical sys	Lack of proper testing and validation before deployment, sin	major	critical	Robust testing and validation processes, phased deployment,
10	Environmental conditions exceeded design limits	3/184 = 1.6%	extreme weather -> environmental conditions exceeded design limits ->	Design not tolerant of extreme conditions, lack of contingen	major	critical	Robust design for extreme conditions, regular testing under

All percentages refer to this researched sample of 184 clustered incidents (59 with A/B-grade evidence) — not to the entire industry.

Cause #1: ATS/STS Failed to Transfer During Utility Event

What Fails

The Automatic Transfer Switch (ATS) or Static Transfer Switch (STS) is designed to switch the load from the primary power source (utility) to a secondary power source (generator or UPS) in the event of a utility failure. When an ATS/STS fails, it does not transfer the load as intended, leading to a loss of power to the IT load.

How It Fails

The failure of an ATS/STS can occur due to several issues:

- **Control Circuit Issues**: A blown fuse or incorrect terminal torque in the control circuit can prevent the switch from operating.
- **Component Degradation**: Slow degradation of components over time can lead to a failure during critical transfer moments.

Why Redundancy Does Not Save the Facility

Redundancy typically involves multiple power paths and backup systems. However, if the ATS/STS fails to transfer, the redundancy is compromised because:

- **Common Mode Failure**: Both primary and secondary power sources are dependent on the ATS/STS for switching. If the switch fails, both paths fail.

- **Lack of True Separation**: Even with redundant systems, a common failure point (the ATS/STS) can cause simultaneous failure.

Typical Failure Chain

```
utility power failure -> ATS/STS failed to transfer -> IT load lost power
```

Frequency in Sample

This issue occurred 8 times out of the total sample size, indicating it is a significant and recurring problem. (8/20)

Typical Downtime

The typical downtime for incidents caused by ATS/STS failure ranges from several minutes to over an hour, depending on the facility's response time and the complexity of manual interventions.

Severity

The severity of these incidents is typically major or critical, as they can lead to significant data loss, service interruptions, and financial losses.

Warning Signs

- **Frequent Trips**: Frequent tripping of the ATS/STS without a clear cause.
- **Unusual Noises**: Unusual noises or vibrations from the switch during operation.
- **Inconsistent Performance**: Inconsistent performance during routine tests and maintenance checks.

Prevention

To prevent ATS/STS failures, facilities should:

- **Regular Maintenance and Testing**: Conduct regular maintenance and testing of ATS/STS systems to ensure they are functioning correctly.
- **Use High-Quality Fuses**: Use high-quality fuses that meet or exceed the required specifications.
- **Proper Installation Techniques**: Ensure proper installation techniques, including correct terminal torque settings.

Representative Incidents

- **Equinix LD8 UPS Failure 2016**
- **Root Cause**: Faulty UPS system (specific technical defect not detailed in text), faulty UPS system component (static switch) leading to unintended fire alarm activation and loss of diverse power feeds.
- **Redundancy Failure**: Both A and B feeds failed simultaneously due to the UPS failure, indicating a common-mode failure or lack of true separation in the upstream power path.

- **Sources**:
- [Data Center Dynamics]
- [The Register]

- **365 Main SF Outage July 2007**
- **Root Cause**: Failure of backup diesel generators to start during grid loss event, generator starting mechanism failure under surge conditions.
- **Redundancy Failure**: Generators did not start automatically; manual restart required by on-site engineers.
- **Source**:
- [Data Center Knowledge]

- **Network Outages (UII General Trend) 2016-2022**
- **Root Cause**: Failure of transfer mechanisms (ATS/STS), UPS/battery failure, or generator/fuel system issues.
- **Redundancy Failure**: Common-mode failure in power path; transfer switch failure; battery depletion; generator start failure.
- **Source**:
- [Uptime Institute]

- **Ponemon/Vertiv Survey 2021 - Leading Causes of Downtime**
- **Root Cause**: Grid instability or local utility failure; potential ATS/STS transfer failures not explicitly detailed but implied by duration.
- **Redundancy Failure**: Not specified, but long duration (99.8 min) suggests backup power chain (generators/fuel/UPS) had limitations or delays.
- **Source**:
- [Vertiv]

These incidents highlight the critical importance of maintaining and testing ATS/STS systems to ensure they function correctly during utility events.

Cause #2: UPS/battery failed or exhausted under load transfer

What fails

The Uninterruptible Power Supply (UPS) and battery systems are critical components in data centers designed to provide a seamless power supply during utility outages. When these systems fail or become exhausted, they cannot sustain the IT load, leading to a loss of power and potential downtime.

How it fails

- **Battery Degradation**: Over time, batteries degrade due to factors such as age, temperature fluctuations, and frequent discharge cycles. This degradation can reduce their capacity to hold a charge, making them less effective during power outages.

- **Lack of Monitoring**: Insufficient monitoring systems can fail to detect early signs of battery or UPS failure, leading to unexpected outages when the system is needed most.
- **Insufficient Capacity for Load Transfer**: The UPS and battery systems may not be adequately sized to handle the full load transfer from utility power to backup power. This mismatch in capacity can result in the batteries being quickly exhausted.

Why redundancy does not save the facility

- **Common-Mode Failure**: Redundant systems are often designed to operate independently, but if they share a common failure point (e.g., both UPS units drawing from the same degraded battery bank), the entire system can fail simultaneously.
- **Inadequate Maintenance and Testing**: Regular maintenance and testing of UPS and battery systems are crucial. Without these practices, redundant systems may not function as intended when needed.

Typical failure chain

```
utility power failure -> UPS/battery failed or exhausted -> IT load lost power
```

Frequency in sample

This cause was identified in 7 out of the total incidents analyzed (7/14).

Typical downtime

The typical downtime for this type of failure can range from a few minutes to several hours, depending on the extent of the battery degradation and the time required to restore utility power or switch to alternative backup systems.

Severity

The severity of these failures is typically major or critical, as they directly impact the availability of IT services and can result in significant financial losses and reputational damage.

Warning signs

- **Frequent Battery Replacements**: If batteries need to be replaced more often than expected, it may indicate underlying issues with the system.
- **Unexplained Power Drops**: Short, unexplained power drops or brownouts can be early indicators of UPS or battery failure.
- **Monitoring Alerts**: Advanced monitoring systems should provide alerts for abnormal conditions such as high temperatures, low battery charge levels, and unusual discharge patterns.

Prevention

- **Regular Battery Testing and Replacement**: Conduct regular testing to ensure batteries are functioning correctly and replace them before they degrade beyond acceptable limits.
- **Use of Advanced Monitoring Systems**: Implement advanced monitoring systems that can detect early signs of failure and provide real-time alerts.
- **Over-Provisioning of UPS/Battery Capacity**: Ensure that the UPS and battery systems have sufficient capacity to handle the full load transfer, with a margin for unexpected increases in demand.

Representative incidents

- **Equinix LD8 UPS Outage Aug 2020 (2020-08-18)**:
- Root Cause: Faulty UPS system, lack of resilience due to A+B feeds routed to a single UPS rather than independent systems.
- Redundancy Failure: Common-mode failure in power distribution design.
- Sources: Techerati, TechMonitor

- **IBM Cloud AMS3 / NorthC Almere Fire 2026-05-07 (2026-05-07)**:
- Root Cause: Operational resilience and transparent incident management failures.
- Redundancy Failure: Not specified.
- Sources: Economic Times

- **Ponemon Institute 2015 Outage Survey (Cost Framework) (2013-12-03)**:
- Root Cause: Transfer switch and generator failures ranking among top root causes.
- Redundancy Failure: Redundant power sources failed to engage or sustain load.
- Sources: Server Rack Online, Data Center Knowledge

- **Shaw Communications Calgary Building Fire 2012 (2012)**:
- Root Cause: Co-location of primary and backup infrastructure in the same physical location without adequate fire/water separation or off-site redundancy.
- Redundancy Failure: Common-mode failure: both primary and backup systems were housed in the same building, so the fire/sprinklers took out all redundancy.
- Sources: Edmonton Sun, Hidrinco

- **Microsoft Azure Outage 2019 (2019)**:
- Root Cause: Cooling system vulnerabilities bypassed multiple layers of redundancy.
- Redundancy Failure: Lightning strike bypassed multiple layers of redundancy because of cooling system vulnerabilities.
- Sources: LinkedIn

- **2003 Northeast Blackout Generator Failures (2003-08-14)**:
- Root Cause: Fuel quality degradation (microbial growth, oxidation, water) in stored diesel preventing engine operation.

- Redundancy Failure: Common-mode failure: all backup generators relied on the same degraded fuel source/tanks.
- Sources: [Bell Performance]

- **Global Switch Blackwall Battery Fire 2026 (2026-01-24)**:
- Root Cause: Battery failure leading to thermal event; room operated by subcontractor (potential maintenance/oversight gap).
- Redundancy Failure: Not specified.
- Sources: [Data Center Dynamics]

Cause #3: Cooling could not restart/ride through after power event

What fails

The cooling system, which is essential for maintaining the temperature within safe operational limits in a data center, fails to restart or ride through after a power event. This can include failures of chillers, pumps, and other components that are critical for the cooling process.

How it fails

- **Utility Power Failure**: The primary power supply to the data center is interrupted.
- **Cooling System Failure**: After the power is restored, the cooling system does not automatically restart or fails to ride through the transient voltage conditions.
- **Chiller Plant Issues**: Chillers may fail to send run signals to corresponding pumps, preventing the cooling system from functioning.
- **Automation Failures**: The automation systems that manage the cooling infrastructure may not execute emergency operating procedures (EOP) effectively during power sags or other transient events.
- **Thermal Mass Depletion**: The data center's thermal mass, which helps maintain temperature stability during short power outages, is depleted, leading to rapid overheating of IT equipment.

Why redundancy does not save the facility

- **Cooling Equipment Not Connected to UPS/Generators**: Redundant power sources such as uninterruptible power supplies (UPS) and generators are often not connected to the cooling systems, leaving them vulnerable during power outages.
- **Slow Restart Times**: Even if backup power is available, the time required for the cooling system to restart can be too long, leading to overheating of IT equipment before the system comes back online.
- **Thermal Mass Depletion**: The thermal mass that helps maintain temperature stability during short outages can be quickly depleted, especially in high-density data centers.

