Swiss Re
Corporate Solutions

Risk Engineering Services

Data Centre Builders Risk (Engineering & Construction) Insurance Brochure

This paper focuses on the hardware infrastructure layer of hyperscale data centres, highlighting the key engineering, construction, and risk considerations that arise throughout the project lifecycle of Builders Risk (CAR/EAR), Engineering & Construction. It explores how evolving design standards, high-density computing, and AI-driven workloads are reshaping traditional approaches to construction and risk management.

From an insurance and risk perspective, the continuous developments in this space are widening the gap between project values and available insurable capacity with the insurance markets, while introducing new exposures related to power generation, cooling systems, supply chains, and natural catastrophe risks among others. As a result, successful project delivery increasingly depends on early collabouration between clients, brokers, and insurers to structure robust, forward-looking risk solutions.

This document provides practical insights and guidance to help stakeholders better understand these challenges, manage emerging risks, and support the safe, efficient, and resilient delivery of next-generation data centre infrastructure.

Content

AI is a major growth catalyst for the data centre boom

The rapid adoption of AI technologies is a primary driver of this expansion, requiring immense computational capacity & the demand to accelerate data centre construction.

Unprecedented scale and market challenges

Data centres have been in existence for several decades to support digital infrastructure, originating as mainframe computer rooms, server rooms, and evolving into dedicated structured facilities during the internet boom of the early 1990s to today's AI-optimised hyperscale campuses, Supercomputing and Neo-cloud infrastructure. What sets today's development apart is the sheer scale and speed of development; projects are larger, timelines are tighter, and the pace of construction is unprecedented. In addition, for these massive construction projects and computing speeds the demand of cooling requirements, footprint requirements, power requirement and water requirements are revolving around.

Construction timelines are aggressive, and projects now routinely reach multi-billion-dollar valuations. This surge has captured the attention of every major insurer, as the market grapples with the sheer magnitude of risk.

Groundup construction costs for a moderate size hyperscale data centre now typically begin at around USD 1 billion and can escalate into the multibilliondollar range, with the largest project observed reaching approximately USD 25 billion.[1,2] In contrast, the maximum builder's risk capacity available in many insurance markets is generally limited to about USD 2.5 billion for construction placements,[3] while certain markets can offer up to USD 5 billion,[4] for operational property coverage. Despite this, client demand for true groundup limits continues to exceed what the market can realistically deploy. This widening gap highlights the growing challenge of structuring effective risk transfer solutions at these unprecedented project scales.

The issue is particularly acute in North America and other regions, where projects also face exposure to catastrophic natural hazards, compounded by construction labour shortages and extended lead times for critical electrical components.

By 2030, McKinsey & Company estimates approximately USD 6.7–7 trillion of global data-centre infrastructure spend, with >40% in the U.S., indicating that 2026 is squarely inside an acceleration phase rather than a peak.[5]

According to McKinsey & Company, global demand for data centre capacity is projected to be more than triple by 2030,[5] driven primarily by the rapid expansion of artificial intelligence workloads, which could account for approximately 70% of total capacity demand. In the U.S., McKinsey & Company estimates that more than 50 GW of additional data centre capacity will be required by the end of the decade, representing over USD 500 billion in infrastructure investment. A substantial development pipeline is underway, persistent constraints in power generation, grid interconnection, and infrastructure delivery pose a growing risk that supply may struggle to keep pace with demand.[5]

Similarly, JLL's,[6] latest outlook confirms that approximately 100 GW of additional global data centre capacity is expected to be delivered between 2026 and 2030, effectively doubling current capacity. This expansion is projected to require up to USD 3 trillion of combined investment, encompassing both real estate asset value creation and tenant IT fitouts. The scale of this buildout implies annual capital deployment in hundreds of billions of dollars through the latter half of the decade, driven by the rapid scaling of AI inference workloads, rising rack densities, and increased adoption of liquid cooling technologies.

Asset management and private equity companies like BlackRock, Blackstone and Brookfield are reinvesting heavily in data centres and has become one of the most powerful forces in global data centre infrastructure partnerships.[7,8,9,10]

Recently, Blackrock catalysed such investment by USD 40 billion acquisition of Aligned Data Centres globally, one large deal in data centres space to date.[7]

Additionally, we noted that many of the tech giants we insure are planning capital expenditures ranging from USD 100–200 billion in 2026 alone, primarily for construction activities.

Power first approach 'Where power leads, data centres follows'

Energy infrastructure has emerged as the ultimate determinant in both project site selection, data centre construction and package of risk (in industrial campus set-up).

As a result, regions with strong gas and power infrastructure are increasingly being viewed as viable locations for AI factories beyond traditional hubs such as Virginia, Dallas, and Chicago. Virginia state in the U.S. alone illustrates this challenge and opportunity, with approximately 40 GW,[11] of statewide data centre demand either under contract or in advanced stages of development including engineering, construction, or executed service lease agreements as of late 2024. Notably, Virginia's installed data centre capacity is already comparable to the total installed capacity across Europe, underscoring both the maturity of the market and the growing pressure to diversify to new regions.

