Understanding systemic cloud events: An overview for insurance professionals

A CYBER INDUSTRY STEERING GROUP WHITE PAPER BY AXIS, GALLAGHER RE, MOODY'S, AND MUNICH RE

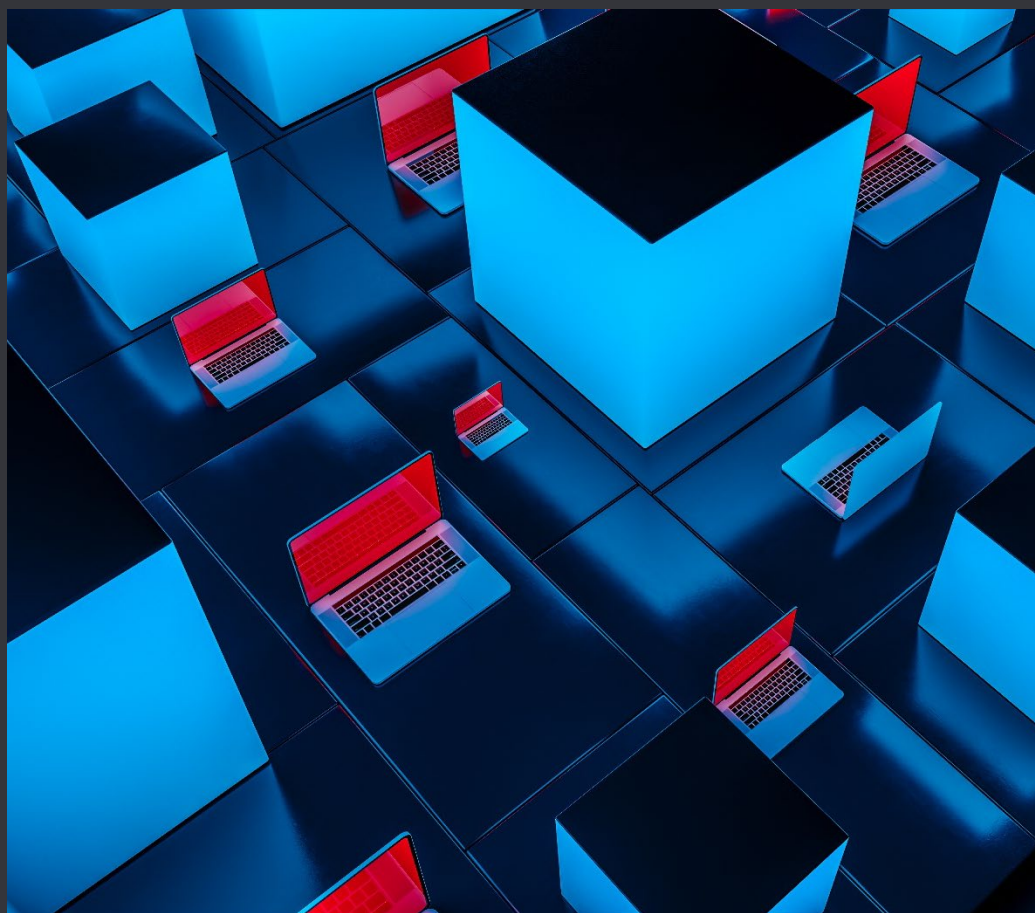

The information contained in this white paper is derived from public sources and general market knowledge. No company-specific price, modelling, customer, or any other information was used to inform this white paper or its conclusions. This white paper is for informational purposes only.

EXECUTIVE SUMMARY

On October 19, 2025, a prolonged outage at AWS's Northern Virginia facilities, which host critical infrastructure for global cloud services, triggered widespread service disruptions, affecting businesses and users across multiple continents. This event demonstrates that cloud outages can potentially pose a significant accumulation risk to the cyber insurance industry. However, within the industry, there can be limited understanding of how the underlying infrastructure is built and operated. A limited visibility can make it harder to assess how failures unfold and how insurable losses may emerge once an outage has taken place.

The underlying infrastructure and operating procedures of cloud technology are highly complex, interdependent, and built with layers of redundancy at scale. As a result, there are numerous ways a cloud outage can manifest. Failures may occur within specific products, individual data centers, or across entire geographic regions. They can happen abruptly or escalate gradually. In some cases, the system may remain technically functional but become effectively unusable due to latency or delays. Restoration timelines also vary depending on the nature of the fault, ranging from immediate recovery to prolonged resolution. Given this diversity in failure modes, it is entirely reasonable that professionals within the cyber insurance industry find it challenging to develop a solid grasp of how cloud infrastructure operates and failures unfold.

For insurance buyers, the impact of a cloud outage is rarely uniform. The same disruption can manifest in vastly different ways depending on the nature of the outage, the design of each company's cloud-based infrastructure, and each company's dependency on that infrastructure. This variability adds another layer of complexity for insurers assessing potential losses. At the same time, it presents a business opportunity: Companies that invest in resilient architectures and recovery strategies are better positioned to weather outages and limit the severity of business interruptions.

The goal of this paper is to provide the reader with a foundational understanding of cloud infrastructure and the mechanisms through which cloud outages occur. By demystifying the technical architecture and offering a structured framework for interpreting outage reports, we aim to bridge the knowledge gap between cloud technology and cyber insurance risk assessment. Readers will explore a structured framework that categorizes cloud outages and is informed by insights from cloud service provider (CSP) status reports. This approach is essential for evaluating potential business interruption scenarios and designing more informed coverage solutions in an increasingly cloud-dependent economy.

INTRODUCTION

The insurance problem

Cyber insurance is complex, encompassing both first-party and third-party protections. Although its origins lie in third-party liability, the primary concern for the (re)insurance industry today centers on the potential for aggregation of losses, most notably from the first-party business interruption component. Identifying what could trigger such aggregation and how to quantify it is not straightforward. Unlike natural catastrophe risk, cyber risk is evolving. It lacks historical depth and has not yet experienced a defining event comparable to Hurricane Andrew, which would validate systemic cyber risk and offer a reference point for calibration. Despite these challenges, there is broad consensus among model vendors and (re)insurers that two types of events are most likely to drive significant insured aggregation: cloud outages and widespread systemic malware or ransomware. Cloud outages are often discussed, and due to the complexity of cloud infrastructure - its scale, interdependence, and operational opacity - understanding the nature of this risk requires careful consideration within the insurance industry. This paper aims to balance technical depth with accessible language, offering a clear explanation of how cloud infrastructure functions and how it can fail.

What is cloud computing?