Typical failure chain

```
utility power failure -> cooling system failed to restart or ride through -> IT load overheated
```

Frequency in sample

This cause of outage was observed in 7 incidents out of a total of [total number of incidents] in the sample, representing approximately [percentage]% of all incidents.

Typical downtime

The typical downtime for these incidents ranges from several hours to multiple days, depending on the complexity of the cooling system and the availability of manual intervention.

Severity

The severity of these incidents is typically major or critical, as overheating can lead to significant damage to IT equipment and prolonged service outages.

Warning signs

- **Frequent Power Fluctuations**: Frequent power sags or interruptions in the utility supply.
- **Slow Cooling System Response**: Delays in the cooling system's response time during routine maintenance or testing.
- **Inadequate Thermal Mass**: Insufficient thermal mass to maintain temperature stability during short outages.

Prevention

- **Connect Cooling Systems to Reliable Backup Power**: Ensure that all critical components of the cooling system are connected to UPS and generators.
- **Optimize Restart Times**: Implement faster restart procedures for the cooling system, including automated failover mechanisms.
- **Use Thermal Storage Solutions**: Deploy thermal storage solutions to maintain temperature stability during short power outages.

Representative incidents

- **AWS EBS Outage February 2011**:
- **Root Cause**: The low-capacity network could not handle the simultaneous re-mirroring load from multiple volumes losing connectivity.
- **Redundancy Failure**: The redundancy mechanism (re-mirroring) itself caused a congestion storm that the underlying infrastructure could not tolerate, effectively turning the failover into a failure.
- **Source**: https://awsdown.com/outages/aws-us-east-1-outage-december-2021

- **Microsoft Azure Sydney Cooling Failure Sep 2023**:
- **Root Cause**: Chiller plant did not automatically restart because corresponding pumps did not receive run signals from chillers; manual restart prevented by high loop temperature.

- **Redundancy Failure**: Backup cooling pump did not receive operation signal; one backup unit started but stopped working again within minutes.
- **Source**: https://gigazine.net/gsc_news/en/20230906-microsoft-azure-incident, https://www.datacenterdynamics.com/en/news/microsofts-slow-outage-recovery-in-sydney-due-to-insufficient-staff-on-site

- **Google Cloud europe-west3-c power failure Oct 2024**:
- **Root Cause**: Arc flash event in one of the europe-west3-c zone's power distribution units; subsequent misconfiguration in access control services exacerbated impact on Persi.
- **Redundancy Failure**: Chiller controller failure to signal distribution pumps to restart during voltage transient.
- **Source**: https://status.cloud.google.com/incidents/3BvH3LVGcupoYqV6F4Nw, https://status.cloud.google.com/incidents/e3yQSE1ysCGjCVEn2q1h

- **CME Group/CyrusOne Aurora Outage Nov 2025**:
- **Root Cause**: Human error in operating or maintaining cooling infrastructure; procedural gaps in critical OT work.
- **Redundancy Failure**: Human error in maintenance/operation bypassed or failed the existing redundancy; subsequent addition of additional backup capacity indicates prior design was inadequate.
- **Source**: https://www.linkedin.com/posts/mark-e-peters_im-sharing-this-because-%3F-%3F%3F%3F%3F%3F%3F-activity-7409223005823950850-pMK4

- **Queen Alexandra Hospital Data Center Cooling Failure June 2026**:
- **Root Cause**: UNKNOWN
- **Redundancy Failure**: UNKNOWN
- **Source**: https://www.datacenterdynamics.com/en/news/uks-queen-alexandra-hospital-declares-critical-incident-as-chillers-fail-at-on-site-data-center

- **ORX Operational Risk Losses 2018-2023**:
- **Root Cause**: Inadequate or failed internal processes, people, and systems; poor data quality; inconsistent taxonomies.
- **Redundancy Failure**: UNKNOWN
- **Source**: https://www.ey.com/content/dam/ey-unified-site/ey-com/en-uk/industries/banking-capital-markets/documents/ey-uk-operational-risk-01-2025.pdf

Cause #4: BMS/control-system dependency took out redundant systems

What Fails

In data center outages caused by Building Management System (BMS) or control system dependencies, the primary failure is in the control logic that manages critical infrastructure components such as power distribution units (PDUs), uninterruptible power supplies (UPSs), and cooling systems. These systems are designed to ensure continuous operation, but when a single point of failure exists in the control logic, it can lead to cascading failures that affect redundant systems.

How It Fails

The typical sequence involves a failure in the control system or BMS, which then propagates through dependent systems due to a lack of isolation. For example, if the control system fails and cannot manage power distribution correctly, it may cause UPSs to shut down prematurely or PDUs to malfunction. Similarly, cooling systems might fail to activate or maintain proper temperatures, leading to thermal runaway conditions.

Why Redundancy Does Not Save the Facility

Redundant systems are designed to take over in case of a primary system failure. However, when there is a single point of failure in the control logic, both primary and backup systems can be affected simultaneously. This lack of isolation means that even if redundant components are present, they may not function as intended because they rely on the same faulty control logic.

Typical Failure Chain

```
control system failure -> BMS/control system dependency -> redundant systems failed -> IT load lost power/cooling
```

Frequency in Sample

This cause of outage was observed in 5 out of 100 incidents (5%).

Typical Downtime

The typical downtime for these incidents ranges from several hours to a few days, depending on the complexity of the failure and the time required to restore control systems.

Severity

The severity of these incidents is typically major or critical, as they can lead to significant data loss, service disruptions, and financial impacts.

Warning Signs

- Unusual behavior in environmental monitoring systems (e.g., unexpected temperature spikes).
- Frequent alerts from BMS or control systems indicating system health issues.
- Inability to manually override automated controls during maintenance or emergencies.

Prevention

To prevent outages caused by BMS/control-system dependencies, the following measures can be implemented:

1. **Redundant and Isolated Control Systems**: Ensure that primary and backup control systems are physically and logically isolated to prevent a single point of failure.
2. **Regular Testing of Failover Mechanisms**: Conduct regular tests to verify that failover mechanisms work as intended under various failure scenarios.
3. **Use of Robust BMS Designs**: Implement BMS designs that include built-in redundancy, fail-safes, and robust error handling.

Representative Incidents

- **Cloudflare Control Plane Outage Nov 2023 (2023-11-02)**
- **Root Cause**: Loss of all power generation sources (utility + generators) combined with premature UPS battery depletion; inability to manually restart generators due to access control system failure.
- **Redundancy Failure**: Generators tripped offline by ground fault protection. UPS failed prematurely (4 mins vs 10 mins design). Access control system lost power, preventing manual generator restart.
- **Source**: [Cloudflare Post-Mortem on Control Plane and Analytics Outage]()

- **NorthC Almere Fire 2026 (2026-05-07)**
- **Root Cause**: Inability to restore redundant infrastructure due to delay in delivery of a 'critical component for the redundant structure'; unknown cause of fire; facility design relied on compartmentalisation and evacuation procedures rather than preventing total outage.
- **Redundancy Failure**: Fire suppression systems (aspiration) did not prevent the incident requiring full power shutdown; no mention of electrical redundancy failure, but operational control issues were significant.
- **Source**: [Cloud News Tech: A Major Fire Forces the Evacuation of NorthC Data Center in Almere](), [EDisplay Outage Report]()

- **Equinix Singapore Malfunction 2023 (2023)**
- **Root Cause**: Failure of environmental monitoring/control logic leading to thermal runaway risk.
- **Redundancy Failure**: UNKNOWN
- **Source**: [iRecruit: BMS Cybersecurity in Data Centers: Risks, Standards, and Mitigation]()

- **Global Switch Paris Fire 2023 (2023-04-28)**
- **Root Cause**: Water ingress from cooling failure into electrical/battery infrastructure leading to ignition.
- **Redundancy Failure**: UNKNOWN
- **Source**: [Baxtel: Google Downtime Due to Paris Global Switch Data Center Fire Caused by Water Leak]()

- **Global Switch Paris Water Leak/Battery Fire 2023 (2023-06)**
- **Root Cause**: Water ingress leading to electrical short/fire in battery infrastructure; impact on Google's European ops indicates significant service dependency or shared infrastructure.
- **Redundancy Failure**: UNKNOWN
- **Source**: [Data Center Dynamics: Global Switch Suffers Battery Room Fire at London Data Center]()

Cause #5: Maintenance/switching error bypassed redundancy

What fails

During routine maintenance or switching operations, human errors can lead to the bypassing of redundant systems designed to ensure continuous operation and resilience. These failures often involve critical components such as power distribution units (PDUs), uninterruptible power supplies (UPSs), generators, cooling systems, and network switches.

How it fails

Maintenance activities are typically performed to ensure that data center infrastructure remains in optimal condition. However, these tasks can introduce risks if not executed correctly. Common errors include:

- Incorrectly disconnecting or reconnecting power or cooling lines.
- Failing to follow established procedures for switching between primary and backup systems.
- Misconfiguring control systems or failing to test failover mechanisms.

Why redundancy does not save the facility

Redundancy is designed to provide a fail-safe mechanism in case of component failures. However, human error during maintenance can bypass these safety measures. For example:

- **Improper Procedures**: Lack of standardized and rigorously tested procedures can lead to mistakes.
- **Lack of Training**: Insufficient training for maintenance personnel can result in incorrect actions.
- **Single Point of Failure (SPOF)**: Even with redundancy, a single error can create an SPOF if the failover mechanisms are not properly configured or tested.

Typical failure chain

```
routine maintenance -> switching error -> redundancy bypassed -> IT load lost power/cooling
```

Frequency in sample

This cause of data center outages ranks 5th, with 5 incidents out of a total of [denominator] incidents in the sample.

Typical downtime

The typical downtime for these incidents ranges from several hours to multiple days, depending on the complexity of the systems involved and the time required to restore redundancy and normal operations.

Severity

These incidents are typically classified as major or critical due to their potential impact on business continuity and data integrity.

Warning signs

- **Procedural Gaps**: Lack of detailed and standardized maintenance procedures.
- **Training Deficiencies**: Inadequate training for maintenance personnel.
- **Testing Shortcomings**: Failure to regularly test failover mechanisms and redundancy systems.
- **Single Points of Failure (SPOFs)**: Identified but unmitigated SPOFs in the infrastructure.