Goldman Sachs Research projects that data centre power demand will grow by 165% by 2030, largely due to AI workloads.[12]

Department of Energy (DOE) estimates that while U.S. data centres now consume 4.4% of the nation's electricity, this proportion could rise to roughly 12% by 2028 under high-growth scenarios that have been speculated by the AI industry.[13]

Effective risk structuring and insurance policy placement for integrated power generation within data centre industrial campuses is critical for all stakeholders across the insurance value chain.

De-risking & building the resilience of data centres through Swiss Re's integrated insurance expertise

At Swiss Re, our solutions are broad, and our expertise is deep. We understand the unique risks involved in constructing and commissioning data centres projects where precision, speed, and resilience matter most. We provide a comprehensive suite of innovative solutions to manage risk across wide range of projects, new constructions, expansion projects, modernisation and retrofitting projects.

As a leading insurer, we support some of the largest and most complex data centre project portfolios globally, ensuring resilience against evolving risks, where scale and footprint has no limitations. Our competence is a combination of all streams: Underwriting, Risk Engineering and Claims Management teams.

Our Builders Risk solutions go beyond traditional coverage. We help you safeguard every phase of your project, from early works to project completion, with protection against delays in the form of Delay in start-up/soft costs natural catastrophes, equipment damage, and more.

With a single point of contact and global expertise, we streamline risk management so you can focus on delivering world-class infrastructure. As a trusted partner for some of the largest and most complex builds worldwide, we combine innovation with reliability to keep your projects on track and your investments secure.

As capacity scales rapidly, operators and investors are seeing heightened exposures including physical damage, Delay in start-up/Business interruption (including utility service interruptions), cyber incidents, supply chain delays, regulatory compliance, and environmental, health & safety risks.

Our programmes are tailored to hyperscale, enterprise, colocation, and edge facilities. We align coverage and risk controls to power density, redundancy design (N+1/2N), location-specific perils, and contractual obligations such as SLAs (Service Level Agreement) and lease requirements.

Risk insights to address key challenges associated with hyperscale AI data centres

1. Design configuration

a. Structural integrity & framework:

Use non-combustible construction materials & assemblies. Ensure compliance with local building codes and standards, especially in the U.S, where every state has some additional provisions other than IBC (International Building Code), preferably 2021/2024.

Maintain structural spans per industry best practices and obtain mandatory structural engineer approval. Increase column depth as required for load-bearing compliance with building code.

Source: Swiss Re

Special precautionary measures are required when adopting prefabricated wall systems, prefabricated modular or containerised construction, pre-engineered metal buildings (PEMBs), and data centre pod-based deployments. While these construction methodologies offer clear advantages in terms of speed of delivery and scalability, they often rely on lightweight, composite, or integrated assemblies, which may not provide the same level of inherent structural and fire resilience as traditional reinforced concrete or heavy structural steel construction. Consequently, these systems typically necessitate additional fire-protection and risk-mitigation measures to achieve an equivalent level of integrity for mission-critical data centre operations/AI factories.

Caution needed with another emerging construction trend involves the use of weatherproof, tent-style structures, also referred to as temporary data centre facilities for rapid deployment. These structures are typically utilised as interim capacity during the construction of permanent data halls, allowing operators to bypass traditional construction timelines and bring data centre capacity online as quickly as possible, predominantly within the U.S..

These tent-style facilities generally comprise steel frames or lightweight aluminium substructures, typically paired with a fire-retardant Class B roof membrane. While this roof classification offers moderate resistance to external flame spread, as referenced by OEM guidance, the membrane and associated roof assemblies are often combustible in nature. Recent large-scale temporary installations and event-driven structures have highlighted the elevated fire exposure inherent in such construction methodologies.

Some OEMs operating in this segment are engineering these facilities to withstand wind speeds of up to 130 mph, providing a degree of resilience against severe weather events. However, notwithstanding these wind-resistance measures, this construction approach is less favourable from a fire-risk perspective, particularly in regions that are wildfire-prone, heat-sensitive, or subject to reduced cooling efficiency during peak summer conditions, where operational curtailment may be required to manage thermal constraints.

It is advised that the client evaluate the inclusion of additional passive fire-protection measures to mitigate the heightened fire risk associated with accelerated delivery schedules and lightweight, rapid-deployment construction systems.

b. Layout considerations:

Data centres are increasingly developed as large industrial campuses, with total insured value (TIV) distributed across multiple buildings and exceptionally high-power densities driven by AI-optimised hardware architectures. The transition from traditional 10–20 kW racks to AI-driven configurations exceeding 100 kW represents a fundamental shift in data centre infrastructure and is now common in production environments.

For typical AI training and inference workloads, rack power densities frequently reach 300–400 kW in advanced liquid cooled deployments. At the leading edge, emerging rackscale and podbased designs aligned with Open Compute Project (OCP) standards are being engineered to approach megawatt scale capacity up to approximately 1 MW per rack or pod in next generation AI factories. These extreme density architectures substantially increase fire severity, thermal stress, and overall loss potential relative to conventional data centre designs.

In view of which, buildings are expected to be constructed in accordance with adopted building and fire codes (IBC, IFC, NFPA as applicable). Within multi-building campuses, effective risk separation should be achieved by maintaining a minimum clear fire break of at least 100 ft (approximately 30 m) between data centre buildings, free of equipment, temporary

structures, or material storage. This separation should be preserved throughout construction and operation to limit fire spread and loss accumulation, consistent with NFPA exterior fire-exposure principles.