CSPs offer a range of services with varying levels of complexity, all designed to meet the diverse IT needs of modern organizations. These providers include global leaders such as AWS, Microsoft Azure, and Google Cloud Platform, as well as regional players like Alibaba and OVHcloud and specialized vendors such as Oracle and IBM. Regardless of their positioning, CSPs ultimately aim to deliver scalable, flexible, and efficient solutions that support enterprise operations.

Cloud computing has changed how organizations access, manage, and scale their computing resources. Although personal and desktop devices remain essential for day-to-day operations, they are often limited in processing power, storage capacity, and expandability. These constraints can hinder their ability to support modern business demands, especially in data-intensive or mission-critical environments.

Before cloud computing became mainstream, organizations relied on “on-premises” systems like clusters and grids to meet their computational and storage needs. These setups pooled resources across machines but required substantial investment and ongoing maintenance. Cloud computing evolved from these approaches by offering on-demand access to large pools of computing, storage, and networking resources hosted in data centers located across multiple regions. This shift allows firms to scale IT capabilities flexibly, manage costs more efficiently, and innovate faster without heavy reliance on physical infrastructure.

Cloud service providers offer different types of cloud computing services, most commonly infrastructure as a service (IaaS), platform as a service (PaaS), and software as a service (SaaS). These models vary in how responsibilities are shared between the provider and the customer. While the provider always manages the core infrastructure, tasks such as networking, system maintenance, and software deployment may fall to the customer depending on the chosen model. As a result, companies using different service models may experience the same cloud outage in very different ways. This variation contributes to the complexity of assessing exposure and highlights the importance of understanding how cloud services are configured and utilized.

In this white paper, we provide an overview of the core components that constitute cloud computing infrastructure. By understanding the cloud’s architecture and operational dynamics, insurers and risk managers can more accurately assess exposure, price policies, and develop robust risk mitigation strategies tailored to the evolving digital ecosystem.

DATA CENTERS, AVAILABILITY ZONES, AND REGIONS

At the foundation of cloud computing are data centers — specialized facilities that house the physical infrastructure powering cloud services. These data centers perform three essential functions: computation, storage, and connectivity.

- **Servers** run applications and process computational workloads.
- **Storage** systems ensure data is safely kept and accessible, whether for short-term use or long-term retention.
- **Connectivity** involves high-speed networking equipment that links servers and storage within the data center and connects them to the internet and private networks. This includes tools for secure access, protection against denial-of-service attacks, and services that help distribute content efficiently.

Together, these components form the backbone of cloud platforms, enabling scalable, resilient, and secure access to digital resources across geographically distributed locations.

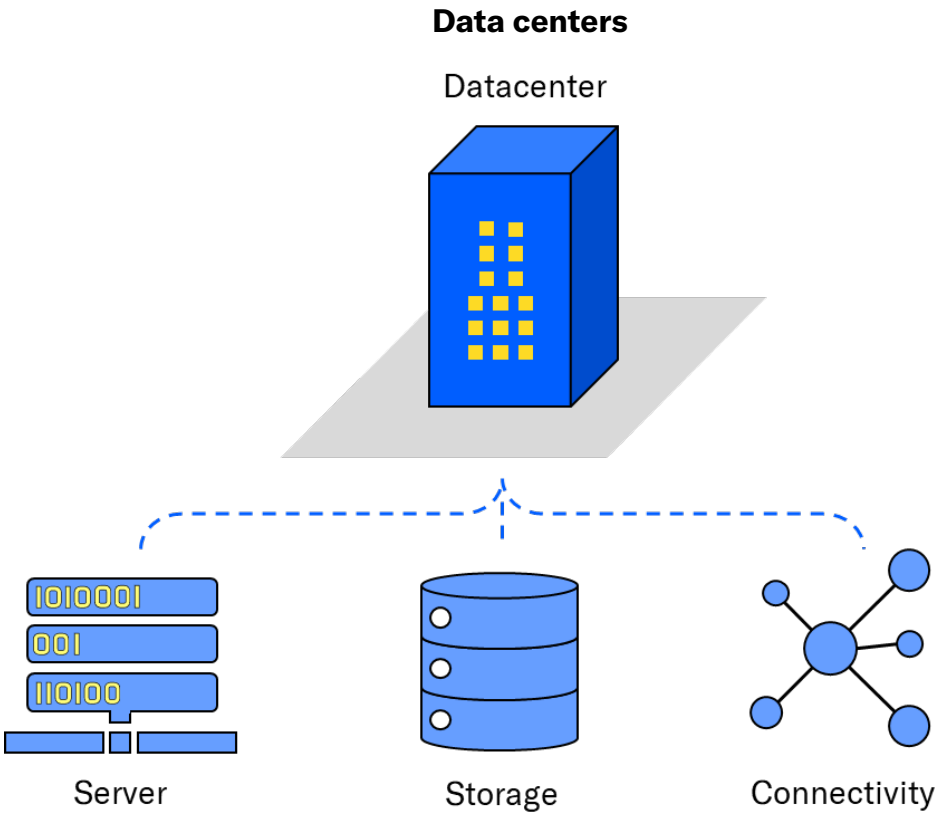

The top image illustrates a building housing a data center. The bottom images depict the primary device types commonly found within a data center: servers (responsible for computation), storage systems (used for managing and storing data), and connectivity devices (which enable communication both within the data center and externally).
Source: Microsoft Azure

To enhance resilience and reduce the risk of service disruption, cloud providers organize their infrastructure using a layered approach. At the first level, multiple data centers are grouped into what are known as availability zones. These zones are typically located several kilometers apart but generally within a 100-kilometer radius. Each availability zone is designed to operate

independently, with its own power supply, cooling systems, and network connections. This separation ensures that a localized issue, such as a power outage or hardware failure, affecting one zone should not impact others. This architectural design is a key element in maintaining high availability and supporting disaster recovery strategies.

