KCC WHITE PAPER

Hyperscale Data Centers
The Opportunity for (Re)Insurers

September 2026

KC &CO

The Innovation and Technology Leader in Weather, Climate, and Catastrophe Risk Modeling

About KCC

Karen Clark & Company (KCC) provides (re)insurers and ILS investors with accurate models and advanced technology for deep insight climate, weather, and catastrophe risk. KCC catastrophe models cover tropical cyclones, extratropical cyclones, severe convective storms, winter storms, wildfires, floods, and earthquakes in over 70 countries. KCC's unique LiveEvents™ application provides real-time information for major weather events as they're unfolding. KCC professionals are globally recognized experts in catastrophe and weather modeling who work with company executives to enhance business strategies, competitive advantage, and financial results.

For more information, please visit www.karenclarkandco.com.

Copyright

Document Date

September 2026

Contact Information

If you have any questions regarding this document, please contact:

Karen Clark & Company
116 Huntington Avenue
Boston, MA 02116
T: 617.423.2800
F: 617.423.2808
Info@karenclarkandco.com

Introduction

As the global economy becomes increasingly dependent on cloud computing and artificial intelligence (AI), data centers are emerging as some of the most critical—and fastest-growing—infrastructure assets.

From isolated server rooms and mainframes, data centers have evolved into sprawling industrial campuses that support cloud computing, data storage, and AI workloads. With each new wave of technological advancement, the scale, sophistication, and strategic value of these facilities accelerate.

This rapid evolution is reshaping how data centers are designed, financed, protected, and insured. As these facilities expand in both number and complexity, they present new opportunities for innovation in risk assessment and insurance product development.

With cloud adoption and AI demand poised to grow for years to come, insurers will require new approaches for assessing and managing these evolving exposures.

This white paper explores the evolution of data centers—and in particular hyperscale data centers—their structural characteristics and exposure to natural hazards, and how the insurance industry can transform these emerging risks into strategic opportunities.

Data Center Evolution

Data centers have been around in some capacity since the 1940s, when the Electronic Numerical Integrator and Computer (ENIAC) was built at the University of Pennsylvania. ENIAC weighed 30 tons and filled a 30x50 foot room. It could process five thousand operations per second (less than the computing power a smartphone uses to capture and save a single selfie today!).

As technology evolved over the next few decades, mainframe computers took over. Mainframes processed data, performed complex calculations, and could run multiple programs at a time, but they were not yet connected to any network.

In the late 1960s, an information transfer network was established in Northern Virginia for communication between the Pentagon and research institutions. This network established the foundation for future internet connectivity and positioned the region as a hub for data transfer—setting the stage for Northern Virginia to become Data Center Alley.

1990s

The Birth of Data Center Alley

One of the first internet exchange points was established in Northern Virginia in 1992, which increased speed and connectivity. This area had low-cost land and electricity, which began to attract internet companies. AOL established a major presence in Ashburn, VA and invested in fiber and power infrastructure, attracting other companies such as Verizon and Telos. The world's largest colocation data center—Equinix—was built in Ashburn in 1999, solidifying the region as a hub for data centers.

2000s

Cloud Computing and the First Hyperscale Data Centers

In the early 2000s, companies shifted from hosting data on-premises to the cloud. This significantly increased the demand for large scale data centers and led to major tech companies opening the first "hyperscale" data centers. Google and Microsoft built their centers in the Pacific Northwest, where there was cheap electricity, tax incentives, and a cool climate. Meanwhile, Amazon built their first hyperscale campus in Data Center Alley, where the majority of the world's internet traffic was passing through.