Internal fire separation within the building is a risk-based approach. Considered, typically, a minimum two-hour fire-rated, non-combustible construction is typically preferred for primary structural elements, including exterior walls and roof assemblies.

Typically, data halls shall be separated from adjacent electrical rooms and mechanical areas by fire-resistance-rated wall assemblies. Two-hour fire-rated walls are commonly accepted as the minimum internal separation where higher ratings are not otherwise required. And, three-hour fire-rated separations/blast walls are preferred, particularly in areas such as electrical rooms housing lithium-ion batteries, due to the increased fire severity and thermal runaway risk.

Where multiple data halls are located within the same structure, a minimum two-hour fire-rated separation between data halls is recommended to limit fire propagation and reduce loss accumulation.

All cable and pipe penetrations through fire-resistance-rated walls must be protected using listed and approved firestop systems that maintain the same fire-resistance rating as the wall assembly. Likewise, ventilation and exhaust ducts penetrating fire-rated barriers shall be equipped with listed fire dampers matching the rating of the penetrated assembly to preserve compartment integrity.

c. Floor & roof considerations:

For AI data centres, it is essential that foundation design, ground-improvement measures, and structural allowances are carefully evaluated and implemented in accordance with project-specific geotechnical recommendations. Given the exceptionally high equipment weights and concentrated loads associated with AI factories, design finalisation should occur as early as possible, to confirm whether deep foundations or soil improvement are required based on site conditions.

Slab-on-grade floor systems are generally preferred over raised floors, as they provide greater structural stability and are better suited to supporting high-density loads and liquid-cooling architectures.

Roof-mounted equipment should be minimised or eliminated where feasible to reduce structural loading, wind uplift exposure, and long-term maintenance risk. Where rooftop installations such as generators, chillers, or other mechanical equipment are unavoidable, roof systems must be designed to safely accommodate heavy equipment loads, maintenance traffic, and environmental impacts including debris, snow, and hail. The use of fire-resistant roof assemblies is recommended to reduce both external fire exposure and internal fire spread.

Data centre roofing systems should also provide robust moisture protection to prevent water ingress that could result in equipment damage or operational downtime. Multi-ply roof membranes are commonly favoured for enhanced durability, while cover boards are recommended to create a stable substrate that improves waterproofing adhesion and resistance to wind uplift.

Rooftop PV installations should be avoided where possible; if implemented, they require careful risk assessment due to added electrical and fire exposure.

d. Load management & heavy equipment:

Structural design should account for both static and dynamic loads, including the substantial weight of cooling networks, manifolds, piping, and plumbing systems. Dynamic loading conditions, such as those arising from natural catastrophe exposures including heavy snow and hail, should be explicitly incorporated into design assumptions.

Hyperscale facilities require robust, large-format support systems for cooling manifolds, often integrated with primary structural columns, and these requirements should be addressed early in layout and structural planning.

Appropriate lifting and installation procedures must be defined and enforced, particularly for modular data centres, pre-engineered metal buildings (PEMBs), prefabricated wall and roof panels, mega-beams, and roof-mounted equipment. These procedures should account for component weight, handling tolerances, sequencing constraints, and site access conditions to minimise construction risk and prevent structural damage.

e. Emerging data centre models beyond conventional terrestrial facilities:

Some unconventional facilities been experimented such as underground data centres, and Floating barge data centres and undersea data centres being deployed on pilot basis by tech giants.

Underground data centres are facilities located in repurposed mines, caverns, or purposebuilt subterranean structures, often hundreds of feet below ground, such as former limestone mines, Cold War bunkers, and hydropowered caverns in North America and Europe. These sites leverage the natural properties of rock formations to support mission critical infrastructure. These structures provide natural protection against extreme weather, seismic events, wildfires, and external threats, with limited access points, however, there are more issues with heat dissipation, scalability, customisation, and lithium-ion batteries underground structures and are not suitable for AI factories.[14]

Underwater data-centre development is gaining momentum across parts of Asia and as a pilot project in Europe, as operators explore subsea deployment to significantly reduce cooling demands through passive ocean-based heat dissipation and to mitigate exposure to natural-catastrophe risks common to land-based facilities. Because the engineering profile of these installations differs fundamentally from conventional data-centre construction, the standards outlined in this publication are not applicable to submerged systems.[15,16,17,18]

Underwater data centres employ a 'subsea module architecture', consisting of sealed, pressure-resistant enclosures designed to withstand hydrostatic pressure, corrosion, long-term marine exposure, and potential anchor-strike hazards. These modules are integrated into the subsea environment and connected to shore through submarine fibre-optic and power cables, routed via engineered, watertight penetrations to maintain structural and environmental integrity.[15,16,17,18]

2. Fire risk exposure (construction & pre-commissioning phases) & indirect smoke damage controls:

Fire protection during the construction phase presents a heightened risk for AI data centres, as life-safety and property-protection systems are not fully commissioned until hot testing and final integration. In phased construction and SIMOPS (simultaneous operations) environments, complete avoidance of sprinkler isolation or system impairment cannot always be guaranteed. Accordingly, the following controls are required to manage elevated fire risk during this period.