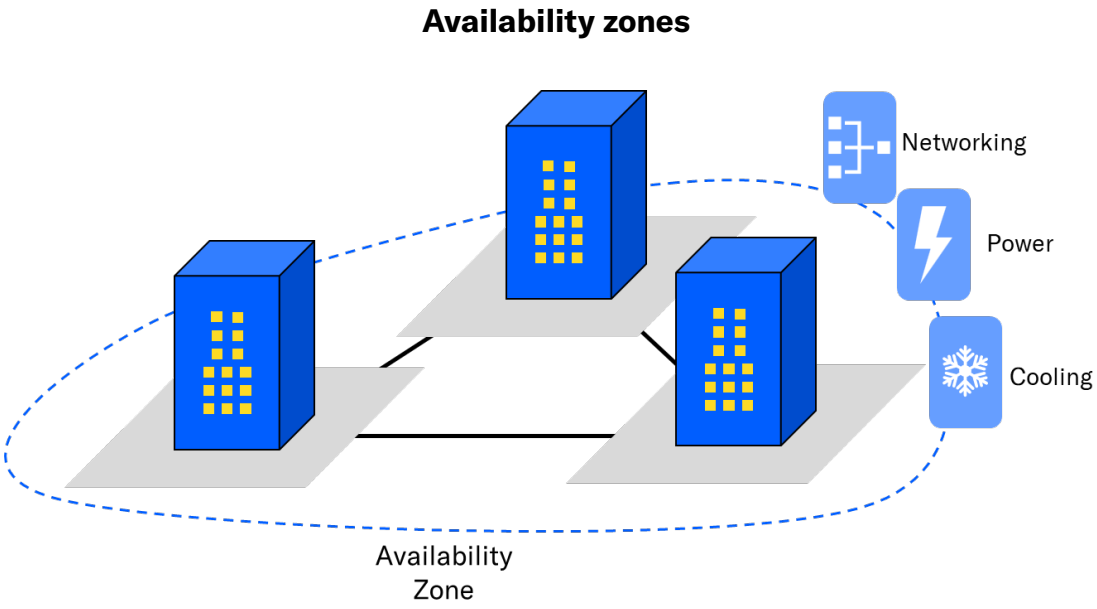

Multiple data centers can be grouped into availability zones, represented by the dashed perimeter line. This grouping enhances resilience and allows for greater specialization to support a wider range of tasks. Typically, data centers within the same availability zone share common infrastructure, including networking, power, and cooling systems as indicated by the relevant symbols.
Source: Microsoft Azure

At the second level of cloud infrastructure, providers divide the globe into regions. Each region consists of multiple availability zones and is geographically dispersed, often spanning different metropolitan areas or even countries. This separation adds another layer of resilience. In the event of a regional outage, some workloads and data can be shifted to other regions to support business continuity. However, the ability to transfer data between jurisdictions is often constrained by regulatory frameworks such as the General Data Protection Regulation (GDPR), which place strict limits on cross-border data movement, and by latency requirements between users and the cloud infrastructure. These constraints must be considered when evaluating the resilience and exposure of cloud-based systems.

Regions

The gray lines on the globe divide the world map into regions, each of which can host one or more availability zones. These regions are interconnected, enabling communication across geographic boundaries.
Source: Microsoft Azure

Beyond regional diversification, some organizations adopt a multi-cloud strategy, meaning they use multiple cloud providers, to further reduce outage risk. This approach limits reliance on any single provider and adds resilience by leveraging distinct infrastructures and operational models. However, effective, multi-cloud deployment also poses notable challenges. Differences in platforms, security protocols, and data handling require careful coordination, and regulatory constraints like GDPR can restrict cross-provider data movement. Multi-cloud strategies can also be expensive and complex to integrate. As a result, only large enterprises with critical availability needs and sufficient resources tend to pursue this path.

CLOUD COMPUTING LAYERS

To deepen understanding of how cloud outages occur and how risk can be assessed, this section introduces the most technical part of the paper. It presents a conceptual model that divides cloud computing infrastructure into four distinct layers. Although simplified, this framework offers insurance professionals a practical lens for evaluating failure points and resilience strategies. At the base is the physical layer, which includes the hardware — servers, storage arrays, and networking equipment — that powers cloud services. In this layer, availability zones are interconnected by optical cables, allowing high-speed, low-latency communication between data centers. Above this is the control layer, which manages and aggregates physical resources. Key components include the hypervisor, which partitions compute and memory across virtual machines, and controllers, which distribute workloads to maintain performance and availability. The middleware layer consists of virtual machines and containerized environments that expose tools, APIs, and runtime services. These support the applications layer, where end-user software resides and interacts with the cloud. This layered model serves as a useful proxy for

understanding how failures can propagate through the system and where mitigation strategies, such as redundancy and load redistribution, are most effective.

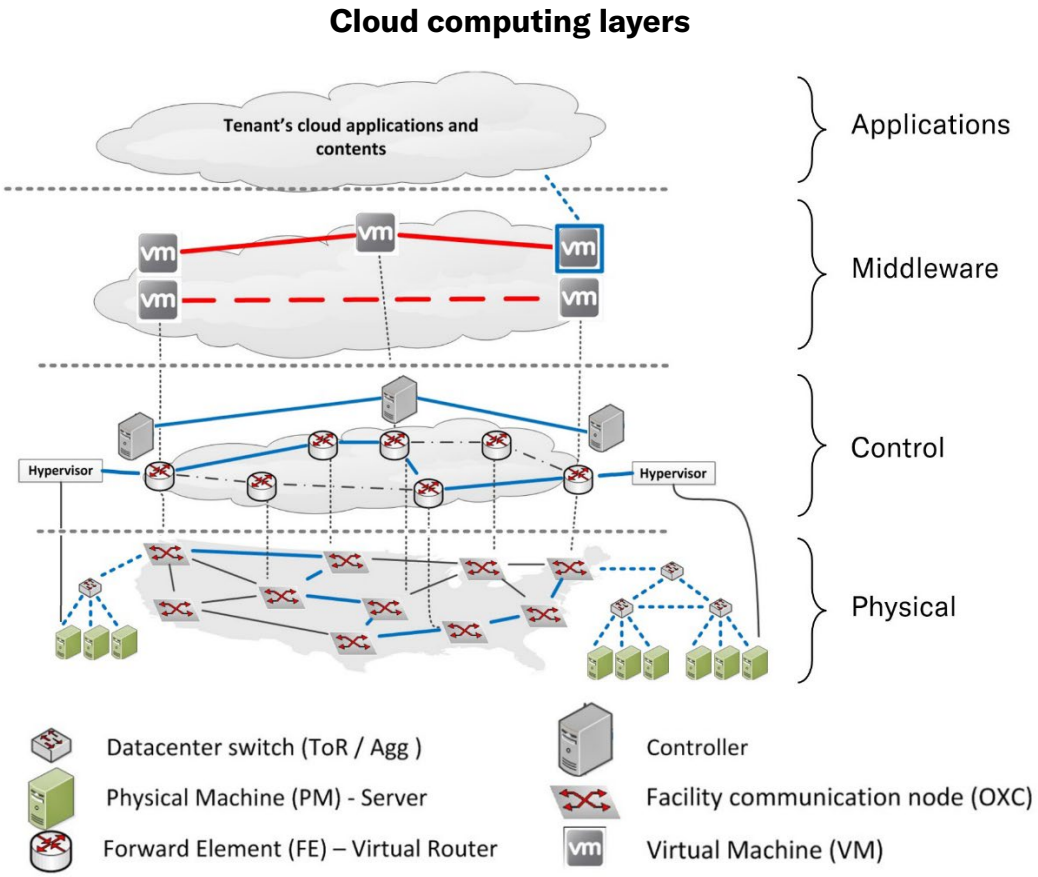

Overlaid on the US map, this figure visually represents the conceptual division of the cloud into four distinct layers: physical, control, middleware, and applications.
Source: "A survey on resiliency techniques in cloud computing infrastructures and applications," *IEEE Communications Surveys & Tutorials*