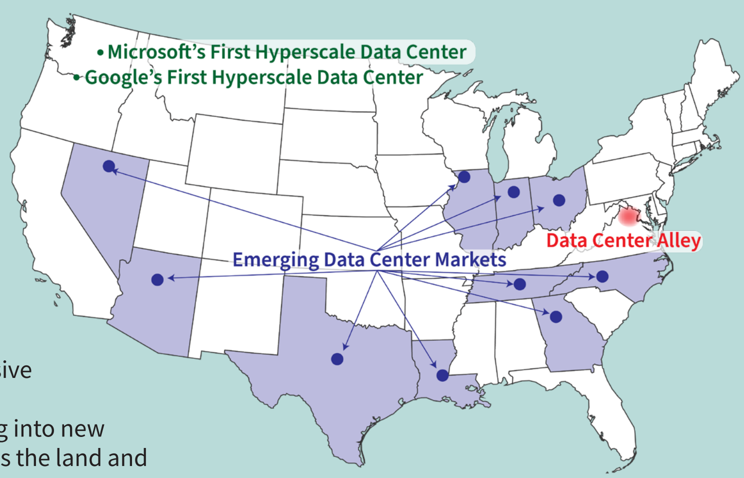

2020s

AI Era

The AI boom has accelerated demand for computational resources, leading to a surge in data center construction—particularly hyperscale data centers that can handle the massive amounts of data required to power AI. Developers are increasingly expanding into new regions across the US interior to access the land and power infrastructure required for these campuses.

Typical Design of Hyperscale Data Centers

A hyperscale data center is a complex industrial campus. The typical layout includes large rectangular server buildings (also known as data halls) supported by on-site substations, cooling infrastructure, backup generators, and even on-site power plants in some cases to ensure continuous operation.

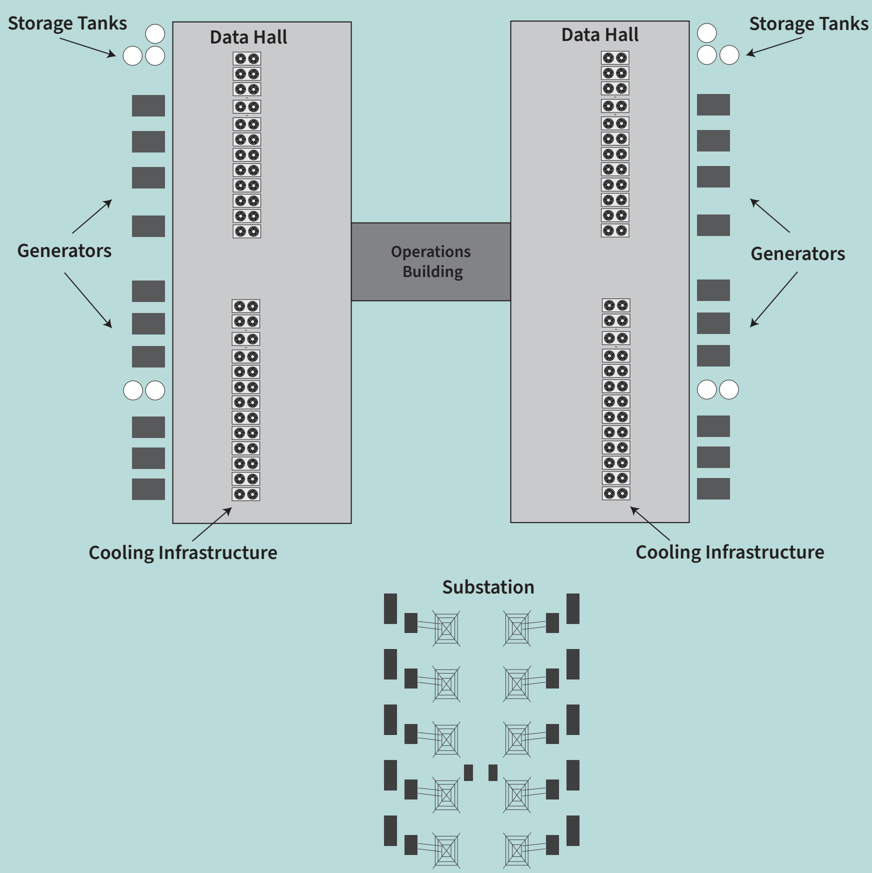

Data Halls


Photo: Chad Davis

A hyperscale campus typically contains several identical low-rise data halls arranged across the site in a standardized layout.


Photo: US DOE

Data halls house the computing hardware, including servers, storage systems, networking equipment, and server racks.

Power Infrastructure

Photo: Wikimedia

On-site substations control voltage from the regional power grid for campus distribution.

Photo: Ryan Sparks

Backup generators can take on the full electrical load almost immediately after a grid failure, maintaining operation until grid power is restored.

Cooling Infrastructure

The main components of the cooling infrastructure include cooling towers, chillers, heat exchangers, and piping networks connected to each data hall.

Storage Tanks


Photo: USDA

Storage tanks at a data center include water or glycol tanks that support cooling operations, as well as fuel tanks that store diesel or natural gas for backup generators.