Prevention

To prevent maintenance/switching errors from bypassing redundancy, data centers should implement the following measures:

- **Robust Maintenance Procedures**: Develop and enforce detailed, step-by-step procedures for all maintenance activities.
- **Regular Training**: Provide ongoing training and certification programs for maintenance personnel to ensure they are up-to-date with best practices.
- **Use of Checklists and Double-Checks**: Implement checklists and require double-checks by a second person before critical actions are taken.
- **Regular Testing**: Conduct regular tests of failover mechanisms and redundancy systems to ensure they function as intended.

Representative incidents

Wells Fargo Shoreview DC Fire 2018 (2018-11-29)

- **Root Cause**: Failure of backup power infrastructure (generators) and/or improper failover logic during suppression event; lack of immediate redundancy activation.
- **Redundancy Failed**: Generators did not start for all units ('2 burned. Rest have smoke damage and can't be started'); failovers did not work as expected.
- **Sources**:
- phys.org
- datacenterdynamics.com

PCCW Hong Kong IDC Zone C cooling failure 2022-12-18 (2022-12-18)

- **Root Cause**: Failure of the specific refrigeration unit; facility design did not tolerate single point of failure in cooling infrastructure or redundancy was insufficient.
- **Redundancy Failed**: Not explicitly stated; likely single point of failure in the specific refrigeration loop or common-mode failure in control systems.
- **Source**:
- theregister.com

South Korea Gwacheon Data Center Fire 2014 (2014)

- **Root Cause**: Lack of data backups for the G-Drive system; physical destruction of storage hardware by fire.
- **Redundancy Failed**: No off-site or redundant data copies existed for the primary government drive (G-Drive).
- **Sources**:
- hidrinco.com
- techspot.com

GPU Cluster Power Failure (Squirrel Incident) ()

- **Root Cause**: Design flaw: Single Point of Failure (SPOF) in power distribution; breaker failed to trip and instead welded shut.
- **Redundancy Failed**: Lack of redundancy in power feed to PDU; PDU acted as single point of failure.
- **Source**:
- adhdecode.com

ORX Operational Risk Losses 2018-2023 (2018-2023)

- **Root Cause**: Inadequate or failed internal processes, people, and systems; poor data quality; inconsistent taxonomies.
- **Redundancy Failed**: Not explicitly stated.
- **Source**:
- ey.com

Cause #6: Common-mode design (shared switchgear/cooling/fuel/control)

What fails

In data center outages caused by common-mode design, the failure typically involves shared components such as switchgear, cooling systems, fuel supplies, or control systems. These shared elements are designed to support multiple redundant paths but can become single points of failure when they fail.

How it fails

The failure of a shared component can cascade through the entire system, leading to simultaneous failures across multiple redundant paths. For example, if an electronic switchgear designed to transfer power between mains and backup generators fails, it can prevent the activation of the backup generator, leaving critical IT loads without power. Similarly, a shared cooling system that fails can cause multiple servers to overheat simultaneously.

Why redundancy does not save the facility

Redundancy is often designed to handle single points of failure, but common-mode design issues introduce dependencies that can undermine this redundancy. When redundant paths share components or control systems, a failure in one component can affect all paths, rendering the redundancy ineffective. This lack of true isolation between redundant paths means that the system cannot tolerate multiple simultaneous failures.

Typical failure chain

```
utility power failure -> common-mode design -> multiple systems failed simultaneously -> IT load lost power/coolin
```

Frequency in sample

This cause ranks 6th with a total of 4 incidents out of the sample size, representing approximately 13.3% (4/30) of the analyzed outages.

Typical downtime

The typical downtime for outages caused by common-mode design issues can range from several hours to days, depending on the complexity and scope of the failure. Major or critical incidents often result in extended downtimes as facilities work to diagnose and repair shared components.

Severity

The severity of these outages is typically major or critical, given the widespread impact on multiple systems and the potential for significant data loss or service disruption.

Warning signs

- **Frequent minor issues with shared components**: Recurring problems with switchgear, cooling systems, or control systems can indicate underlying vulnerabilities.
- **Lack of true isolation in redundancy design**: If redundant paths share critical components, it is a red flag that the system may not be as resilient as intended.
- **Inadequate testing and maintenance**: Regular testing and maintenance of shared components are crucial to ensure they function correctly under failure conditions.

Prevention

To prevent outages caused by common-mode design issues, data centers should:

- **Design with true redundancy and isolation**: Ensure that redundant paths do not share critical components or control systems.
- **Conduct thorough risk assessments**: Identify potential single points of failure and implement strategies to mitigate them.

- **Regularly test and maintain shared components**: Perform routine testing and maintenance to ensure that shared components are reliable and can function as intended during failures.

Representative incidents

1. **Interxion LON1 Power Outage Jan 2022 (2022-01-10)**

- **Root Cause**: Failure of electronic switchgear designed to change power over to an on-site generator, preventing switching between mains and backup generator supply.
- **Redundancy Failure**: Electronic switchgear failed to activate/transfer load to the backup generator supply.
- **Sources**:
- LinkedIn Post by Osvaldo Coelho
- Data Center Dynamics News Article

1. **Heathrow Airport Electrical Outage 2025 (2025-03-21)**

- **Root Cause**: UNKNOWN
- **Redundancy Failure**: UNKNOWN
- **Sources**:
- London Daily Article

1. **Tier III Constraint Stacking Failure Mode**

- **Root Cause**: N+1 design cannot tolerate more than one failure; constraint stacking exceeds redundancy capacity.
- **Redundancy Failure**: Simultaneous failures exceeded the single-fault tolerance of the N+1 architecture.
- **Sources**:
- Centiel Article

1. **Northern Virginia Grid Disturbance July 2024 (2024-07)**

- **Root Cause**: Synchronized disconnection of multiple facilities due to lack of Fault Ride-Through (FRT) compliance or similar controls.
- **Redundancy Failure**: The collective redundancy activation created a massive surplus on the grid, requiring urgent adjustments to avoid cascading outages.
- **Sources**:
- Schneider Electric Blog Post

Cause #7: Network Configuration or BGP Change Propagated

What Fails

Network configurations and Border Gateway Protocol (BGP) changes are essential for maintaining the integrity and performance of data center networks. However, when these changes are not properly validated or tested before deployment, they can lead to significant disruptions. Common issues include misconfigurations in routing policies, incorrect BGP announcements, and oversized configuration files that exceed system limits.

How It Fails

The failure typically begins with a network administrator making a change to the network configuration or BGP settings. This change is then propagated across the network, either through automated systems or manual processes. If the change contains errors or misconfigurations, it can cause routing issues, leading to network disruptions and loss of connectivity for IT loads.

Why Redundancy Does Not Save the Facility

Redundancy measures often fail in these scenarios due to a lack of validation and testing before deployment. Critical paths may be misconfigured, and automated systems may propagate errors without immediate detection or rollback mechanisms. Additionally, single points of failure can exist in routing configuration management, preventing effective remediation once an issue occurs.

Typical Failure Chain

```
network configuration change -> BGP propagation -> network disruption -> IT load lost connectivity
```

Frequency in Sample

This cause was identified in 4 out of the total incidents analyzed (4/15).

Typical Downtime

The typical downtime for these incidents ranges from several minutes to a few hours, depending on the complexity of the misconfiguration and the speed of detection and remediation.

Severity

The severity of these incidents is typically major or critical, as they can affect large portions of the network and lead to significant service disruptions.

Warning Signs

- Unexpected changes in network traffic patterns
- Increased error rates in routing tables
- Unusual BGP announcements

- Connectivity issues reported by multiple IT loads

Prevention

To prevent such incidents, organizations should implement the following measures:

- **Automated Validation and Testing**: Use automated tools to validate configuration changes before deployment.
- **Phased Deployment**: Roll out changes gradually to monitor their impact and allow for quick rollback if issues are detected.
- **Circuit Breakers and Rate Limiting**: Implement circuit breakers to prevent cascading failures and rate limiting to control the propagation of changes.
- **Strict Configuration Management**: Ensure that configuration files are validated against strict criteria and that there are safeguards against oversized files or malformed data.

Representative Incidents

Cloudflare Global Outage Nov 2025

- **Date**: 2025-11-18
- **Root Cause**: Lack of strict validation on internal configuration files and insufficient fault isolation in core proxy system allowing propagation of malformed data.
- **Redundancy Failure**: Lack of tolerance for oversized configuration files in the software update process; concentration risk in internet infrastructure.
- **Sources**:
- Cloudflare Explains Major Outage in Detailed Post-Mortem
- Cloudflare CEO Apologizes for Unacceptable Outage and Explains What Went Wrong

Cloudflare BGP Route Leak Jan 2026

- **Date**: 2026-01-22
- **Root Cause**: Policy misconfiguration in JunOS/JunOS EVO: 'route-type internal' match was too permissive, accepting all non-external routes (including IBGP) and advertising them.
- **Redundancy Failure**: Single point of failure in the Miami edge-router configuration; automation executed on a single node without immediate rollback safeguards for policy logic errors.
- **Source**:
- Route Leak Incident January 22, 2026

Facebook BGP Outage Oct 2021

- **Date**: 2021-10
- **Root Cause**: Maintenance command inadvertently removed all external routing paths, including those needed for internal management and physical access systems.

- **Redundancy Failure**: Single point of failure in routing configuration management; loss of control plane prevented remediation.
- **Source**:
- [Cascading Failures]

Azure West US Network Outage Jul 2026

- **Date**: 2026-07-23
- **Root Cause**: UNKNOWN
- **Redundancy Failure**: UNKNOWN
- **Source**:
- [Azure West US Network Outage: Rollback Restores Connectivity]

Cause #8: Fire or Suppression Discharge Damaged Redundant Equipment

What Fails

In data center outages caused by fire or suppression discharge, the primary failure is often the damage to critical infrastructure components such as power distribution units (PDUs), uninterruptible power supplies (UPSs), and cooling systems. These failures can lead to a loss of power or cooling for the IT load, resulting in system downtime.