SERVICES

A cloud service is a software-based capability offered by a cloud service provider to help users perform specific tasks or achieve operational goals within a cloud environment. These services summarize the complexity of the underlying infrastructure, delivering functionality through APIs, user interfaces, or automation tools. Depending on their design and purpose, cloud services may span multiple layers of the cloud computing infrastructure (CCI). For example, some interact directly with virtual machines while others manage access or data storage. Services are commonly grouped into three categories: identity and access management, storage, and compute. Understanding these service types is essential since they are frequently referenced in cloud outage reports and often indicate which parts of the infrastructure were affected.

UNDERSTANDING OUTAGE REPORTS

Despite their engineered resilience, cloud-based systems remain vulnerable to service interruptions, known as cloud outages, which can result in significant business interruption. For organizations that rely on continuous access to cloud services, even short periods of downtime can disrupt operations, impact revenue, and trigger cascading effects across supply chains and customer engagement. Understanding how these outages unfold is essential for insurance professionals assessing exposure and designing coverage for digital infrastructure risks.

This section presents the conceptual centerpiece of the paper: a new diagram that builds on the four-layer CCI model introduced earlier. While the CCI framework explains how cloud services are constructed and organized, the new diagram focuses on how outages develop and propagate. Each node in the diagram, depicted as a circle, represents a conceptual hub (a stage in the life cycle of a cloud outage) from its initial formation to its impact on end-user applications. Gray bars indicate whether two nodes are connected, signifying the possible presence of a direct causal relationship.

How outages develop and propagate

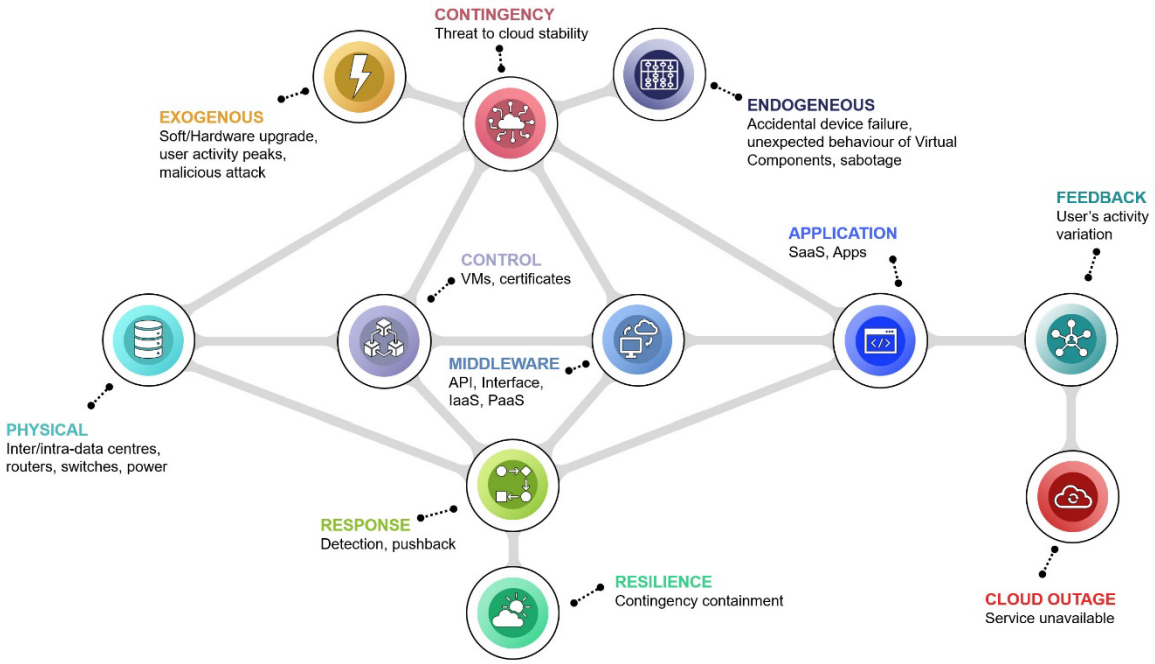

Beginning at the top of the diagram, "contingencies" whether "exogenous" (external events like natural disasters or cyberattacks) or "endogenous" (internal failures such as misconfigurations or hardware faults), can affect one or more layers of the cloud architecture. The "response" node represents the reactive mechanisms (whether automatic or with human intervention) that are activated to address a given contingency. Typically, the containment of these events is successful as indicated by the "resilience" node. However, the contingency may not be immediately resolved and may propagate across layers until end users experience service degradation or disruption, as indicated by the "feedback" node. These events can potentially escalate into full-scale service disruptions, thus determining a "cloud outage."

In such cases, cloud service providers typically issue structured incident reports. These reports detail the duration of the outage, calculated from the start time to the moment the last client no longer experiences disruption, as well as the affected regions and impacted services. This

standardized reporting format helps stakeholders, including insurers, assess the scope and severity of the event and understand which parts of the infrastructure were compromised.

Overall, this layered progression offers a structured lens for interpreting outage reports and evaluating systemic risk, helping insurers better understand the nature and scope of cloud-related business interruptions.

HISTORICAL OUTAGE EXAMPLES

To demonstrate how the conceptual framework can be applied in practice, we include and analyze two cloud outage events. First, we study excerpts from the AWS outage report titled "Summary of the Amazon Kinesis Event in the Northern Virginia (US-East-1) Region," published on November 25, 2020, followed by the AWS outage report titled "Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-East-1) Region," published on November 1, 2025. These reports serve as real-world examples of how a cloud outage progresses through the infrastructure layers previously described. Each excerpt will be mapped to the relevant conceptual hub in the diagram, illustrating how disruptions originate, propagate, and ultimately affect end-user services. Because outage reports are typically written retrospectively, they allow for reverse analysis, starting from the observed impact and tracing back to the root cause. This approach helps insurance professionals understand the sequence of events and assess the nature and extent of business interruption resulting from cloud failures.