Data Halls

Each data hall operates as a self-contained unit with its own mechanical, electrical, and fire protection systems. Typically designed with a modular structure, new data halls can be added to the campus quickly as computing demand grows.

Data halls are typically designed and constructed in accordance with local building codes, which establish minimum standards for resilience against region-specific natural hazards such as high winds, floods, earthquakes, heavy snow loads, and wildfire exposure. Depending on the facility's location, these requirements may influence design elements such as structural design, building elevation, roof attachment systems, floodproofing measures, seismic bracing, or fire-resistant construction materials.

The most common structural systems include tilt-up or precast concrete wall panels and steel framing, allowing for large, open interior spaces. Roofs are often constructed with built-up systems or metal panels over insulated decking, designed to handle moderate or high equipment loads and resist water infiltration.

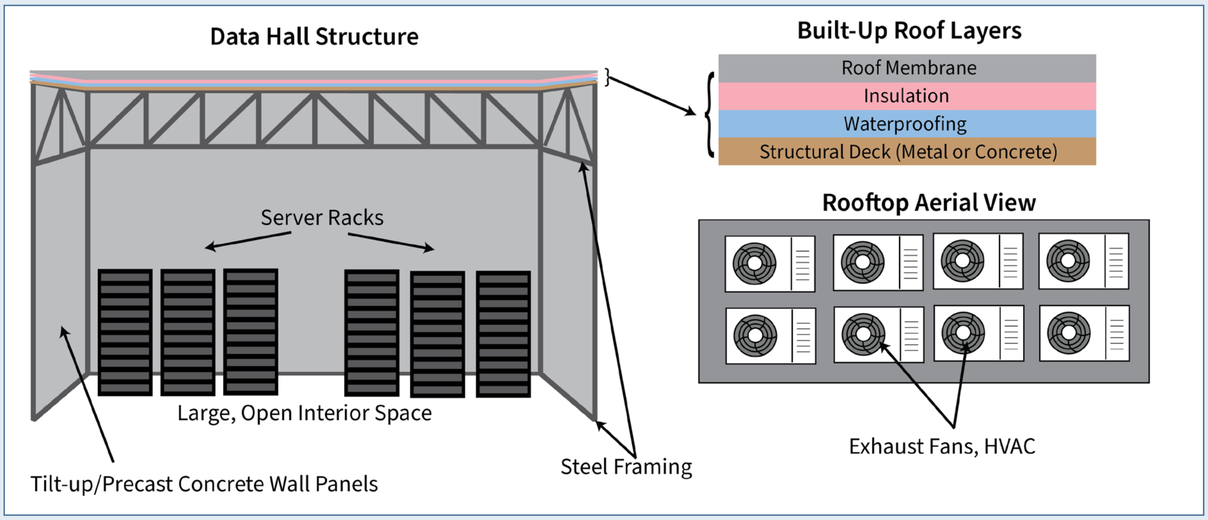

Power Infrastructure

Power infrastructure at a data center typically includes a substation and backup generators. Substations consist of transformers, switchgear, and circuit breakers to ensure the electricity is at the proper voltage and to protect against faults.

The electricity supply for a data center can be fully grid supplied, fully on site, or a combination of the two. In recent years, the share of US hyperscale campuses incorporating on-site power plants has increased significantly, mainly driven by delays in connecting to the grid and the need to maintain a reliable power supply.

On-site power plants most commonly rely on natural gas, although some campuses are beginning to incorporate renewable sources such as solar and wind. These renewable energy plants require on-site battery energy storage systems (BESS) to store and release electricity due to the intermittent nature of solar and wind power.

Cooling Infrastructure

There are two primary types of cooling infrastructure at a data center: air cooling and liquid cooling.

Air Cooling

Air cooling is the most common type of cooling setup. Cool "supply" air is delivered to the server racks, where it absorbs heat from the IT equipment. The warmed "return" air is then collected and routed back to the cooling system, which can be installed either on the roof or adjacent to the data hall, with rooftop equipment being most common.

Some air-cooled data centers use a chilled-water plant to cool the supply air. The chiller transfers the heat to the cooling tower water loop, where the heat is extracted through evaporation and released outside. Once the water is cooled by the cooling tower, it returns to the chiller. The air that is cooled by the chiller is pumped back into the data hall.

In cooler climates, facilities can opt for "free cooling" systems, in which cool air is pulled in from outside and then, after being heated by the equipment, is ejected back outside.