a. AI data centres are increasingly adopting lithium-ion batteries for UPS backup due to their advantages over VRLA systems, including rack-integrated battery backup units and the co-location of batteries with other critical infrastructure. In this context, Computational Fluid Dynamics (CFD) analysis is no longer optional but a critical design requirement to validate thermal-runaway behaviour, gas dispersion, smoke control, ventilation effectiveness, and firefighter safety. Without CFD, fire-risk controls rely on assumptions that are increasingly invalid at AI-factory scale; accordingly, CFD analysis is strongly recommended.

b. Fire Prevention for risks associated with immersion cooling fluids needs to be well planned. The closed and sealed nature of immersion cooling system setups mitigates the risk of dust accumulation and potential for electrical fires, while the dielectric liquid used in immersion cooling systems is electrically non-conductive, however additional fire suppressants needed. However hardware failures/leaks could lead to pool fire.
Current mitigation strategies rely primarily on passive controls, including sealed enclosures, fluid containment within the rack or tank footprint, and structural separation. Active monitoring, leak detection, and faultprediction mechanisms remain underdeveloped relative to the risk profile, and that immersion cooling systems require formal HAZOP studies, fluidspecific firerisk assessment, and integration with facilitylevel emergencyresponse planning.[19]

c. In AI and hyperscale data centres, high cooling airflow dilutes smoke and can delay conventional fire detection, increasing the risk of undetected incipient fires. We recommend deploying and validating veryearly, highsensitivity aspirating smoke detection under real operating conditions and ensuring all alarms, including suspected false activations are promptly investigated rather than ignored to prevent escalation and avoid catastrophic loss.

d. Early activation of fire mains, site hydrants during construction.

e. Temporary Fire Protection measures shall be provided & maintained. This includes temporary fire detection systems where permanent systems are not installed and installation of adequate Portable/Wheeled extinguishers appropriate to the risk class as per NFPA 13.

f. Early activation of fire protection prior to Energisation of electrical system and commencement of hot testing. Energisation without active fire protection shall not be permitted.

g. Deploy VESDA in completed data halls, H2 detection in Battery/electrical rooms awaiting fit outs.

h. Manage Hot Work, Combustibles, & Temporary Utilities. Enforce hot work permits, fire watches, and combustible control procedures following OSHA (Occupational Safety and Health Administration) and NFPA51B guidelines. Keep temporary wiring and lighting safe, grounded, and removed promptly when unused.

i. Ensure the construction fire safety teams are fully staffed, trained & deployed prior to general contractor mobilisation and the emergency response plan is fully prepared.

j. Electrical & HVAC Cut-off/isolation Mechanism to be formalised. Prior to hot testing, electrical isolation and emergency shutoff mechanisms shall be installed and verified. HVAC shutdown and smokecontrol interlocks must also be functional to prevent smoke migration and secondary damage across data halls and adjacent spaces in the event of an incipient fire.

k. For multi-storey data centre buildings, verify that standpipe systems are provided for fire-hose connections in accordance with applicable fire-code requirements. Confirm that the local fire brigade has an established emergency response plan, has conducted site familiarisation visits, and is equipped with appropriate access capabilities. Particular attention should be given to ladder reach limitations, as aerial apparatus in many U.S. jurisdictions is typically limited to approximately 75 ft, which may constrain roof-level firefighting and rescue operations.

l. FireResistant Construction Materials and Temporary Works are supposed to be well handled. Scaffolding used for multi-storey structures should comply with best-in-class safety standards, preferably utilising metal frameworks, particularly for data centre construction worldwide. Use only fire resistance (FR)material for construction separations like curtains etc.

m. Isolation from SIMOPS (Simultaneous Operations) is mandatory for phased construction or expansion projects where completed assets are brought into operation while adjacent works are ongoing. Operational areas must be fully segregated from active construction zones, and only fire-retardant containment curtains or modular fire-resistant (FR) barriers shall be used to maintain separation without introducing additional combustible materials.

3. Water/liquid leak damage risk

The transition from air cooling to liquid-based cooling solutions is becoming inevitable in AI data centres. AI servers equipped with high-performance accelerators (such as GPUs/TPU's/NPU's) used for large language model (LLM) training and inference generate heat levels that exceed the practical limits of air cooling which is challenging and not well suitable for loads beyond 100KW+ which is common with AI/hyperscale data centres today.

As a result, these servers are designed with liquid-cooling interfaces and depend on an integrated ecosystem of manifolds, cooling distribution units (CDUs), and external heat-rejection systems. Multiple liquid-cooling heat-rejection architectures are in existence for specific AI server or cluster requirements.

Large-scale and AI-focused data centres can consume up to approximately 5 million gallons of water per day, largely driven by cooling demands, with consumption levels comparable to a small city.

Mitigation measures

a. Establish a comprehensive Water Damage Prevention Plan to mitigate the heightened risk from liquid-cooled infrastructure in hyperscale designs, Hyperscale facilities handle massive liquid volumes under pressure and rely on dense equipment layouts, so even small leaks can escalate rapidly and cause major damage.
   Modern hyperscale data centres increasingly rely on water-cooled or liquid-cooled systems, often using water/glycol mixtures, to support high-density racks and chip-level cooling. These systems require extensive pumping infrastructure, large-diameter supply-and-return manifolds, complex piping networks, and hot-aisle/cold-aisle containment to move coolant to and from the chips (NPU/TPU/GPU/CPU's & accelerators sitting in rack servers).

b. Lessons learned from recent construction-phase losses show that even a single equipment failure, such as a valve rupture, weld failure, pump seal failure, or manifold defect, can result in hall-level flooding, causing millions in damage to raised floors, electrical equipment, and in-progress installations.