Cloud outage on November 25, 2020

- Exogenous contingency affecting control layer: "The trigger, though not root cause, for the event was a relatively small addition of capacity that began to be added to the service at 2:44 AM PST, finishing at 3:47 AM PST."
    - Since hardware failure was not mentioned, it is likely that the additional capacity only affected the control layer.
- Response and start time: "At 5:15 AM PST, the first alarms began firing for errors on putting and getting Kinesis records. Teams engaged and began reviewing logs."
    - In contingency situations, teams of experts are typically engaged to analyze and address the issue, often by reviewing error logs to identify the root cause. Kinesis is a foundational AWS offering, conceptually sitting in the middleware layer, used for real-time data streaming and analytics.
- Middleware and application layer: "There were a number of services that use Kinesis that were impacted as well."
    - The report highlights multiple affected services that rely directly on Amazon Kinesis, demonstrating the cascading effects of the outage.
- Feedback and cloud outage: "As a result, Cognito customers experienced elevated API failures and increased latencies for Cognito User Pools and Identity Pools, which prevented external users from authenticating or obtaining temporary AWS credentials."
    - Among other issues, users were affected at the identity and access management level.
- End time: "By 4:15 PM PST, the majority of these issues had been resolved."
    - User access was restored gradually as services recovered and systems stabilized.

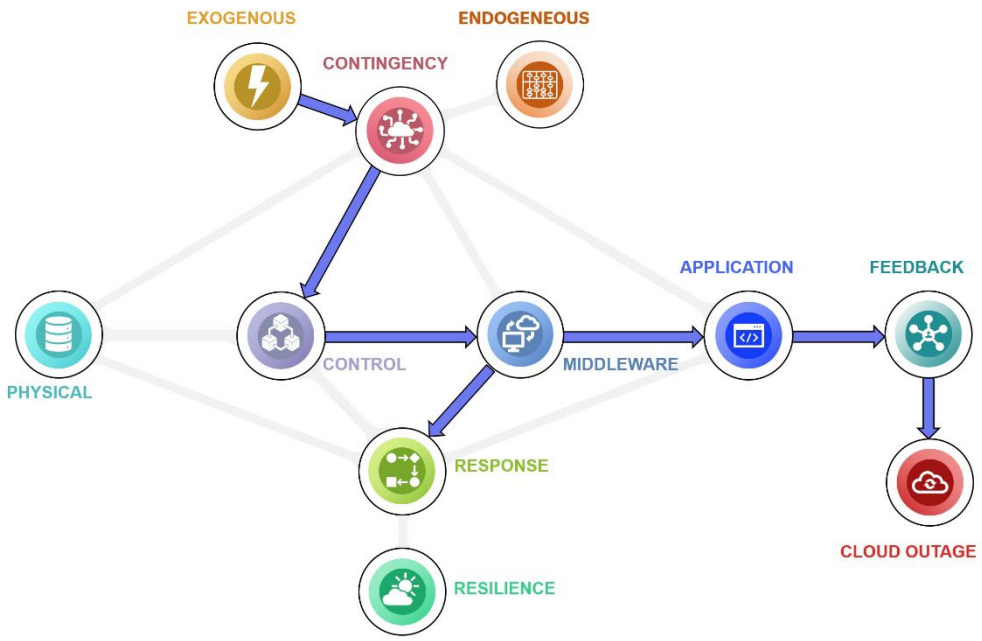

Cloud outage on October 19, 2025

- Endogenous contingency affecting control layer: "The incident was triggered by a latent defect within the service's automated DNS management system that caused endpoint resolution failures for DynamoDB."
    - DynamoDB is an AWS service that provides a highly scalable and low-latency database solution, delivering high-speed performance regardless of the number of users. As a core component of the storage category, many other services (potentially residing in different layers) depend on its availability to perform data-related tasks.
- Response and start time: "When this issue occurred at 11:48 PM PDT, all systems needing to connect to the DynamoDB service in the N. Virginia (US-East-1) region via the public endpoint immediately began experiencing DNS failures and failed to connect to DynamoDB... Engineering teams for impacted AWS services were immediately engaged and began to investigate."
- Middleware and application layers: "Between 11:48 PM PDT on October 19 and 1:50 PM PDT on October 20, customers experienced increased EC2 API error rates, latencies, and instance launch failures in the N. Virginia (US-East-1) region."
    - "The delays in network state propagations for newly launched EC2 instances also caused impact to the Network Load Balancer (NLB) service and AWS services that use NLB. Between 5:30 AM and 2:09 PM PDT on October 20, some customers experienced increased connection errors on their NLBs in the N. Virginia (US-East-1) region."
    - EC2 is an AWS service that allows users to launch as many or few virtual servers as needed, configure security and networking, and manage storage. Meanwhile, the NLB orchestrates the distribution of traffic across multiple EC2 instances. During the incident, EC2 experienced failures immediately after DynamoDB, and NLB was impacted a few hours later. Both EC2 and NLB conceptually reside within the same layer, resulting in an intra-layer cascade effect. The report also highlights that a total of 141 services were impacted.

- Feedback and cloud outage: "Customers experienced… error rates, latencies, and instance launch failures."
    - "Between October 19 at 11:51 PM PDT and October 20 at 1:25 AM PDT, AWS customers attempting to sign into the AWS Management Console using an IdAM user experienced increased authentication failures due to underlying DynamoDB issues in the N. Virginia (US-East-1) region."
    - The outage manifested as intermittent and slow service availability, along with an inability to sign in.
- End time: "By 2:25 AM, all DNS information was restored, and all global tables replicas were fully caught up by 2:32 AM. Customers were able to resolve the DynamoDB endpoint and establish successful connections as cached DNS records expired between 2:25 AM and 2:40 AM."
    - "Recovery started at 12:01 PM PDT with full EC2 recovery occurring at 1:50 PM PDT. During this period, new instance launches failed with either a 'request limit exceeded' or 'insufficient capacity' error."
    - "Between 8:31 AM and 9:59 AM, STS API error rates and latency increased again as a result of NLB health check failures. By 9:59 AM, we recovered from the NLB health check failures, and the service began normal operations."
    - According to AWS's own outage event summary, the outage on the CSP side ended at 2:20 PM PT on October 20, 2025.
    - Although the underlying issue was resolved, it took additional time for the service to become reliably accessible to other services in higher layers. This delay not only defines the outage's end time, which occurred significantly later than the initial resolution, but also highlights the gradual nature of recovery.