Liquid Cooling

Some facilities utilize liquid cooling, in which the water is circulated around the equipment, rather than cooled air. This is less water intensive because the water is kept in a closed loop rather than evaporating into the environment. Liquid or hybrid systems rely less on rooftop equipment, with more cooling infrastructure located on the ground.

Liquid- or hybrid-cooling approaches are common in hotter and more humid locations, such as parts of Texas, which generally provide fewer favorable hours for air-cooling operation.

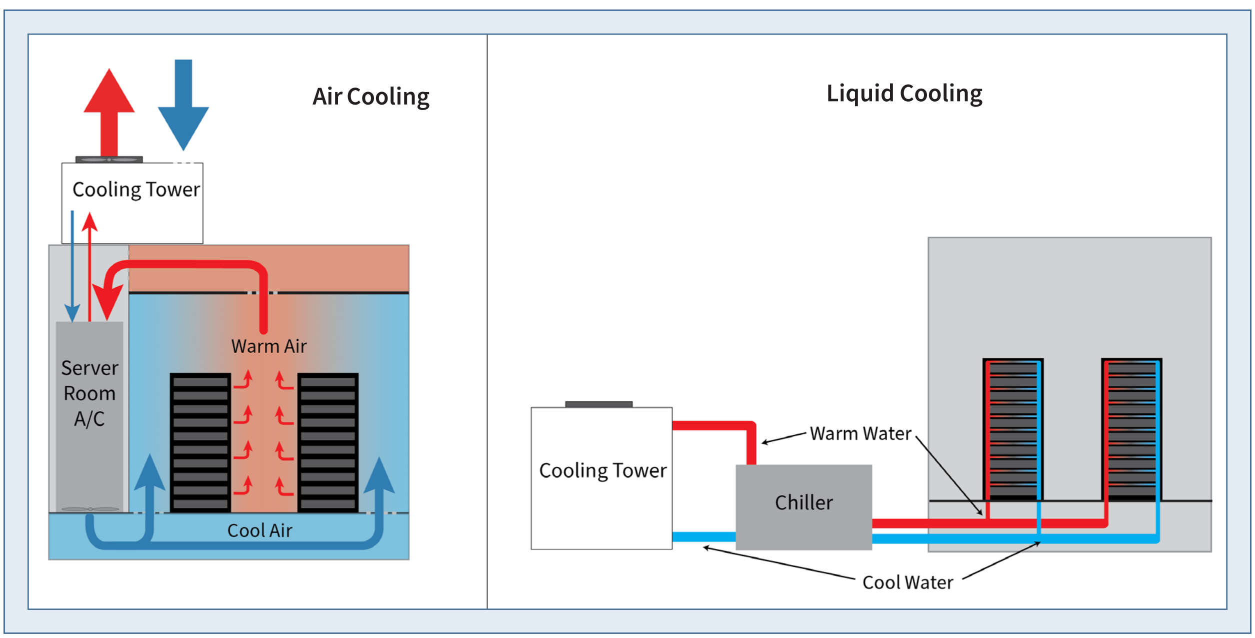

Relative Vulnerability by Structure Type and Peril

The diversity of the various structures creates a complex risk profile in which damage mechanisms and loss drivers are not uniform across the hyperscale campus. Vulnerability depends on site-specific factors, including design standards, mitigation measures, and operational redundancy.

None of these fall neatly into the traditional occupancy codes in the catastrophe models. KCC models have been updated to account for the unique characteristics and peril-specific vulnerabilities of each structure.

Flood

Operations and support buildings are typically the most vulnerable to flooding because they often have more doors, windows, and ground-level access points, increasing the likelihood of water intrusion. Data halls generally have fewer openings but contain high-value equipment. Even shallow flooding can damage servers and lead to costly downtime. Cooling and power infrastructure are highly sensitive to water, and minor flooding can trigger cascading failures and force partial or full shutdowns. Flood risk can be reduced through elevation of critical structures, such as data halls, electrical rooms, and generators, above design flood levels. While critical buildings and equipment are often elevated in flood-prone areas, campus-wide elevation is relatively uncommon. Additional mitigation measures include flood barriers and enhanced site drainage.

What insurers need to know: elevation of critical equipment and structures, flood management systems, and flood barriers.