To reduce this risk, contractors should implement:

- A formal Water Damage Prevention Plan (WDPP) outlining controls for testing, isolations, inspections, and emergency shutdown procedures during construction.
- Mandatory installation of temporary and permanent leak-detection sensors, automatic shutoff mechanisms, drip trays, and containment barriers before system commissioning.
- Controlled hydrostatic or pressure testing with supervision, water-free testing methods where allowable (e.g., pneumatic testing with strict safety protocols), and stepwise system activation.

c. Implement robust temporary drainage, enclosure, and weatherproofing controls during construction:
   Ensure temporary roofing, guttering, pumps, and diversion channels are in place and tested. All wall penetrations, façade openings, and rooftop cutouts should be sealed or covered ahead of forecasted weather events. Construction teams must actively monitor weather conditions and secure all vulnerable points before storms, as several industry losses have occurred due to unprotected openings allowing rainwater intrusion into near-complete data halls.

d. Monitor weather forecasts and secure all temporary or permanent building openings in advance of storm events. Where the building core and shell are not fully enclosed and significant temperature drops are anticipated, temporary building heating shall be provided to prevent freezing of fireprotection sprinkler systems and domestic water piping.

4. Equipment handling & storage:

Builders Risk policies typically cover the core and shell of the facility and extend to electrical and mechanical infrastructure, including cooling systems, fibre networks, power cabling, powergeneration equipment, and the associated piping and pump networks. Much of this highvalue equipment is delivered on a justintime basis for commissioning. Where delivery or installation timelines are extended, secure and climatecontrolled indoor/ weatherproofed storage should be planned to protect these components during receipt, staging, installation, and tooling.

Additionally, vibration and shock protection should be used during transportation, handling, and installation to minimise the risk of mechanical or electronic damage.

Servers and computing equipment are not typically installed under, or covered by, a Builder's Risk policy, as such policies generally protect the structure and permanent building systems rather than operational IT contents.

5. Supply chain & lead time risks (delay in start-up)

Pre-plan for critical equipment leads times. The demand for electrical equipment and Power gen equipment (if critical load power-up is provided from said equipment), requires the purchase order to be placed during early engineering stages and the project schedule to be in sync based on these committed purchase orders. Long lead times been seen for utility scale transformers, Gas turbines/diesel generators, MV Switchgear, UPS systems, Large Chillers/Cooling units where design changes are required for noise control.

Identify alternative suppliers for critical components wherever required.

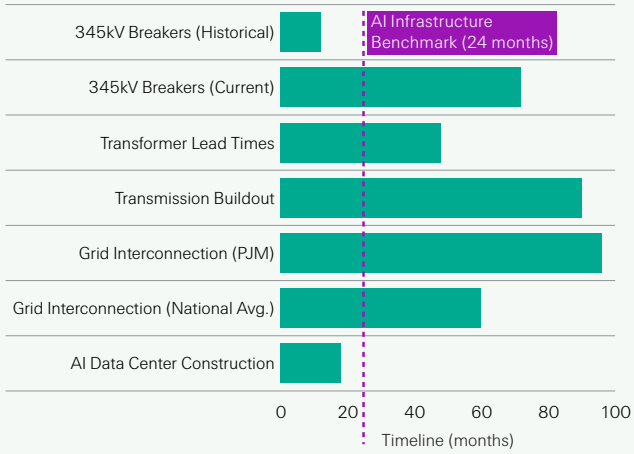

Source: Swiss Re

We are increasingly seeing multi-billion-dollar data centre submissions that include on-site power generation, and lead times for gas turbines can now extend up to 5–7 years due to limited global manufacturing capacity. As a result, some markets are considering used or refurbished/used turbines equipment to meet schedule demands. Please note that mixing used equipment into a project can restrict coverage under London Engineering Group (LEG) conditions, particularly for hot testing, which is limited or excluded for used machinery.

In such cases, it is recommended that on-site power generation assets be submitted under a separate policy, rather than bundled within the main data centre builders risk submission. Additionally, we have observed situations where temporary or leased mobile generation units are installed to accelerate initial power availability for the facility, and these also require separate coverage planning due to capacity availability in the market.

Develop robust project scheduling and contingency planning aligned with vendor procurement orders and committed delivery dates.

Data centre construction requires careful scheduling discipline, especially as hyperscale facilities depend on long-lead equipment such as generators, switchgear, transformers, UPS systems, and liquid-cooling infrastructure. Project timelines should always include appropriate schedule float to account for supply-chain uncertainty and engineering changes.

The rapid evolution of the AI sector shifting from CPU-based compute to GPUs and, increasingly, to supercomputer-class systems can lead to material changes in infrastructure requirements during the project lifecycle. This may result in revisions to planned density, power capacity, cooling strategy, or overall footprint. Builders Risk planning should therefore anticipate that design modifications may occur mid-project, potentially affecting both construction sequencing and final commissioning dates. It is therefore critical that the project has a robust risk register and management of change system to make sure engineering changes are approved by the appropriate level of authority and tracked through to completion of the issued for construction documentation.