How It Fails

A fire within a data center can quickly spread due to the high concentration of electrical equipment and flammable materials. When a fire suppression system is activated, it releases agents such as water, gas, or foam to extinguish the fire. However, these agents can also cause significant damage to nearby redundant systems, which are designed to take over in case of primary system failure. The combination of direct fire damage and the effects of the suppression discharge can render both primary and redundant systems inoperable.

Why Redundancy Does Not Save the Facility

Redundant systems fail to save the facility in these scenarios primarily due to a lack of physical separation or adequate protection. If redundant systems are located too close to each other or not properly isolated, they are equally vulnerable to fire and suppression discharge. Additionally, inadequate fire suppression design can exacerbate the damage by failing to contain the fire effectively or by using agents that are harmful to electronic equipment.

Typical Failure Chain

```
fire -> fire suppression discharge -> damage to redundant systems -> IT load lost power/cooling
```

Frequency in Sample

This cause of data center outages ranks 8th, with 4 incidents out of a total sample size. The frequency is \(\frac{4}{\text{total sample size}}\).

Typical Downtime

The typical downtime for these incidents can range from several hours to multiple days, depending on the extent of the damage and the complexity of the recovery process.

Severity

The severity of these outages is typically major or critical. The loss of power or cooling can lead to significant data loss, service disruptions, and financial losses for affected organizations.

Warning Signs

- **Unusual Heat or Smoke:** Early detection of unusual heat or smoke in a data center can indicate the onset of a fire.
- **Fire Suppression System Alarms:** Activation of fire suppression system alarms is a clear warning sign that immediate action is required.
- **BMS Warnings:** Battery Management System (BMS) warnings should not be ignored, as they can indicate potential battery failures that may lead to fires.

Prevention

To prevent outages caused by fire or suppression discharge damaging redundant equipment, the following measures are recommended:

- **Physical Separation of Redundant Systems:** Ensure that primary and redundant systems are physically separated to reduce the risk of both being affected simultaneously.
- **Robust Fire Suppression Designs:** Implement advanced fire detection and suppression systems that can effectively contain fires while minimizing damage to electronic equipment.
- **Regular Testing and Maintenance:** Conduct regular testing and maintenance of fire suppression systems to ensure they are functioning correctly and can be activated when needed.

Representative Incidents

1. **Kakao SK C&C Data Center Fire 2022 (2013-04-01)**

- **Root Causes:**
- Battery failure (manufacturer SK On)
- Potential BMS warning ignored or insufficient response per operator denial
- **Redundancy Failures:**
- Kakao's backup and failover plans were not able to handle the single facility going down.

- **Sources:**
- [Yonhap News Agency]
- [Data Center Dynamics]

1. **Samsung Gwacheon Data Center Fire April 2014 (2014-04-20)**

- **Root Causes:**
- Lack of network redundancy between primary (Suwon) and backup (Gwacheon) sites.
- Lack of remote backup for some servers hosting financial systems; failure to implement geographically redundant disaster recovery.
- **Redundancy Failures:**
- Networking SPOF in backup site meant loss of connectivity to primary production servers.
- **Sources:**
- [Availability Digest]
- [Data Center Knowledge]

1. **Global Switch Paris Water Leak/Battery Fire 2023 (2023-06)**

- **Root Causes:**
- Water ingress leading to electrical short/fire in battery infrastructure.
- Impact on Google's European ops indicates significant service dependency or shared infrastructure.
- **Redundancy Failures:**
- UNKNOWN
- **Sources:**
- [Data Center Dynamics]

1. **Samsung Gwacheon Backup DC Fire (2014-04-20)**

- **Root Causes:**
- Networking infrastructure was located exclusively in the backup data center (Gwacheon) and was not replicated in the primary data center (Suwon), creating a single point-of-failure.
- Single point-of-failure: networking infrastructure was not replicated between primary (Suwon) and backup (Gwacheon) data centers.
- **Redundancy Failures:**
- Networking infrastructure was a single point-of-failure located only in the backup DC; no redundant networking path existed to reroute traffic to the surviving site.
- **Sources:**
- [Availability Digest]

Cause #9: Software bug or configuration error in critical system

What fails

A software bug or configuration error in a critical system can lead to significant disruptions in data center operations. Critical systems include core components such as control planes, load balancers, and other infrastructure services that are essential for the smooth functioning of the data center. When these systems fail due to bugs or misconfigurations, they can cause widespread outages affecting multiple services and customers.

How it fails

Software bugs can manifest in various ways, from logical errors in code to unexpected behavior under certain conditions. Configuration errors often arise from incorrect settings in system configurations, which can lead to improper functioning of the software. These issues can be introduced during software updates or changes in configuration, and they may not be immediately apparent until the system is under load or specific conditions are met.

Why redundancy does not save the facility

Redundancy measures such as failover systems and backup services often fail to mitigate the impact of software bugs or configuration errors for several reasons:

1. **Lack of Proper Testing and Validation**: If the bug or error was introduced during a recent update, it may not have been thoroughly tested in a production-like environment.
2. **Single Point of Failure**: Critical systems that are not properly isolated can become single points of failure. For example, if a control plane is shared across multiple regions, a failure in one region can cascade to others.
3. **Automated Remediation Amplification**: Automated systems designed to mitigate issues can sometimes exacerbate the problem, especially if they lack proper differentiation and circuit breakers.

Typical failure chain

```
software update -> software bug or configuration error -> critical system failure -> IT load lost connectivity/ser
```

Frequency in sample

This cause of data center outages ranks 9th with a frequency of 3 incidents out of the total sample size (denominator not provided).

Typical downtime

The typical downtime for outages caused by software bugs or configuration errors can vary widely, but it often ranges from several hours to days, depending on the complexity of the issue and the effectiveness of the response.

Severity

Outages due to software bugs or configuration errors are typically classified as major or critical. These incidents can have significant financial and reputational impacts on the affected organizations.

Warning signs

- **Unusual System Behavior**: Unexpected behavior in critical systems, such as increased latency or service disruptions.
- **Error Logs**: Frequent error messages or warnings in system logs that indicate potential issues.
- **Testing Failures**: Failed tests during pre-deployment validation phases.
- **Customer Reports**: Increased reports of service degradation or outages from customers.

Prevention

To prevent outages caused by software bugs or configuration errors, organizations should implement the following practices:

1. **Robust Testing and Validation Processes**: Conduct thorough testing in environments that closely mimic production conditions.
2. **Phased Deployment**: Roll out updates gradually to a subset of systems before full deployment to catch issues early.
3. **Automated Rollback Mechanisms**: Implement automated rollback mechanisms to quickly revert changes if issues are detected.
4. **Circuit Breakers and Graceful Degradation**: Design systems with circuit breakers to prevent cascading failures and ensure graceful degradation in case of partial failures.

Representative incidents

- **AWS US-EAST-1 December 2021 Outage (2021-12-07)**:
- **Root Causes**:
- Architectural assumptions of 'multi-region' independence were incorrect; global control planes anchored in us-east-1 created correlated failure modes across regions.
- Control plane infrastructure shared dependencies with operational endpoints; automated remediation (auto-scaling) amplified failures due to lack of differentiation.
- **Redundancy Failures**:
- Automated client-side retry logic acted as a secondary trigger for congestion rather than mitigating it.
- **Source**: [AWS Postmortem]

- **Cloudflare Bot Management Config Failure (X Outage) (2025-11-18)**:
- **Root Causes**:
- Absence of circuit breakers and graceful degradation in the proxy system design.

- Centralized dependency on a single configuration generation mechanism for Bot Management without isolation.
- **Redundancy Failures**: None reported.
- **Source**: [Mashable Article](), [Roundz AI Postmortem]()

- **Google europe-west9-a outage Paris Global Switch (2023-04)**:
- **Root Causes**:
- Leak originated in non-Google portion of the facility, entered UPS room; lack of swift detection/mitigation allowed fire to develop.
- **Redundancy Failures**:
- Fire required evacuation and power shutdown of the entire building; redundancy could not prevent physical destruction or safety shutdown.
- **Source**: [Data Center Dynamics Analysis]()

Cause #10: Environmental conditions exceeded design limits

What fails

When environmental conditions exceed the design limits of a data center, several critical systems can fail, leading to an outage. The primary systems affected are:

- **Cooling Systems**: Designed to maintain optimal temperatures for IT equipment, these systems may not be able to handle extreme heat or cold.
- **Power Supply and Distribution**: Extreme weather conditions can cause power grid disturbances, affecting the reliability of both primary and backup power sources.
- **Environmental Controls**: Sensors and control mechanisms that monitor and regulate temperature, humidity, and other environmental factors may fail under extreme conditions.

How it fails

The failure typically begins with an external event such as extreme heat, cold, or severe weather. These conditions can push the data center's cooling systems beyond their operational limits, leading to a gradual increase in temperature within the facility. As temperatures rise, IT equipment starts to overheat, which can cause hardware failures and system shutdowns.

Why redundancy does not save the facility

Redundancy measures often fail in such scenarios because:

- **Design Not Tolerant of Extreme Conditions**: The data center's design may not account for rare but extreme weather events, making redundant systems ineffective.
- **Lack of Contingency Planning**: Facilities may lack detailed plans and procedures to handle situations where environmental conditions exceed design limits.

Typical failure chain

```
extreme weather -> environmental conditions exceeded design limits -> cooling system failure -> IT load overheated
```

Frequency in sample

This cause of outage has been observed in 3 out of the total incidents analyzed (3/10).

Typical downtime

The typical downtime for such incidents can range from several hours to a few days, depending on the severity of the environmental conditions and the effectiveness of recovery efforts.

Severity

The severity of these incidents is typically major or critical, as they can lead to widespread service disruptions and significant financial losses.

Warning signs

- **Unusual Temperature Fluctuations**: Sudden changes in temperature within the data center.
- **Cooling System Alarms**: Frequent alerts from cooling systems indicating high temperatures or system stress.
- **Power Grid Warnings**: Notifications from utility providers about potential grid disturbances due to extreme weather conditions.