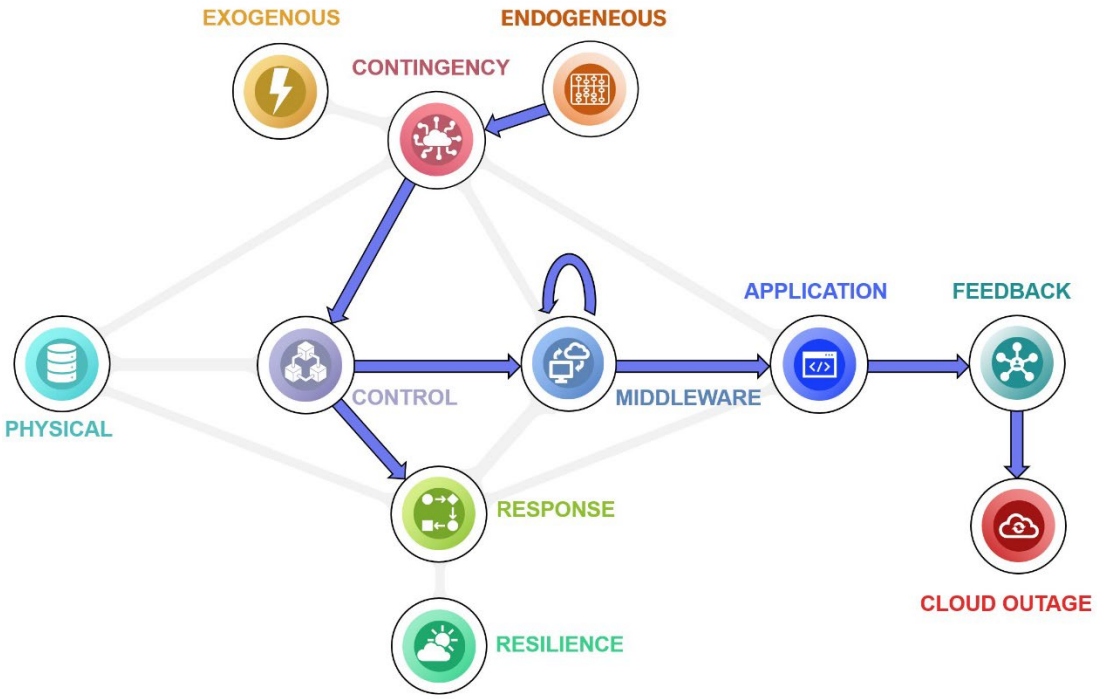

Summarizing cloud outages

The information from the reports above can be distilled in the following tabular framework. This approach allows for the derivation of statistics for those interested in modeling cloud outages.

Start time	End time	Affected regions	Affected service layers	Affected service
November 25, 2020 5:15 AM PT	November 25, 2020 4:15 PM PT	Northern Virginia (US-East-1)	Control and middleware	Kinesis, Amazon Cognito, CloudWatch, Auto Scaling, Lambda, CloudWatch Events, EventBridge, Elastic Container Service, Elastic Kubernetes Service
October 19, 2025 11:48 PM PT	October 20, 2025 2:20 PM PT	Northern Virginia (US-East-1)	Control and middleware	DynamoDB, EC2, Network Load Balancer, Lambda, Elastic Container Service, Kubernetes Service, Fargate, Connect, Security Token Service, Management Console, Redshift, Support Console, Managed Workflows for Apache Airflow, Outposts life cycle operations

CONCLUSION

Cloud infrastructure is inherently complex, built from interdependent layers of hardware, software, and services designed to deliver scalable and resilient computing. However, as this paper demonstrates, there are multiple ways in which cloud systems can fail, each with distinct technical characteristics and business implications. The impact of a cloud outage is not uniform across organizations. For instance, firms experience disruptions differently depending on how their infrastructure is architected, which cloud services they rely on, and the extent to which their revenues are directly dependent on cloud performance. While many providers offer robust resilience features, including regional failover and multi-cloud strategies for the most advanced users, uncertainty remains. Even with these safeguards, cloud outages retain the potential to cause material business interruption. Our aim has been to offer a concise introduction to cloud infrastructure and a conceptual framework for understanding how outages emerge and propagate, helping to interpret outage reports and assess exposure more effectively.

GLOSSARY

Application layer: Topmost layer in a cloud computing or network architecture where end-user applications and software reside

Application programming interface (API): A set of tools and protocols that allow different software applications to communicate with each other and are often used to interact with cloud services

Availability zone: A grouping of one or more data centers within a cloud region, designed to operate independently with separate power, cooling, and networking to ensure high availability

Business interruption (BI): A type of insurance coverage that compensates for lost income and operational costs due to disruptions, such as cloud outages

Certificates: Digital documents used to establish trust between entities by verifying the identity of a server, user, or device

Cloud computing: The delivery of computing services (such as servers, storage, databases, networking, and software) over the internet, enabling on-demand access to resources

Cloud outage: A disruption in cloud services that can result in downtime, latency, or unavailability of resources for end users

Cloud service provider (CSP): A company that offers cloud computing services, such as AWS, Microsoft Azure, or Google Cloud Platform

Cognito: An identity platform for web and mobile applications that provides user signup, sign-in, and access control

Compute: A category of cloud services that provides processing power for running applications and workloads

Containerization: A form of virtualization where applications run in isolated environments called containers (such as Docker)

Contingency: An exogenous (arising from external factors) or endogenous (originating from within the organization) event that may disrupt normal cloud operations

Control layer: The layer in cloud infrastructure responsible for managing and aggregating physical resources, including hypervisors and workload controllers

Data center: A physical facility housing servers, storage systems, and networking equipment that powers cloud services

Disaster recovery: Strategies to restore IT systems after a disruptive event (for example, replicating data and applications across multiple regions or leveraging backup and failover solutions to ensure business continuity and minimize downtime)

Domain Name System (DNS): A system that translates human-readable domain names (like www.example.com) into IP addresses that computers use to identify each other on the network

DynamoDB: NoSQL database service provided by AWS

Elastic Compute Cloud (EC2): An AWS service that provides scalable virtual servers for running applications

Endogenous contingency: Internal failures within a cloud system, such as misconfigurations or hardware faults

Exogenous contingency: External events that disrupt cloud systems, such as natural disasters or cyberattacks

Failover: A backup operational mode in which system functions are automatically transferred to a standby system in the event of a failure