Because hyperscale builds involve multi-billion-dollar capital flows, full funding is often staged and not always secured at project inception. When establishing Delay in Start-Up (DSU) exposures, it is important to adopt realistic assumptions regarding commissioning timelines and cash-flow conditions. This helps ensure DSU values are proportional, thereby limiting unrealistic Probable maximum Loss (PML) estimations. Additionally, provisions for tolling or adjusting DSU exposures should be discussed upfront with insurers to accommodate future design revisions, phasing changes, or capacity adjustments expected later in the project.

6. Quality control & testing

a. Conduct rigorous commissioning tests for power, cooling, and Critical load redundancy systems based on Tier-I/II/III/IV configurations.

b. Typically, builders risk has not been requested to cover fitouts (Servers & computing equipment CPU's/GPU's/ Supercomputers frameworks installations) or hot testing. If that's not the case, it should be clearly stated in the Submission for insurance.

c. Hot testing for data centres typically covers Core & shell, Infrastructure such as Cooling network, mechanical equipment and electrical equipment.

d. FAT (Factory Acceptance Test) and SAT (Site Acceptance Test) for critical systems shall be followed.

7. Electrical isolation measures and smoke-damage risk control during construction:

a. Ensure clearly defined electrical isolation points. During construction, all temporary and permanent electrical systems should incorporate well-marked and easily accessible cut-off points. These isolation locations should allow rapid shutdown in the event of an electrical fault, equipment malfunction, or emergency response action.

b. Assess BESS and lithium-ion battery risks early in the project. If Battery Energy Storage Systems or lithium-ion UPS units are included in the build, the project team should evaluate risks such as thermal runaway, off-gassing, and potential ignition.

c. Design battery and high-energy electrical rooms with enhanced protection. Rooms intended to house BESS, large UPS banks, or lithium-ion racks should follow industry norms for segregation, typically minimum 2-hour fire-rated construction, appropriate ventilation pathways, and, where applicable, explosion-mitigation or pressure-relief measures to manage deflagration risks. Ideally these elements should be located against an external wall with fire separation to the rest of the building.

d. Integrate early detection and monitoring systems. To support safe commissioning and reduce electrical incident severity, incorporate advanced detection technologies such as thermal sensors, off-gas monitoring, and high-sensitivity smoke detection. These systems provide early warning of abnormal conditions and support safer energisation of critical systems.

e. Even incipient or localised fire events in data centres can generate significant volumes of corrosive and conductive smoke, which may migrate rapidly through high-airflow environments and cause widespread damage to sensitive electronic equipment. Such smoke contamination can result in prolonged outages, equipment degradation, and long-term operational and reliability issues.
To mitigate these risks, HVAC systems should be integrated with fire detection and alarm systems to automatically modify or reduce airflow upon detection of smoke or fire conditions. Best practices include the use of dedicated HVAC systems for data halls, smoke or fire dampers on shared ductwork, and the maintenance of positive pressurisation in critical equipment rooms to limit smoke ingress and support effective detection and suppression.

8. Power & Energy Risks

Data centres are constructed/developed with power first approach. Power generation is now interfaced with data centre campuses in many builders risk submissions.

Historical context of interface of power generation into Data centres

- Pre 2015: Traditional BackupPowerOnly Data Centres
- 2015–2020: First Signs of Grid Strain
- 2020–2023: Early OnSite Generation & Microgrids
- 2023–2025: AIDriven Power Volatility & New Technologies
- 2025–2026: Turbine Shortages, Energy Storage Expansion
- 2026+: Future State & Risks for Insurers

Modern Data Centres are driving a shift toward onSite and hybrid power generation. As no grid in any part of the world has been designed for taking in such massive surge of power requirements, it is to be noted that the grids are not designed for taking such massive loads. While data centre projects might not be stating the capacity limitations of these factors in submissions about this, but you are still covering hot testing in builder's risk which may expose the electrical system interface and any fit outs (if any) to this voltage dip and surge with grid instability.

a. Grid Capacity Bottlenecks for New Data Centres. This trend is predominant in U.S.

Grid/ region	States most impacted	Near-term capacity headroom (Data centre loads)
PJM	VA, OH, PA, MD, NJ	Extremely limited
ERCOT	TX	Very limited in key load hubs – DFW, Austin, and Houston experiencing volatility and stress
CAISO	CA	Very limited – Transmission constraints and permitting delays
ISO-NE	MA, CT, RI	Limited /seasonal – Winter reliability exposure and constrained generation additions
NYISO	NY	Severely constrained (downstate) – NYC and Long Island dominate capacity limitations
MISO	IL, IN, OH, MI, MN, IA, WI, MO	Moderately constrained, tightening rapidly – Transmission congestion and queue backlogs rising due to data centre and industrial load
SPP	KS, OK, NE, AR, MO, TX (panhandle)	Moderate today, tightening – Historically ample generation, but data centre clustering and transmission deliverability are becoming constraints

Source: Swiss Re

b. Transmission infrastructure challenges:

Transmission capacity constraints and rapidly rising power-load requirements are emerging as critical bottlenecks for the data-centre industry, particularly as hyperscale's accelerate deployment timelines. In parallel, aging transmission infrastructure in many regions is increasingly unable to accommodate new large-scale loads without significant upgrades, further limiting near-term connectivity options.