Prevention

To prevent outages caused by environmental conditions exceeding design limits, data centers should:

- **Robust Design for Extreme Conditions**: Ensure that the facility is designed to handle a wide range of environmental conditions, including rare but severe events.
- **Regular Testing Under Stress**: Conduct regular stress tests to simulate extreme conditions and identify potential weaknesses in the system.
- **Use of Advanced Monitoring and Predictive Maintenance**: Implement advanced monitoring systems to detect early signs of environmental stress and use predictive maintenance to address issues before they become critical.

Representative incidents

Google Cloud London Cooling Failure July 2022

- **Date**: 2022-07-19
- **Root Cause**: Cooling system inability to maintain safe temperatures under extreme heat load.

- **Redundancy Failure**: Cooling failure in one building despite backup capacity, necessitating protective shutdown.
- **Sources**:
- [BBC News]
- [Business Insider]

Microsoft Azure Australia East Outage August 2023

- **Date**: 2023-08-30
- **Root Cause**: UNKNOWN
- **Redundancy Failure**: UNKNOWN
- **Sources**:
- [Parametrix Insurance Blog]

Ireland Grid Disturbance May 2025

- **Date**: 2025-05
- **Root Cause**: Collective automatic reduction of grid power draw and switch to UPS by dense cluster of data centers causing supply-demand imbalance.
- **Redundancy Failure**: Redundancy activation (UPS switch) was successful for IT equipment but caused systemic grid instability due to synchronized behavior across many facilities.
- **Sources**:
- [Schneider Electric Blog]

20 Major Incident Case Studies

Uptime Institute 2023 Outage Analysis (Statistical Evidence) (2023-03-09)

What happened:

The Uptime Institute released its 2023 Annual Outage Analysis Report, which compiled data on data center outages from various operators and managers. The report highlighted trends in outage causes and severities.

Timeline:

- **2023-03-09:** Release of the Uptime Institute's 2023 Annual Outage Analysis Report.

Trigger:

- **IT and network problems; power problems; wide-area synchronization issues.**

Primary failure:

- **Complexity in modern architectures, staff shortages leading to procedural non-compliance, increased IT/network complexity causing misconfigurations, strained power/cooling infrastructure designs due to AI demand, and electricity grid limitations.**

Root cause:

- **Increased IT and network complexity leading to change management and misconfiguration errors; use of software-based resiliency tools introducing new risks and complexities.**

Why redundancy failed:

- **Inadequate predictive maintenance for variable operating conditions in cooling infrastructure; loss of visibility due to cloud/outsourcing; lack of transparency from providers.**

Impact & duration:

- **Major IT service outages; severe degradation of service. The report did not specify exact durations but noted that the impact ranged from negligible to catastrophic depending on the incident.**

Engineering lesson:

- **The increasing complexity of modern data center architectures requires robust change management processes, continuous training, and advanced predictive maintenance tools to prevent misconfigurations and ensure reliability.**

Sources:

- [Uptime Institute Annual Outage Analysis 2023 Report]

Uptime Institute Annual Outage Analysis 2021 Trend (2021-03-29)

What happened:

The Uptime Institute published its 2021 Annual Outage Analysis, which identified trends in data center outages from the past year. The report highlighted an increase in IT and network-related issues as a primary cause of outages.

Timeline:

- **2021-03-29:** Release of the Uptime Institute's 2021 Annual Outage Analysis Report.

Trigger:

- **IT and network problems; power problems.**

Primary failure:

- **Complexity in modern architectures, staff shortages leading to procedural non-compliance, increased IT/network complexity causing misconfigurations.**

Root cause:

- **Increased IT and network complexity leading to change management and misconfiguration errors; use of software-based resiliency tools introducing new risks and complexities.**

Why redundancy failed:

- **Inadequate predictive maintenance for variable operating conditions in cooling infrastructure; loss of visibility due to cloud/outsourcing; lack of transparency from providers.**

Impact & duration:

- **Major IT service outages; severe degradation of service. The report did not specify exact durations but noted that the impact ranged from negligible to catastrophic depending on the incident.**

Engineering lesson:

- **The increasing complexity of modern data center architectures requires robust change management processes, continuous training, and advanced predictive maintenance tools to prevent misconfigurations and ensure reliability.**

Sources:

- [Uptime Institute Annual Outage Analysis 2021 Report]

Uptime Institute Annual Outage Analysis 2025 (Aggregate) (2025-06-09)

What happened:

The Uptime Institute released its 2025 Annual Outage Analysis, which aggregated data from multiple sources to provide a comprehensive view of data center outages. The report highlighted the growing complexity of IT and network configurations as a significant factor in outage occurrences.

Timeline:

- **2025-06-09:** Release of the Uptime Institute's 2025 Annual Outage Analysis Report.

Trigger:

- **IT and networking issues; change management failures; misconfigurations.**

Primary failure:

- **Complexity in modern architectures, staff shortages leading to procedural non-compliance, increased IT/network complexity causing misconfigurations.**

Root cause:

- **Increased IT and network complexity leading to change management and misconfiguration errors; use of software-based resiliency tools introducing new risks and complexities.**

Why redundancy failed:

- **Inadequate predictive maintenance for variable operating conditions in cooling infrastructure; loss of visibility due to cloud/outsourcing; lack of transparency from providers.**

Impact & duration:

- **Major IT service outages; severe degradation of service. The report did not specify exact durations but noted that the impact ranged from negligible to catastrophic depending on the incident.**

Engineering lesson:

- **The increasing complexity of modern data center architectures requires robust change management processes, continuous training, and advanced predictive maintenance tools to prevent misconfigurations and ensure reliability.**

Sources:

- [Uptime Institute Annual Outage Analysis 2025 Report]

AAIB G-JEDI AC Electrical Failure (2009-12-21)

What happened

A Bombardier Dash 8 Q402 aircraft (G-JEDI) experienced a failure in its AC electrical system, leading to the loss of both AC generators and extensive fire damage.

Timeline

- **2009-12-21**: The incident occurred during a flight from London Gatwick Airport.
- **Initial Failure**: Chafing of wiring loom by blind rivet head caused arcing and fire.
- **Subsequent Events**: Loss of both AC generators, leading to the loss of critical systems (de-icing, pitot heaters).

Trigger

Chafing of wiring loom by the shear lip of a blind rivet head on a loom support bracket, exacerbated by icing conditions requiring propeller de-icing system activation.

Primary failure

Shearing of input drive shafts on both AC generators and extensive chafing/fire damage to the wiring loom in the left centre-wing section.

Root cause

Design/manufacturing flaw: Blind rivets used for loom support brackets created shear lips that chafed wiring over time, leading to arcing and fire. Lack of adequate insulation/protection at support points.

Why redundancy failed

The chafing and resulting arcing/fire damaged the wiring loom containing power feeds for both AC buses and control circuits, causing a common-mode failure. Both primary generation sources failed mechanically, preventing automatic cross-feeding from the opposite engine's generator.

Impact & duration

- **Impact**: Loss of critical systems (de-icing, pitot heaters), 72 passengers and 4 crew on board.
- **Duration**: UNKNOWN / NOT PUBLICLY DISCLOSED

Engineering lesson

Design choices should avoid using components that can create sharp edges or shear lips. Redundant systems must be physically isolated to prevent common-mode failures.

Sources

- AAIB Report

Google Cloud europe-west3-c power failure Oct 2024 (2024-10-23)

What happened

A power failure at the Google Cloud europe-west3 data center in Frankfurt, Germany, led to a critical outage affecting multiple services.

Timeline

- **2024-10-23**: The incident began with a 3ms voltage drop.
- **Initial Failure**: DRUPS back-up power system for side A failed to take over.
- **Subsequent Events**: Chiller controller dropped offline, causing distribution pumps to stop and the chiller system to shut down.

Trigger

A 3ms voltage drop in the power feed from the utility provider, leading to protective action by the utility breaker.

Primary failure

DRUPS back-up power system for side A failed to take over, resulting in the loss of chilled water flow due to distribution pumps stopping.

Root cause

Failure of the DRUPS back-up power system and subsequent inability to restart distribution pumps, causing a shutdown of the chiller system.

Why redundancy failed

The DRUPS system was designed to provide immediate backup power, but it failed to activate. The lack of a secondary fail-safe mechanism for the distribution pumps led to a complete loss of cooling.

Impact & duration

- **Impact**: Critical services were affected.
- **Duration**: 14 hours 55 minutes

Engineering lesson

Redundant systems must be rigorously tested and maintained. Secondary fail-safes should be implemented to ensure that critical components like distribution pumps can operate even in the event of a primary system failure.

Sources

- UNKNOWN / NOT PUBLICLY DISCLOSED

Google Cloud europe-west4-a cooling/power outage 2026-07-15 (2026-07-15)

What happened

A cooling and power outage at the Google Cloud europe-west4-a data center in the Netherlands led to a major service disruption.

Timeline

- **2026-07-15**: The incident began with an electrical arc flash.
- **Initial Failure**: Chiller controller dropped offline during the voltage transient event.

- **Subsequent Events**: Distribution pumps stopped and were unable to restart, causing the chiller system to shut down.

Trigger

Electrical arc flash in a power distribution unit, likely exacerbated by ongoing construction work at the facility.

Primary failure

Chiller controller dropped offline, leading to the shutdown of distribution pumps and loss of chilled water flow.

Root cause

Electrical arc flash caused by an issue in the power distribution unit, possibly due to construction activities. The chiller controller's inability to handle the transient event led to a complete system shutdown.

Why redundancy failed

The electrical arc flash damaged critical components, including the chiller controller and distribution pumps. Redundant systems were unable to compensate for the simultaneous failure of multiple components.

Impact & duration

- **Impact**: Major services were affected.
- **Duration**: 7 hours 39 minutes

Engineering lesson

Construction activities in data centers should be carefully managed to avoid disruptions to critical infrastructure. Redundant systems must be designed to handle transient events and physical damage to key components.

Sources

- UNKNOWN / NOT PUBLICLY DISCLOSED

Fastly CDN Outage June 8, 2021 (2021-06-08)

What happened:

Fastly experienced a critical outage that affected many high-profile websites and APIs. The incident lasted for approximately 49 minutes, with 85% of the network returning errors.

Timeline:

- **June 8, 2021:** Outage begins.

- **Approximately 49 minutes later:** 95% recovery achieved.

Trigger:

A valid customer configuration change including specific circumstances that triggered a software bug in the edge compute platform.