General Data Protection Regulation (GDPR): A European Union regulation that governs data protection and privacy, often restricting cross-border data movement

Hypervisor: Software that creates and manages virtual machines by partitioning physical hardware resources

Identity and access management (IdAM): A framework of policies and technologies for ensuring that the right individuals have access to the right resources in a cloud environment

Infrastructure as a service (IaaS): A cloud computing model in which the provider manages the core infrastructure (such as servers, storage, and networking) and the customer handles operating systems, applications, and data

Kinesis: An AWS service for real-time data streaming and analytics

Latency: The delay between a user's action and the response from a cloud service, often measured in milliseconds

Load balancer: A service that distributes incoming network traffic across multiple servers to ensure reliability and performance

Middleware layer: The layer in cloud infrastructure that includes virtual machines, containerized environments, and runtime services, all of which support the applications layer

Multi-cloud strategy: The use of multiple cloud service providers to reduce reliance on a single provider and enhance resilience

Multi-region strategy: The use of data centers across multiple regions and potentially CSPs to reduce reliance on a single geographical location and enhance resilience

Network load balancer (NLB): A type of load balancer that distributes traffic across multiple servers at the network layer, often used for high-performance applications

Physical layer: The foundational layer of cloud infrastructure, consisting of hardware such as servers, storage arrays, and networking equipment

Platform as a service (PaaS): A cloud computing model in which the provider manages the infrastructure and runtime environment, allowing customers to focus on application development

Redundancy: The duplication of critical components or functions in a system to increase reliability and resilience

Region: A geographically distinct area containing multiple availability zones, designed to provide resilience against outages

Runtime services: Cloud-based tools that provide the necessary resources and support for running applications or workloads (handling execution, scaling, monitoring, managing dependencies, and more), allowing developers to focus on application logic without worrying about the underlying infrastructure

Security Token Service (STS): An AWS service that provides temporary security credentials for accessing AWS resources

Server: A device that provides resources (storage, computational, connectivity) to other machines over a network

Service degradation: A temporary reduction in the performance or quality of a cloud service (such as slower response times, limited functionality, and reduced capacity) that falls short of a complete loss of service but may still impact user experience or business operations

Service disruption: A temporary loss of functionality in a cloud service that may or may not escalate into a full-scale outage

Software as a service (SaaS): A cloud computing model in which the provider delivers fully managed software applications to customers over the internet

Storage: A category of cloud services that provides scalable and secure data storage solutions

Technical debt: The implied cost of additional work caused by choosing an easier or faster solution with lower standards (like maintainability, feature completeness, scalability, and more) instead of a more robust approach

Traffic orchestration: The process of managing and directing network traffic to ensure optimal performance and reliability

Virtualization: The process of creating virtual versions of physical computing resources (such as servers, storage, or networks) to allow multiple virtual environments to run on a single physical system

Virtual machine: A software-based emulation of a physical computer that allows multiple operating systems to run on a single physical machine

Zero downtime: A design goal in cloud computing to ensure continuous availability of services without any interruptions.

REFERENCES

AWS, "Summary of the Amazon DynamoDB Service Disruption in the Northern Virginia (US-East-1) Region," November 1, 2025

AWS, "Summary of the Amazon Kinesis Event in the Northern Virginia (US-East-1) Region," November 25, 2020

AWS Health Dashboard, "Service health," October 20, 2025

Colman-Meixner, C.; Develder, C.; Tornatore, M.; and Mukherjee, B.; "A survey on resiliency techniques in cloud computing infrastructures and applications," *IEEE Communications Surveys & Tutorials*, 18(3), pp.2244-2281, 2016

Microsoft Azure, "What are availability zones?"

ACKNOWLEDGEMENTS

This white paper was developed by AXIS, Gallagher Re, Moody's, and Munich Re as part of the Cyber Industry Steering Group (ISG) initiative. The authors wish to thank the broader ISG membership for their thoughtful review, constructive feedback, and insights shared during steering group meetings, which have materially strengthened this work. Although this paper reflects the perspectives of the named authors, it has been informed by the collective expertise of the ISG community.

What is the Cyber Industry Steering Group?

The Cyber Industry Steering Group is an industry initiative launched by Moody's, alongside cyber market participants, to help facilitate the growth of the global cyber insurance market.

It brings together leading members of the insurance community — including (re)insurers, brokers, technology partners, and other stakeholders — to advance understanding of the cyber threat landscape, build resilience, improve analytics and modeling, and attract new sources of capital to meet increasing customer needs.

This group believes that by advancing the industry's understanding and quantification of cyber risk, this will help to attract capacity and support the market's growth.

This group is conducting important primary research, developing improved cyber risk assessment tools, and leading a series of market education initiatives.

Find out more about the Cyber Industry Steering Group here.

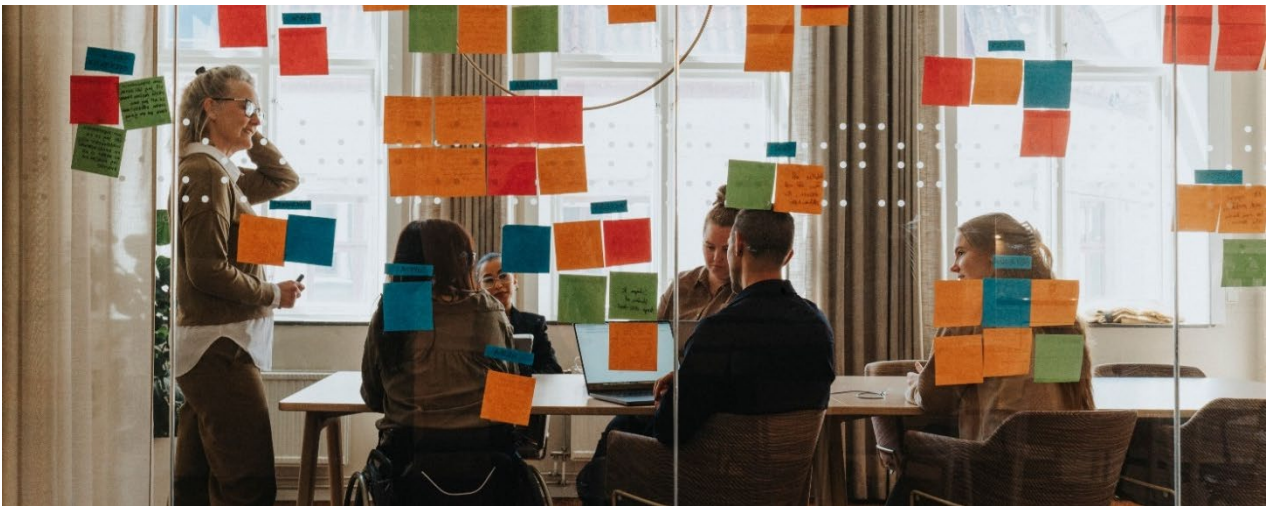

The ISG is a cyber industry steering group comprised of stakeholders collaborating with the collective goal of supporting the sustainable growth of the cyber insurance market. The ISG members did not establish any legal partnership, joint venture, or similar for the purposes of the ISG or producing this paper and no ISG member is constituted the agent of another or otherwise authorised to act on another ISG member's behalf. Without limiting the foregoing, no ISG member shall have any responsibility for any act or omission of any other ISG member.