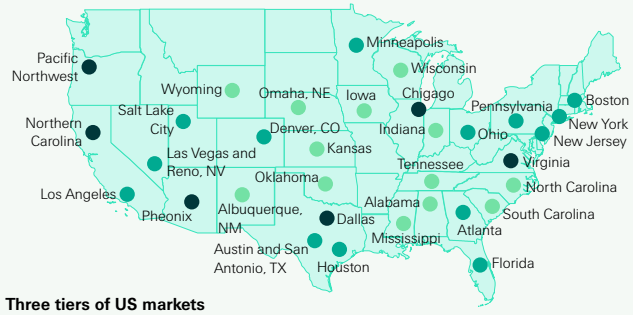

Three tiers of US markets

Primary markets
Large existing demand of more than ~800 MW

Secondary markets
Relatively smaller demand but typically high growth

Emerging markets
Recent hyperscale activity because of cheap and sustainable or cleaner power with negligible co-location presence

Source: McKinsey & Company[5]

c. From Grid-dependent to on-site generation:
Another way of dealing with this stress and local restrictions, with regulatory bills being passed in certain jurisdictions, to address grid constraints and rising energy demands, data centre operators are increasingly moving toward onsite and hybrid power solutions as well. With this approach, we are not only insuring 'data centre', but rather we are covering broader risk which includes power generation, on site substations and additional electronics. These include natural gas turbines, reciprocating engines, fuel cells, renewable generation, and integrated energystorage systems such as largescale battery installations.

d. Data centres are also evolving from passive energy consumers into active grid participants, using loadflexibility mechanisms that help stabilise local grids, way forward, while so far that's not the case the way they are operating, but it's expected in near future, the power generation assets from data centres may pump back power into the grid too, during extreme scenario's like winter storms, etc.

e. AI workloads, especially training, and inference create scenarios where power demand becomes extraordinarily high, followed by sudden collapses in load. This creates a "local blackhole effect" on the grid. This volatility requires technologies that can buffer and stabilise both the facility and the surrounding grid. To manage these extremes, data centres increasingly require formal energystorage systems not only for facility resilience but also for gridlevel stability.

This issue is becoming a central topic across the industry. Groups such as the Open Compute Project (OCP) are now analysing the load profiles, voltage sag events, and transient spikes associated with AI operations. As models move from training to inference at scale, these fluctuations are expected to intensify, creating opportunities for a wider range of storage technologies, which can tolerate repeated highstress cycles and near Zero trips.

To navigate this, we are increasingly seeing the use of equipment that can absorb buffering, stabilising fast power swings.
The designs include & utilise non-rotating equipment synchronous condensers and specialised power- electronics at the data centre campus level to meet demand profiles and hybrid supercapacitors for fast voltage support at the rack level. Even a minor voltage disturbance can create major operational issues, so resilience is now being engineered at every layer of the power chain:

- Dynamic damping equipment such as synchronous condensers
- BESS for load ramp stabilisation
- Statcom, Loadcells
- Supercapacitors installed at the rack
- Demandcharge management or storage outside the building
- Dedicated battery rooms for longduration support. Battery backup units at server rack
- Distributed energy resources located throughout the facility

Technologies that had little relevance to data centres in the past are now entering the sector rapidly with more electrical & electronic equipment seeping in at high voltages. Operators are no longer relying solely on traditional UPS systems' resiliency and reliability are becoming multitiered, integrated functions.

Also, BESS (Battery energystorage design) could be new norm and now must handle sharp current spikes and equalise disturbances quickly either in those campuses or in the region with high density of data centres.

In regions like the Dallas – Fort Worth corridor, where data centre construction continues aggressively, these issues are already visible. Industry analyses and white papers, including recent work by Microsoft, show patterns of voltage sag, inrush currents, and spike behaviour that closely resemble frequencyregulation applications requiring fastacting energy resources to absorb or supply power instantly.

f. Another challenge that is predominant lately since 2025 is 'Gas turbine shortages' in the markets. Natural gas is just dominating data centre power, with large, highefficiency combinedcycle gas turbines backordered for 5–7 years, data centre developers are turning to mobile generators, aeroderivative turbines, reciprocating engines, and refurbished/used units to accelerate deployment. Also, for immediate power supply on site, they are rather going for 'leased equipment' mainly Aero derivate machines to ramp-up the project schedule to bring the system online as often combined cycle construction cannot be ramped up at fast pace. While, these options are more expensive per MWh, and more polluting than baseload CCGTs, but the urgency to

capture AIdriven revenue outweighs those drawbacks.

g. Lithiumion batteries are increasingly being adopted in data centres due to their higher energy density, longer operational lifespans, and lower maintenance requirements compared with traditional leadacid systems. They are most used in UPS applications, where they support critical loads and help maintain continuous uptime during outages or periods of grid instability.

However, despite their advantages, concerns remain, particularly around safety, especially when systems are designed for 3 to 4-hour backup durations. At the same time, new alternative battery chemistries are emerging in the market, prompting data centre operators to reassess longterm energy storage strategies.