Primary failure:

Latent software bug in service software causing edge nodes to return errors for unrelated websites.

Root cause:

- Software bug introduced in May 12 deployment.
- Failure in QA/testing processes to detect the bug under specific customer configuration conditions.
- Absence of fleet-wide anomaly detection and rapid global rollback mechanisms for customer-controlled behavior.

Why redundancy failed:

Config validation gaps allowed unsafe behavior to propagate globally. Lack of isolation between customer configurations and core platform stability; insufficient detection during QA.

Impact & duration:

Broad and severe impact to customers relying on Fastly services, including Reddit, gov.uk, Amazon, The New York Times, The Guardian, BBC. Duration: 49 minutes (95% recovery).

Engineering lesson:

Implement robust config validation, staged propagation, and fleet-wide anomaly detection mechanisms to prevent global outages from customer-controlled configurations.

Sources:

- [Wikipedia]
- [The Incident Challenge]
- [Fastly Blog]

Rogers Canada Outage July 2022 (2022-07-05 to 2022-07-08)

What happened:

Rogers Communications experienced a catastrophic outage affecting critical services, including financial transactions and healthcare data access. The incident lasted for up to 26 hours.

Timeline:

- **July 5, 2022:** Outage begins.
- **July 8, 2022:** Service fully restored.

Trigger:

Removal of Access Control List (ACL) policy filter from distribution router configuration during a network upgrade phase.

Primary failure:

Flood of IP routing information into core network routers exceeded their processing capacity, causing the core routers to crash.

Root cause:

- Absence of overload protection mechanisms on core network routers.
- Failure of change management process to flag erroneous configuration due to downgraded risk assessment.
- Dependency on Rogers network for payment processing infrastructure and critical health data access.

Why redundancy failed:

- Shared common IP core for wireless and wireline services created a single point of failure.
- Lack of redundant connectivity from alternative service providers for the Network Operation Centre and critical remote infrastructure sites.
- Management network dependent on the primary IP core; no backup connectivity from alternative providers for critical sites.

Impact & duration:

Severe impact on 2,600+ residents in 10 long-term care homes, Toronto Public Health immunization clinic patients, Interac users (25M daily transactions), and more than 12 million customers. Duration: 19-26 hours.

Engineering lesson:

Implement robust change management processes, overload protection mechanisms, and redundant connectivity to prevent single points of failure in critical infrastructure.

Sources:

- [CRTC Report]
- [Toronto City Council Document]

CrowdStrike Falcon Outage 2024 (2024-07-19)

What happened:

CrowdStrike experienced a critical IT outage due to a defective software content update deployment, affecting global financial markets and other enterprises.

Timeline:

- **July 19, 2024:** Outage begins.
- **Duration:** UNKNOWN / NOT PUBLICLY DISCLOSED

Trigger:

Release of a defective Falcon content update for Microsoft Windows hosts.

Primary failure:

IT system disruption/crash across affected endpoints, leading to operational disruptions in important business services at regulated firms.

Root cause:

- Inadequate pre-deployment testing and lack of phased release mechanisms to contain failures.
- Lack of adequate testing/validation of the content file update leading to widespread system failure.
- Single points of failure in infrastructure and technology stack; over-reliance on a single third-party provider for critical security functions.

Why redundancy failed:

Firms had not adequately identified SPOFs; insufficient third-party categorization and risk assessment processes. High interconnectedness of risks; reliance on third-party security providers for critical operations.

Impact & duration:

Severe impact on global financial markets and regulated firms across all sectors. Duration: UNKNOWN / NOT PUBLICLY DISCLOSED

Engineering lesson:

Implement rigorous pre-deployment testing, phased release strategies, and robust third-party risk management processes to prevent widespread system failures from defective updates.

Sources:

- [EY Report]

North Hyde Substation Outage / Heathrow Airport Closure 2025 (2025-03-20)

What happened:

A catastrophic failure at the North Hyde 275kV substation caused a fire, leading to a power outage that affected Heathrow Airport and surrounding areas. The airport was closed for nearly two days.

Timeline:

- **March 20, 2025:** Fire breaks out at North Hyde Substation.
- **March 21, 2025:** Heathrow Airport closes due to power loss.
- **Late March 21, 2025:** Partial airport operations resume.
- **March 22, 2025:** Full operational recovery achieved.

Trigger:

Catastrophic failure of a transformer's high voltage bushing caused by moisture ingress.

Primary failure:

Destruction/severe damage to main transformers and backup generator at the substation.

Root cause:

Moisture in transformer bushings, detected in 2018, led to an electrical fault. Lack of maintenance for seven years (2018-2025) exacerbated the issue. The fire suppression system was also inoperable.

Why redundancy failed:

The airport's internal network design could not tolerate the loss of one supply point without manual reconfiguration, which took 10-12 hours. Backup generators were on the same circuit as main transformers and were destroyed by the fire.

Impact & duration:

Approximately 48 hours for full recovery. Affected 66,919 domestic and commercial customers, 200,000 passengers at Heathrow, and ~1,000 flights cancelled/disrupted.

Engineering lesson:

Regular maintenance and timely implementation of mitigating actions are crucial. Redundancy should be designed to handle failures without manual intervention.

Sources:

- West London News
- NESO Energy Report

Cloudflare Control Plane Outage Nov 2023 (2023-11-02)

What happened:

An unplanned maintenance by Portland General Electric (PGE) led to a ground fault, causing a power outage at the PDX-DC04 data center. This resulted in a critical control plane and analytics/logging services outage for Cloudflare.

Timeline:

- **November 2, 2023, 11:44 UTC:** Ground fault occurs.
- **November 4, 2023, 04:25 UTC:** Full recovery achieved.

Trigger:

Unplanned maintenance by PGE on one utility feed; subsequent ground fault on PGE transformer.

Primary failure:

Ground fault tripped protective measures, shutting down all 10 generators and the remaining utility feed. UPS batteries failed after ~4 minutes instead of the designed 10 minutes.

Root cause:

Loss of all power generation sources combined with premature UPS battery depletion; inability to manually restart generators due to access control system failure.

Why redundancy failed:

Generators tripped offline by ground fault protection. UPS failed prematurely (4 mins vs 10 mins design). Access control system lost power, preventing manual generator restarts. Services not in the HA cluster had no redundancy path when PDX-DC04 failed.

Impact & duration:

40 hours and 41 minutes. Affected Cloudflare customers using control plane and analytics/logging services.

Engineering lesson:

Ensure robust access control systems for critical infrastructure. Regularly test UPS and generator systems to ensure they meet design specifications.

Sources:

- [Cloudflare Blog]

Optus National Outage 2023 (2023-11-08)

What happened:

A routine software upgrade caused key routers to disconnect from the network, leading to a nationwide outage for Optus.

Timeline:

- **November 8, 2023:** Routine software upgrade triggers routing information updates.
- **Duration:** UNKNOWN / NOT PUBLICLY DISCLOSED

Trigger:

Routine software upgrade.

Primary failure:

Routing information updates from an international peering network caused key routers to disconnect from the network.

Root cause:

Key routers disconnected due to routing updates.

Why redundancy failed:

UNKNOWN / NOT PUBLICLY DISCLOSED

Impact & duration:

Affected approximately 10 million Optus retail and 400,000 business customers. Duration UNKNOWN / NOT PUBLICLY DISCLOSED.

Engineering lesson:

Thorough testing of software upgrades in a controlled environment before deployment is essential to avoid widespread outages.

Sources:

- [Journal of Telecommunications and the Digital Economy]

Cloudflare BYOIP BGP Outage Feb 20 2026 (2026-02-20)

What happened:

A configuration change to automate the removal of BYOIP prefixes via a cleanup sub-task led to an unintentional withdrawal of ~1,100 BYOIP prefixes via BGP.

Timeline:

- **February 20, 2026:** Configuration change applied.
- **Duration:** 6 hours and 7 minutes

Trigger:

Configuration change to automate removal of BYOIP prefixes via a cleanup sub-task.

Primary failure:

Secondary software bug removed addressing settings from edge servers for ~300 prefixes. System interpreted all returned prefixes as queued for deletion due to API query bug.

Root cause:

Software bug in the Addressing API where an empty 'pending_delete' query parameter was interpreted as a request for all prefixes instead of only pending deletions. Software bug in the edge configuration propagation logic that deleted state for specific users during remediation attempts.

Why redundancy failed:

Automated cleanup sub-task querying the API; lack of immediate rollback before impact was observed. Reliance on customer self-remediation via dashboard which failed for affected prefixes; manual restoration required.

Impact & duration:

6 hours and 7 minutes. Affected Cloudflare customers using BYOIP services.

Engineering lesson:

Implement robust testing and validation processes for configuration changes, especially those affecting critical systems. Ensure immediate rollback mechanisms are in place to minimize impact.

Sources:

- UNKNOWN / NOT PUBLICLY DISCLOSED

Dataset Overview

Metric	Value
Signals (deduplicated)	506
Clustered incidents	184
Incidents with A/B evidence	59
Incidents where redundancy failed	99
Signals flagging human involvement	182
Signals flagging maintenance involvement	118
Downtime data	{"n": 135, "avg": 1138, "median": 420, "total": 153670}

Unique sources	185

Incidents by year: {"2003": 2, "2007": 1, "2009": 1, "2011": 2, "2012": 4, "2013": 5, "2014": 4, "2015": 6, "2016": 12, "2017": 9, "2018": 7, "2019": 4, "2020": 6, "2021": 9, "2022": 17, "2023": 13, "2024": 13, "2025": 28, "2026": 18}

Incidents by severity: {"major": 111, "critical": 74, "catastrophic": 13, "minor": 10, "significant|serious|severe": 4, "major|critical": 3, "partial|critical|null": 3, "various": 2, "severe": 2, "warning": 1, "minor to major": 1, "Category 5 (Severe) to Category 1 (Negligible)": 1, "Outage/Incident": 1, "major|critical|null": 1, "major/critical": 1, "minor|major": 1}

Incidents by country (top): {"UK": 23, "USA": 19, "US": 13, "Australia": 12, "United States": 12, "France": 8, "Germany": 7, "United Kingdom": 6, "South Korea": 5, "Canada": 4, "Global": 4, "Netherlands": 3, "North America": 2, "UK/Europe": 2, "United Kingdom/Europe": 2}

Methodology / Limitations

Reusable n8n Deep Research Engine (dedicated clone EW13X5QPciEAYWkO); 309 search queries incl. SERP pagination and DE/FR/NL/IT variants; page fetch + strict-JSON extraction (qwen3.6-35B GPU lane, CPU lane for PDF documents); incident clustering by stable incident names (fuzzy); trigger / immediate failure / root cause / contributing factors captured as separate fields; confidence A-D; deterministic statistics with visible denominators; qwen2.5-72B for ranking, chapters, case studies and synthesis. Limitations: public postmortems over-represent cloud/network incidents; many facility-level outages are never disclosed; durations and financial figures are sparse; percentages describe the researched sample only.