© 2026 Moody's Analytics, Inc. and its affiliates; Gallagher Re Inc and its affiliates; AXIS Specialty U.S. Services, Inc. and its affiliates; Münchener Rückversicherungs-Gesellschaft Aktiengesellschaft and its affiliates; all rights reserved.

AUTHORS

For the purposes of this white paper, any and all reference to the authors shall mean:

AXIS
AXIS Capital Holdings Limited is a publicly listed global specialty insurer and reinsurer organized under the laws of Bermuda, with its principal executive offices at **29 Richmond Road, 3rd Floor, Pembroke HM 08, Bermuda**. **AXIS Capital Holdings Limited** is subject to supervision by the Bermuda Monetary Authority, which acts as the Group Supervisor for the AXIS group of companies. AXIS conducts insurance and reinsurance business through various licensed subsidiaries in multiple jurisdictions and, in certain countries, including the United States, AXIS insurance and reinsurance entities may operate subject to applicable regulatory restrictions.

Dr. Marco Rohrer
Marco.Rohrer@axiscapital.com

Moody's
Moody's Analytics is a provider of data, analytics, research, software, and advisory solutions for risk management and decision making and operates as a business unit of **Moody's Corporation**, a company organized under the laws of the **United States**. Moody's Analytics operates through licensed and authorized entities in multiple jurisdictions worldwide and provides analytics and advisory services to financial institutions, insurers, corporates, and public sector market participants. Moody's Analytics maintains its principal offices in the United States at **7 World Trade Center, 250 Greenwich Street, New York, NY 10007**.

Dr. Zoltan Tuza
Zoltan.Tuza@moodys.com
Dr. Andrea Pizzoferrato
Andrea.Pizzoferrato@moodys.com

Gallagher Re
Gallagher Re is the global reinsurance broking and advisory business of Arthur J. Gallagher & Co., a company organized under the laws of the **United States**. Gallagher Re operates through licensed and authorized brokerage entities in multiple jurisdictions worldwide and provides reinsurance placement, advisory, and analytics services to insurance and reinsurance market participants. **Gallagher Re** maintains its principal offices in the **United Kingdom at The Walbrook Building, 25 Walbrook, London EC4N 8AW.**

Ed Pocock
ed_pocock@gallagherre.com

Munich Re
Münchener Rückversicherungs-Gesellschaft Aktiengesellschaft in München ("**Munich Re**," Commercial Register Munich Number: HRB 42039) is supervised by the German Bundesanstalt für Finanzdienstleistungsaufsicht (BaFin). **Munich Re** is a reinsurance company organized under the laws of **Germany** with its Registered Office at Koeniginstrasse 107, 80802 Munich. In some countries, including in the United States, **Munich Re** holds the status of an unauthorized reinsurer.

Dr. Matthias Killiches
mkilliches@munichre.com
Dr. Martin Sigl
msigl@munichre.com

ALL INFORMATION CONTAINED HEREIN IS PROTECTED BY LAW, INCLUDING BUT NOT LIMITED TO, COPYRIGHT LAW, AND NONE OF SUCH INFORMATION MAY BE COPIED OR OTHERWISE REPRODUCED, REPACKAGED, FURTHER TRANSMITTED, TRANSFERRED, DISSEMINATED, REDISTRIBUTED OR RESOLD, OR STORED FOR SUBSEQUENT USE FOR ANY SUCH PURPOSE, IN WHOLE OR IN PART, IN ANY FORM OR MANNER OR BY ANY MEANS WHATSOEVER, BY ANY PERSON WITHOUT THE AUTHORS' PRIOR WRITTEN CONSENT. THE INFORMATION HEREIN IS PROVIDED FOR INFORMATIONAL PURPOSES ONLY AND DOES NOT CONSTITUTE AND SHOULD NOT BE CONSTRUED AS PROFESSIONAL ADVICE. FOR CLARITY, NO INFORMATION CONTAINED HEREIN MAY BE USED TO DEVELOP, IMPROVE, TRAIN OR RETRAIN ANY SOFTWARE PROGRAM OR DATABASE, INCLUDING, BUT NOT LIMITED TO, FOR ANY ARTIFICIAL INTELLIGENCE, MACHINE LEARNING OR NATURAL LANGUAGE PROCESSING SOFTWARE, ALGORITHM, METHODOLOGY AND/OR MODEL.

All information contained herein is obtained by the Authors from sources believed by it to be accurate and reliable. Because of the possibility of human or mechanical error as well as other factors, however, all information contained herein is provided "AS IS" without warranty of any kind. To the extent permitted by law, the Authors and their directors, officers, employees, agents, representatives, licensors and suppliers disclaim liability to any person or entity for any indirect, special, consequential, or incidental losses or damages whatsoever arising from or in connection with the information contained herein or the use of or inability to use any such information, even if the Authors or any of their directors, officers, employees, agents, representatives, licensors or suppliers is advised in advance of the possibility of such losses or damages, including but not limited to any loss of present or prospective profits. To the extent permitted by law, the Authors and their directors, officers, employees, agents, representatives, licensors and suppliers disclaim liability for any direct or compensatory losses or damages caused to any person or entity, including but not limited to by any negligence (but excluding fraud, willful misconduct or any other type of liability that, for the avoidance of doubt, by law cannot be excluded) on the part of, or any contingency within or beyond the control of, the Authors or any of their directors, officers, employees, agents, representatives, licensors or suppliers, arising from or in connection with the information contained herein or the use of or inability to use any such information. NO WARRANTY, EXPRESS OR IMPLIED, AS TO THE ACCURACY, TIMELINESS, COMPLETENESS, MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OF ANY ASSESSMENT, OTHER OPINION OR INFORMATION IS GIVEN OR MADE BY THE AUTHORS IN ANY FORM OR MANNER WHATSOEVER.