Earthquake

Cooling infrastructure is among the most vulnerable components during earthquakes, as chillers, piping, and storage tanks can shift, deform, or sustain connection failures due to ground shaking. Data halls also face high risk, due to the concentration of servers, network equipment, and support systems. Strong shaking can cause server racks and electrical equipment to move, overturn, or become disconnected, potentially disrupting operations even when the building itself remains structurally intact.

Common mitigation measures include seismic anchorage, structural bracing, flexible utility connections, and seismic design standards for equipment and support systems. In regions with elevated earthquake risk, additional measures may be used to isolate critical equipment from ground motion and improve operational resilience.

What insurers need to know: seismic mitigations, internal baffles, and tank position relative to ground.

Hail

Most data center buildings have relatively low vulnerability to hail due to robust roof assemblies and limited exterior openings. However, infrastructure located outside buildings—including rooftop mechanical equipment, cooling towers, and exposed piping—can be damaged by hail impacts. While such damage may be localized, it can reduce cooling capacity and disrupt operations in otherwise undamaged data halls. Hail can also damage roof coverings and equipment enclosures, increasing the potential for water intrusion. Protective enclosures and impact-resistant panels can mitigate hail-related damage. Campuses with on-site solar installations face additional exposure, as photovoltaic panels are particularly vulnerable to hail. Impact-rated panels can reduce, but not eliminate, this risk.

What insurers need to know: roof hail impact resistance, hail impact protection, rooftop equipment density, roof leakage protection, and opening type.

Wind

Cooling infrastructure is the most vulnerable element to wind, as high winds can induce uplift, overturning, and lateral loads on cooling towers and rooftop equipment. Enhanced anchorage and structural bracing are the primary mitigation strategies. Data halls and operations buildings generally face moderate wind risk, though extreme wind events can damage roof systems, particularly where rooftop equipment creates weak attachment points. Wind-driven rain can further exacerbate losses. Severe tornado damage remains possible. Although data halls commonly use precast concrete construction, the scale of hyperscale campuses makes it impractical to build facilities with heavily reinforced concrete walls and roofs specifically designed to withstand violent tornadoes. However, the probability of a direct hit from a violent tornado at a specific location is relatively low.

What insurers need to know: tank position relative to ground, rooftop equipment density, roof leakage protection, fan braking systems, and opening type.

Freezing and Snow

Because water-based cooling systems are common, cooling infrastructure is particularly vulnerable to freezing that can lead to burst pipes, which can damage equipment and interior spaces in data halls. Insulation, heat tracing, and protection of exposed piping are essential in cold climates. Snow generally presents low risk because data center roofs are engineered for local snow loads.

What insurers need to know: cooling architecture, freezing protection for thermal power plants, and freezing protection for cooling infrastructure.

Wildfire

Data center campuses typically have low wildfire vulnerability due to defensible space and limited on-site vegetation. However, vegetation management and separation from surrounding wildlands remain important. Smoke and embers can enter buildings through ventilation systems, so facilities in wildfire-prone regions may install enhanced filtration and operation controls to reduce potential impacts.

What insurers need to know: vegetation management, perimeter fire breaks, separation between critical assets, wildfire ember protection, and wildfire smoke protection.

Business Interruption and Redundancy

Physical damage leading to operational downtime is a key concern, because even short outages can have significant financial consequences. An important consideration in the construction of data centers is redundancy. Data centers are designed for continuous operation, and even brief disruptions can result in significant financial losses.

Resilience is achieved through multiple layers of redundancy built into the facility infrastructure. Redundancy design includes additional power and cooling capacity, independent distribution paths, and on-site backup generators that can assume the full electrical load if grid power is lost. Fuel storage allows generators to continue operating during extended outages, while redundant cooling systems help maintain safe operating temperatures if individual components fail. Together, these measures significantly reduce the likelihood that a power, cooling, or utility failure will lead to business interruption.

The industry often references the data center Tier classification system defined by the Uptime Institute, shown below.