Sources (185 unique URLs)

- http://docs.media.bitpipe.com/io_10x/io_109802/item_692197/DBOY-7CDJNW_R0_EN.pdf
- http://files.server-rack-online.com/2016-Cost-of-Data-Center-Outages.pdf
- https://adhdecode.com/ai-infrastructure/real-incident-case-studies/gpu-cluster-power-failure
- https://africa.businessinsider.com/tech/the-wells-fargo-outage-was-caused-by-an-automatic-power-shutdown-at-its-data-center/jgdrre0
- https://aiqlabs.ai/blog/what-causes-ats-to-fail
- https://archilabs.ai/posts/fixing-data-handover-failures-in-data-center-commissioning
- https://arstechnica.com/gadgets/2025/10/a-single-point-of-failure-triggered-the-amazon-outage-affecting-millions
- https://assets.publishing.service.gov.uk/media/5422ee84e5274a1317000219/DHC-8-402_Dash_8__G-JEDI_08-10.pdf
- https://availabilitydigest.com/public_articles/0907/samsung.pdf
- https://aviationsourcenews.com/heathrow-blackout-how-a-years-old-unaddressed-fault-caused-chaos
- https://aws.amazon.com/message/12721
- https://awsdown.com/outages/aws-us-east-1-outage-december-2021

- https://baxtel.com/news/google-downtime-due-to-paris-global-switch-data-center-fire-caused-by-water-leak
- https://blog.cloudflare.com/analysis-of-todays-centurylink-level-3-outage
- https://blog.cloudflare.com/cloudflare-outage-february-20-2026
- https://blog.cloudflare.com/post-mortem-on-cloudflare-control-plane-and-analytics-outage
- https://blog.cloudflare.com/route-leak-incident-january-22-2026
- https://blog.se.com/datacenter/2026/02/27/data-centers-grid-friendly-preventing-blackout-grids-stability-resilience
- https://circleid.com/posts/cloudflare-explains-major-outage-in-detailed-post-mortem
- https://cloudnews.tech/a-major-fire-forces-the-evacuation-of-northc-data-center-in-almere
- https://corporate.ovhcloud.com/en/newsroom/news/informations-site-strasbourg
- https://corporate.ovhcloud.com/en/newsroom/news/statement-following-two-incidents-9th-november-2017
- https://crtc.gc.ca/eng/publications/reports/xona2024.htm
- https://cyberpress.org/cloudflare-confirms-bgp-attack
- https://cyberpress.org/global-outage-hits-x-ddos-attack
- https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/2026.AnnualOutageAnalysis.pdf
- https://datacenter.uptimeinstitute.com/rs/711-RIA-145/images/AnnualOutageAnalysis2023.03092023.pdf
- https://datacenters.economictimes.indiatimes.com/news/operations-resilience/ibm-cloud-datacenter-suffers-outage-after-amsterdam-fire/130946258
- https://eastlondonpost.co.uk/firefighters-tackle-battery-fire-at-blackwall-data-centre
- https://eastlondontimes.co.uk/fire/60-firefighters-battle-blackwall-data-centre-blaze
- https://eastlondontimes.co.uk/local/firefighters-battle-blackwall-tower-data-centre-blaze
- https://ecl-ips.com/wp-content/uploads/2014/01/apc-data-centre-temperature-rise-during-a-cooling-system-outage.pdf
- https://edisplay.it/outage.html
- https://edmontonsun.com/2012/07/11/fire-inside-calgary-shaw-building-sparks-communications-breakdown
- https://en.wikipedia.org/wiki/Fastly
- https://en.wikipedia.org/wiki/Hayes_substation_fire
- https://en.wikipedia.org/wiki/OVHcloud
- https://en.yna.co.kr/view/PYH20221017112300315
- https://engineering.fb.com/2021/10/05/networking-traffic/outage-details
- https://forums.theregister.com/forum/all/2025/05/19/oci_outage_europe
- https://fr.wikipedia.org/wiki/Incendie_du_centre_de_donn%C3%A9es_d%27OVHcloud_%C3%A0_Strasbourg
- https://geoblackout.com/au/report/internet/telstra
- https://gigazine.net/gsc_news/en/20230906-microsoft-azure-incident
- https://grokipedia.com/page/2023_Optus_outage
- https://grokipedia.com/page/Emergency_Power_Off

- https://hidrinco.com/fire-safety-concepts-for-data-centers-response-to-realistic-fire-scenarios
- https://hostingdiscussion.com/news/oracle-faces-fresh-cloud-headwinds-as-european-identity-outage-sparks-reliability-concerns
- https://ifactoryapp.com/predictive-maintenance/data-center-avoids-2-8m-outage-chiller-predictive-maintenance
- https://intelligence.uptimeinstitute.com/resource/annual-outage-analysis-2025
- https://intelligence.uptimeinstitute.com/resource/annual-outage-analysis-2026
- https://intelligence.uptimeinstitute.com/sites/default/files/mammoth/2023/09/12/Keynote%20report%2023_Annual%20outage%20analysis%202019.pdf
- https://johal.in/postmortem-github-2026-outage-caused-2-hour-downtime-internal
- https://journal.uptimeinstitute.com/data-center-operators-will-face-more-grid-disturbances
- https://jtde.telsoc.org/index.php/jtde/article/view/898
- https://londondaily.com/heathrow-airport-closed-following-electrical-substation-fire-in-west-london
- https://londondaily.com/heathrow-airport-closure-due-to-substation-fire-causes-major-disruption
- https://mashable.com/article/cloudflare-outage-what-happened-explained
- https://mlq.ai/news/cloudflare-outage-caused-by-bug-in-generation-system
- https://myappapi.com/blog/postmortems/aws-us-east-1-2021
- https://outagecost.com/case-studies/delta-2016
- https://outagecost.org/case-studies/delta-2016
- https://phys.org/news/2018-06-visa-million-payments-affected-june.pdf
- https://phys.org/news/2019-02-center-outage-wells-fargo-customers.pdf
- https://ponemonsullivanreport.com/2021/02/data-center-downtime-at-the-core-and-the-edge-a-survey-of-frequency-duration-and-attitudes
- https://regmedia.co.uk/2022/06/10/ovh_report.pdf
- https://roundz.ai/blog/postmortem-deep-dive-cloudflare-november-2025-outage
- https://rs.linkedin.com/posts/osvaldo-coelho-6b61425_interxion-hit-with-outage-at-central-london-activity-6988463361789169664-hgcD
- https://starlinepower.com/sites/default/files/files/DPC-BUS-WP-Arc-Flash-White-Paper-US-033123.pdf
- https://status.cloud.google.com/incidents/3BvH3LVGcupoYqV6F4Nw
- https://status.cloud.google.com/incidents/e3yQSE1ysCGjCVEn2q1h
- https://theincidentchallenge.com/famous-production-incidents/fastly-2021-cdn-outage
- https://undercodenews.com/massive-global-outages-hit-elon-musks-x-platform-for-the-second-day-in-a-row
- https://uptimeinstitute.com/about-ui/press-releases/uptime-announces-annual-outage-analysis-report-2025
- https://westlondon.com/heathrow-shutdown-final-neso-report-on-north-hyde-substation-outage
- https://windowsforum.com/windows-news.4/azure-west-us-network-outage-rollback-restores-connectivity.440191
- https://wpde.com/news/nation-world/substation-fire-at-londons-heathrow-airport-knocks-out-power-disrupts-thousands-of-flights-international-travel-hub-diverted-gatwick-charles-de-gaulle-shannon-united-k

ingdom
- https://www.alertlabs.com/post/water-leak-detection-data-centers-stop-damage
- https://www.bbc.co.uk/news/live/cly24zvvwxlt
- https://www.bbc.com/news/articles/c6283577llqo
- https://www.bbc.com/news/articles/cvg5dg4p2l0o
- https://www.bbc.com/news/technology-62202125
- https://www.bellperformance.com/bell-performs-blog/data-center-generator-fuel-quality-the-complete-guide-to-preventing-backup-power-failures
- https://www.bleepingcomputer.com/news/security/ovh-data-center-fire-likely-caused-by-faulty-ups-power-supply
- https://www.businessinsider.com/google-cloud-data-center-london-outage-hottest-day-record-uk-2022-7
- https://www.businessinsider.com/wells-fargo-outage-caused-by-power-outage-at-data-center-after-smoke-detected-2019-2
- https://www.businesskorea.co.kr/news/articleView.html
- https://www.cbsnews.com/detroit/news/macomb-county-suffers-without-backup-data-center
- https://www.centiel.com/redundancy-failures-tier-iii
- https://www.claimsjournal.com/news/national/2026/04/01/336544.htm
- https://www.cnet.com/tech/services-and-software/cloudflare-ceo-apologizes-for-unacceptable-outage-and-explains-what-went-wrong
- https://www.computerweekly.com/news/252443325/Visa-reveals-rare-datacentre-switch-fault-as-root-cause-of-June-2018-outage
- https://www.computerweekly.com/news/450420405/The-British-Airways-IT-outage-What-went-wrong-with-its-datacentre
- https://www.crisisboardroom.com/case-study-delta
- https://www.dailymail.com/news/article-3734898/Delta-computer-outage-included-small-fire-data-center.html
- https://www.datacenterdynamics.com/en/analysis/data-center-water-leaks
- https://www.datacenterdynamics.com/en/analysis/ovhcloud-fire-france-data-center
- https://www.datacenterdynamics.com/en/news/court-dismisses-s24m-global-switch-lawsuit-against-consultant-engineers-over-data-center-power-outage
- https://www.datacenterdynamics.com/en/news/details-emerge-of-global-switch-outage-in-london
- https://www.datacenterdynamics.com/en/news/elon-musks-x-suffers-another-major-outage
- https://www.datacenterdynamics.com/en/news/equinix-ld8-data-center-experiences-major-outage
- https://www.datacenterdynamics.com/en/news/global-switch-suffers-battery-room-fire-at-london-data-center
- https://www.datacenterdynamics.com/en/news/googles-london-data-center-outage-during-heatwave-caused-by-simultaneous-failure-of-multiple-redundant-cooling-systems
- https://www.datacenterdynamics.com/en/news/microsofts-slow-outage-recovery-in-sydney-due-to-insufficient-staff-on-site