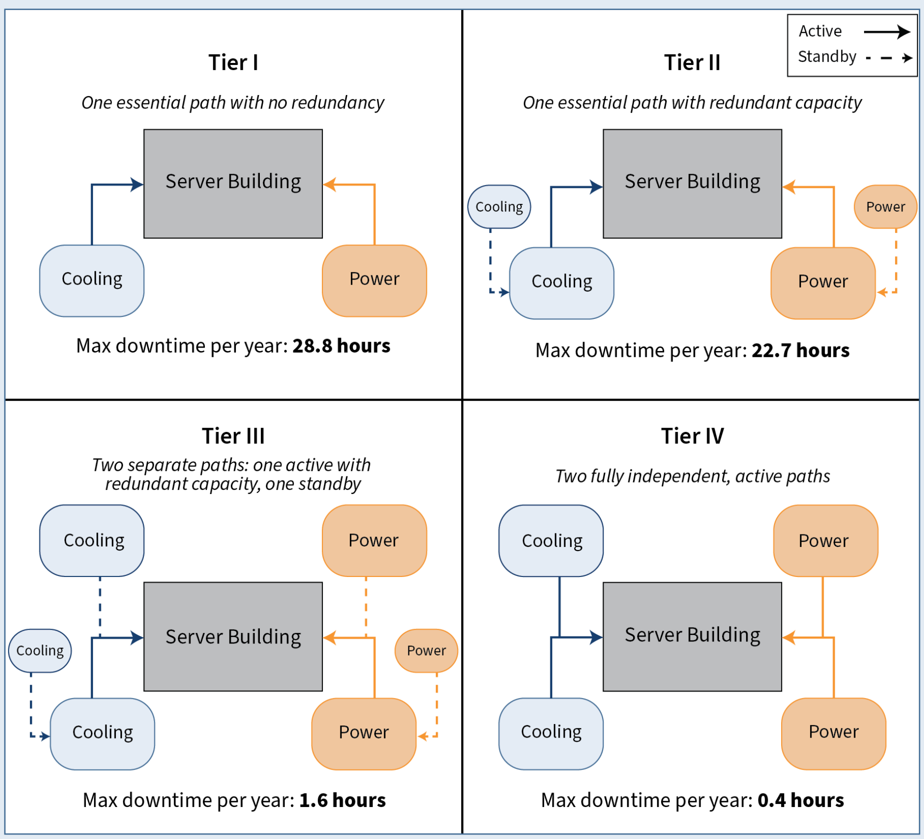

Tier classifications for data center redundancy as defined by the Uptime Institute. The max downtime per year is an industry estimate based on each tier's redundancy design.

Higher-Tier data centers are designed to allow smaller amounts of downtime in a given year, with most hyperscale facilities designed to Tier III standards. A smaller proportion operate at Tier IV—this level of redundancy is expensive, which limits adoption. However, hyperscale data centers often incorporate additional layers of resilience, typically through geographic redundancy and data replication.

Geographic redundancy distributes workloads across physically separated data centers and regions, each with independent power, cooling, and network infrastructure to reduce the risk of correlated failures.

Data replication continuously copies information across these locations, minimizing data loss during disruptions. These operators can quickly shift workloads to unaffected sites or regions during an outage, significantly reducing business interruption.

Insuring Hyperscale Data Centers

Hyperscale data centers differ materially from traditional commercial properties, and insurance for these facilities has been limited to date because of the uncertainty about the risks. Discussions with hyperscale data center operators indicate that some companies are choosing to self-insure because of high premiums and limited coverage.

KCC advanced catastrophe models can provide the information needed to accurately price and manage risk for these complex properties.

KCC engineers have developed new occupancy codes and engineering-based vulnerability functions specifically for hyperscale data centers. These functions capture the characteristics of individual campus structures and their vulnerability to different natural hazards.

As illustrated throughout this paper, the characteristics that most strongly influence vulnerability vary by peril. An example insurer checklist provided on the following pages outlines the key engineering characteristics for modeling data center vulnerability.

In practice, detailed structure-level data may not always be available. To address this limitation, KCC engineers have developed vulnerability functions that can represent the risk of a typical hyperscale campus. Because hyperscale campuses share many common design characteristics, representative campus-level vulnerability functions have been developed.

These functions combine the vulnerabilities of individual structure types into a single campus-level representation and provide a practical alternative when detailed information on individual structures is unavailable.

Example Insurer Data Center Checklist

Feature	Description	Options
Redundancy Level	Level of backup capacity for power, cooling, and other critical systems based on Uptime Institute criteria	☐ Tier I ☐ Tier II ☐ Tier III ☐ Tier IV
On-Site Power Plant	Source of electricity generation located within the campus	☐ None ☐ Thermal ☐ Wind ☐ Solar
Cooling Architecture	Method used to remove heat from IT equipment	☐ Air ☐ Liquid ☐ Hybrid
Elevation of Critical Equipment and Structures	Height of critical assets above design flood levels	☐ No ☐ Yes (with different heights)
Flood Management Systems	Infrastructure that controls, stores, or removes floodwater from the site	☐ No ☐ Pumped Drainage Systems ☐ Floodwater Storage Systems
Flood Barriers	Physical protection designed to prevent floodwaters from reaching critical assets	☐ No ☐ Yes (with different heights)
Seismic Mitigations	Protective measures that reduce damage from earthquake shaking	☐ No ☐ Seismic Bracing ☐ Seismic Anchorage ☐ Rated Raised Floor ☐ Base Isolation

Feature	Description	Options
Internal Baffles	Tank features that limit fluid movement during seismic events	☐ No ☐ Yes
Tank Position Relative to Ground	Placement of storage tanks above, below, or at grade level	☐ Underground ☐ On Ground ☐ Elevated
Roof Hail Impact Resistance	Ability of the roof system to withstand hail impacts	☐ UL 2218 Class 1 ☐ UL 2218 Class 2 ☐ UL 2218 Class 3 ☐ UL 2218 Class 4 ☐ FM Class VSH
Hail Impact Protection	Protective measures that shield exposed cooling equipment from hail damage, such as hail-resistant equipment or hail screens	☐ No ☐ Hail-Resistance Equipment ☐ Hail Screens/Guards
Rooftop Equipment Density	Extent of mechanical and electrical equipment installed on the roof	☐ Low ☐ Medium ☐ High
Roof Leakage Protection	Measures that prevent roof damage from resulting in water intrusion	☐ No ☐ Protected Roof Penetrations ☐ Leak Sensors ☐ Buffer Space Above Critical Spaces
Opening Type	Resistance of doors, windows, and other openings to windborne debris	☐ Impact Rated ☐ Not Impact Rated
Rooftop Equipment Mounting	Method used to secure rooftop equipment against wind loads	☐ No ☐ Structural Anchorage ☐ Engineered Roof Curbs ☐ Ballasted Systems

Feature	Description	Options
Fan Braking Systems	Controls that prevent cooling tower fans from free spinning during high winds	☐ No ☐ Yes
Freezing Protection for Thermal Power Plants	Measures that prevent cold-weather damage to power-generation equipment, such as thermal insulation and freeze inhibitors	☐ No ☐ Yes
Freezing Protection for Cooling Infrastructure	Measures that prevent freezing damage to cooling systems, such as insulated pipes, heat-tracing cables, and leak detection systems	☐ No ☐ Yes
Vegetation Management	Extent of combustible vegetation reduction around critical facilities	☐ No ☐ Periodic ☐ Intensive
Perimeter Fire Breaks	Features that limit wildfire spread into the campus	☐ No ☐ Partial ☐ Continuous
Separation Between Critical Assets	Spacing between critical structures to reduce fire exposure and building-to-building burning	☐ Wide ☐ Moderate ☐ Tight
Wildfire Ember Protection	Features that prevent embers from igniting vulnerable components, such as ember-resistant vents or fire-classified roofs	☐ No ☐ Ember-Resistant Vents ☐ Fire-Classified Roofs
Wildfire Smoke Protection	Measures that reduce smoke intrusion into critical structures, such as smoke filters or smoke-reducing ventilation systems	☐ No ☐ Smoke Filters ☐ Smoke-Reducing Ventilation

Conclusions

As hyperscale data centers continue to expand in size, value, and complexity, they represent a growing opportunity for the (re)insurance industry.

More than $550 billion has been invested in hyperscale data center construction in the US since the first facilities were built in the early 2000s—at least 85 percent of the data center market is now hyperscale. Over the past five years alone, more than $130 billion in new construction value has been added. When IT equipment is included, the total investment during that period exceeds $400 billion, creating an increasingly valuable concentration of insured assets.

As these facilities grow in scale and complexity, accurately differentiating risk becomes increasingly important.

Catastrophe models provide a framework for quantifying these risks by integrating engineering characteristics, hazard exposures, and structure-level vulnerabilities. With detailed exposure data and appropriate modeling approaches, insurers can better differentiate risk across facilities and support more informed underwriting, pricing, and risk management decisions.

As demand for cloud computing and AI continues to drive investment in digital infrastructure, the need for insurance protection will grow. Insurers that develop a detailed understanding of hyperscale data center vulnerability will be better positioned to capitalize on the opportunities offered by this expanding market.