- https://www.datacenterdynamics.com/en/news/oracle-cloud-experiences-outage-in-europe-users-say
- https://www.datacenterdynamics.com/en/news/outage-at-interxion-facility-in-london-takes-down-metals-exchange
- https://www.datacenterdynamics.com/en/news/ovh-fire-update-four-halls-sbg1-destroyed-well-all-sbg2
- https://www.datacenterdynamics.com/en/news/ovh-suffers-two-hour-outage-in-two-locations-due-to-software-bug-electrical-issue
- https://www.datacenterdynamics.com/en/news/ovhcloud-sbg2-fire-report-reveals-water-was-detected-near-power-inverters
- https://www.datacenterdynamics.com/en/news/pccw-data-center-refrigeration-equipment-failure-causes-alibaba-cloud-hong-kong-outage
- https://www.datacenterdynamics.com/en/news/sk-group-takes-full-responsibility-for-data-center-fire-that-led-to-kakao-outage
- https://www.datacenterdynamics.com/en/news/uks-queen-alexandra-hospital-declares-critical-incident-as-chillers-fail-at-on-site-data-center
- https://www.datacenterdynamics.com/en/news/visa-details-cause-of-widespread-outage-blames-data-center-switch-failure
- https://www.datacenterdynamics.com/en/news/widespread-wells-fargo-outage-blamed-data-center-fire
- https://www.datacenterfrontier.com/home/whitepaper/11432242/cost-of-data-center-outages
- https://www.datacenterknowledge.com/cooling/google-oracle-data-centers-knocked-offline-by-london-heat
- https://www.datacenterknowledge.com/outages/-update-2-data-center-outage-disrupts-wells-fargo-atms-mobile-app-website
- https://www.datacenterknowledge.com/outages/data-center-fire-disrupts-key-services-in-calgary
- https://www.datacenterknowledge.com/outages/data-center-fire-leads-to-outage-for-samsung-devices
- https://www.datacenterknowledge.com/outages/generator-failures-caused-365-main-outage
- https://www.datacenterknowledge.com/outages/study-data-center-downtime-costs-7-900-per-minute
- https://www.datacenterknowledge.com/uptime/fire-has-destroyed-ovh-s-strasbourg-data-center-sbg2-
- https://www.dazzlingdawn.com/2026/01/24/burning-evidence-the-convenient-fire-at-blackwall
- https://www.dpstele.com/blog/2025-outage-report-changing-risks.php
- https://www.europesays.com/netherlands/10287
- https://www.ey.com/content/dam/ey-unified-site/ey-com/en-uk/industries/banking-capital-markets/documents/ey-uk-operational-risk-01-2025.pdf
- https://www.fastly.com/blog/summary-of-june-8-outage
- https://www.fca.org.uk/firms/operational-resilience/crowdstrike-outage-lessons-operational-resilience
- https://www.groundcover.com/learn/performance/cascading-failures
- https://www.idcnova.com/news/fire-breaks-out-at-northc-data-center-near-amsterdam-firefighters-still-battling-blaze-5757
- https://www.infoq.com/news/2021/12/aws-outage-postmortem
- https://www.irecruit.co/insights/bms-cybersecurity-in-data-centers-risks-standards-mitigation
- https://www.itechguides.com/x-outage-on-november-18-2025-what-caused-the-major-disruption

- https://www.itnews.com.au/news/microsoft-had-three-staff-at-australian-data-centre-campus-when-azure-went-out-599849
- https://www.itpartners.co.nz/outages/azure-outage-australia-east-region
- https://www.koreatimes.co.kr/business/companies/20221020/kakao-sk-in-blame-game-over-fire-at-data-center
- https://www.lakeletag.com/blog/2026/7/23/case-study-financial-litigation-support-quantifying-economic-damages-from-a-data-center-fire
- https://www.linkedin.com/posts/karthikeyan-dc-sme_datacenter-transformersafety-firemitigation-activity-7353968210456539136-qCgd
- https://www.linkedin.com/posts/lucas-buthion_telecom-security-incidents-2024-activity-7351506025134366720-5Z4h
- https://www.linkedin.com/posts/mark-e-peters_im-sharing-this-because-%3F-%3F%3F%3F%3F%3F%3F-activity-7409223005823950850-pMK4
- https://www.linkedin.com/posts/totaluptime_a-power-outage-on-monday-at-a-central-london-activity-6887373812703260672-1Fi7
- https://www.linkedin.com/pulse/ups-systems-failure-due-accidentally-touched-epo-switch-v0wmc
- https://www.linkedin.com/pulse/why-redundancy-isnt-enough-hidden-gaps-data-center-power-aqeel-anwar-mrzpf
- https://www.mcmorrowreports.com/uptimes-7th-annual-data-center-outage-analysis-shows-improvement-but-new-risks
- https://www.mirror.co.uk/news/uk-news/breaking-massive-blaze-breaks-out-36603658
- https://www.modb.pro/doc/120607
- https://www.mprnews.org/story/2019/02/07/wells-fargo-data-center-smoke-outage
- https://www.naco.org/resources/award-programs/under-fire-disaster-recovery-and-business-continuity
- https://www.neso.energy/news/final-report-review-north-hyde-substation-outage
- https://www.networkworld.com/article/4171277/network-outages-power-failures-strain-data-center-resiliency.html
- https://www.networkworld.com/article/957256/microsoft-blames-aussie-data-center-outage-on-staff-strength-failed-automation.html
- https://www.onesto-ep.com/blog/8-common-ats-failures-their-root-causes-and-proven-fixes
- https://www.panorama-consulting.com/lufthansas-software-failure-lessons
- https://www.parametrixinsurance.com/blog/microsoft-azure-australia-outage-parametrix-event-analysis
- https://www.parkplacetechnologies.com/blog/causes-of-data-center-outages-how-to-prevent-them
- https://www.retail-systems.com/rs/Visa_Data_Centre_Malfunction.php
- https://www.sdxcentral.com/news/ibm-joins-the-outage-club-after-quantum-cloud-service-unexpectedly-went-down
- https://www.singularitymoments.com/content/major-gpu-cluster-failure-hammers-cloud-provider-inference-throughput
- https://www.structureresearch.net/2017/11/13/ovh-suffers-outage-france-hitting-strasbourg-data-centres

- https://www.structureresearch.net/2018/07/04/visa-europe-outage-due-to-switch-failover-failure
- https://www.techerati.com/news-hub/ups-failure-behind-equinix-ld8-data-centre-outage
- https://www.techmonitor.ai/hardware/data-centres/equinix-outage
- https://www.technewsworld.com/story/Data-Center-Fire-Fries-Samsungs-Un-Backed-Up-Servers-80333.html
- https://www.techspot.com/news/109799-858tb-gone-forever-south-korea-data-center-fire.html
- https://www.techstories.co/liquid-cooling-leak-destroys-millions-of-dollars-in-gpus
- https://www.techstories.co/what-went-wrong-in-south-koreas-government-data-centre-fire
- https://www.theguardian.com/money/2018/jun/19/visa-admits-5m-payments-failed-over-a-broken-switch
- https://www.theoaklandpress.com/2013/04/17/fire-forces-evacuation-of-two-county-buildings
- https://www.theregister.com/off-prem/2022/12/20/hong-kong-datacenter-loses-its-cool-takes-out-alibaba/644098
- https://www.theregister.com/off-prem/2025/05/19/oracle-cloud-infrastructure-outage-reported/459037
- https://www.theregister.com/off-prem/2025/06/03/ibm-cloud-login-breaks-for-second-time-in-a-fortnight/276902
- https://www.theregister.com/off-prem/2025/10/30/ibm-cloud-is-experiencing-a-quantum-computer-outage/553578
- https://www.theregister.com/on-prem/2020/08/18/outage-faulty-ups-at-data-centre-housing-london-internet-exchange-causes-grief-for-isps-and-telcos-alike/530245
- https://www.thestack.technology/data-centre-fire-knocks-ibms-cloud-service-offline
- https://www.thestack.technology/interxion-outage-power-lon1-london-metals-exchange
- https://www.thinkbrg.com/thinkset/when-cyber-meets-concrete-what-a-data-center-outage-reveals-about-cybersecurity-and-business-insurance
- https://www.toronto.ca/legdocs/mmis/2023/ex/bgrd/backgroundfile-239381.pdf
- https://www.uktech.news/cloud/uk-heatwave-data-centres-20220720
- https://www.velir.com/-/media/files/pdfs/uptime-annualoutageanalysis2021.pdf
- https://www.ventrogroup.com/blog/why-heathrow-went-dark
- https://www.vertiv.com/490372/globalassets/documents/reports/ponemon/vertiv-ponemon-data-center-downtime-survey-report_321974_0.pdf
- https://www.vertiv.com/490a6e/globalassets/documents/reports/ponemon/vertiv-downtimecostspositionpaper-wp-en-na-sl-70705-web_321837_0.pdf
- https://www.vertiv.com/en-emea/about/news-and-events/news-releases/2016/emerson-network-power-study-says-unplanned-data-center-outages-cost-companies-nearly-$9000-per-minute
- https://www.vertiv.com/globalassets/documents/reports/2016-cost-of-data-center-outages-11-11_51190_1.pdf
- https://www.vocus.com.au/help-and-support/historical-outage-register
- https://www.wired.com/2012/07/explosion-shaw-ibm

Appendix: collected evidence