With all above intrusion of power generation or overlap of power generation into mainstream 'Data centre', it is more needed than before to understand the that additional insurance implications or builders risk and underwriting during reviews:

- Increased fire, explosion, and mechanicalfailure exposure from dense onsite gas generation
- Higher DSU/businessinterruption volatility due to fasttracked builds and nonstandard equipment
- Increase in challenges with hot testing
- Dutycycle stress on machinery designed for peaking or temporary operation
- Usage of refurbished/leased/used equipment
- Near Zero trips
- Uptime loss

For Builders Risk to properly insure a modern data centre, the full suite of power-generation and power-stability systems design such as single line diagram needs to be finalised and must be included in submission, especially because they are essential during hot testing, which is a high-risk activity.

The insurers would be able to help the client understand the risk better as power generation is comparatively high-risk occupancy which is now getting interfaced into mainstream data centre builders risk submissions, if required separating these placements could be seen as increasing common practice and sometimes necessary because of capacity constraints, specialised risk profiles, and different insurer appetites.

9. Non-traditional and emerging mechanical risk drivers

Heat recovery and waste heat reuse are transforming data centres from isolated energy consumers into active components of municipal energy systems, but at the cost of increased mechanical complexity, new risk interfaces, and heightened operational dependency that must be addressed through design, risk engineering, and insurance review.

Modern data centres have traditionally relied on core mechanical systems such as HVAC chillers, cooling towers, fuel storage tanks, and turbomachinery supporting standby power generation. However, evolving sustainability regulations and decarbonisation policies – particularly across Europe, are accelerating the adoption of additional mechanical infrastructure designed to capture, upgrade, and reuse waste heat, rather than simply rejecting it to the atmosphere.

Data centres inherently generate substantial quantities of low-grade waste heat, primarily as warm air or warm liquid exhaust from IT equipment. Historically, especially in air-cooled facilities, this heat was discharged outdoors with no secondary use.

The common approach involves air-to-water or liquid-to-liquid heat exchangers, which transfer low-grade heat into water loops. These loops can then supply local heating networks or municipal district-heating systems, often with the assistance of heat pumps to raise temperatures to usable levels. This model is now supported by multiple operational and near-term examples across Europe, reflecting a broader trend toward integrating data centres into urban energy systems.

The International Energy Agency also notes that waste heat from data centres could supply up to 300 TWh of space heating by 2030 if harnessed at scale.[20]

Because of these programmes, many European data centres now incorporate heat pumps, water loop systems, large heat exchanger banks, and thermal transfer equipment as integral parts of the mechanical plant. These systems are typically located on ground floors or in mechanical yards/basements alongside traditional equipment such as fuel tanks, pumps, and generator infrastructure.

From a risk perspective, this additional equipment introduces new machinery breakdown scenarios

- Heat pump systems and large heat exchangers add mechanical complexity and increase potential points of failure.
- Water loop and district heat interfaces create risks associated with leaks, corrosion, thermal stress failures, and integration with third party infrastructure.
- Co-location of heat recovery equipment with fuel tanks and backup generation systems requires careful review of spacing, ventilation, fire protection, and containment strategies.
- Maintenance obligations increase, as waste heat systems are year-round mechanical assets rather than emergency only systems like backup generators.

10. Natural Catastrophe Risk:

While it's common for data centres to select technology and design from a perspective of optimal power, water and network availability for uptime requirements, conducting comprehensive Natural catastrophe evaluation due diligence is critical for insurance as the capacity limit in the markets is restricted.

Also, it is important to note that data centres are often located in large clusters, potentially creating accumulation challenges for insurance coverage alongside challenges for local electricity demand.

Perform site-specific assessments for all relevant natural hazards (e.g., seismic, flood, wind, wildfire). Develop a corresponding action plan detailing required design, construction, and operational mitigation measures.

a. Earthquake-Zone Mitigation: Prefer sites outside major fault zones when feasible. Where seismic risk is present, design structures in accordance with local building codes and applicable ASCE seismic design standards. Incorporate structural reinforcement, base isolation, or damping systems as required for seismic resilience. Anchor generators, chillers, racks, UPS equipment, and tanks per seismic restraint requirements.

b. Flood-Hazard Mitigation: Avoid sites located in mapped floodplains. If flood exposure is unavoidable, elevate critical electrical, mechanical, and IT infrastructure above at least the 100-year flood level. Do not rely solely on levee or flood-control systems unless they have been inspected and verified for integrity within the past decade.

c. Hurricane, Tornado, and High-Wind Mitigation: Avoid high-risk wind corridors where practical. Reinforce roofs, walls, and building envelope components for uplift, lateral load, and impact resistance by local wind-design criteria. Recognise that tornado winds create significant uplift forces; if heavy rooftop equipment (e.g., generators, chillers) is unavoidable, ensure it is wind-rated, securely anchored, and installed per manufacturer and code requirements. Consider sloped or aerodynamic roof designs; flat roofs are more vulnerable to uplift. Minimise projections, overhangs, irregular geometries, and architectural features that can increase wind turbulence or debris hazards.

d. Wildfire-Risk Mitigation: Evaluate surrounding vegetation, climate, and topography to determine wildfire exposure. Establish defensible space around the facility and use fire-resistant landscaping and materials as required by the Authority Having Jurisdiction (AHJ). Utilise non-combustible exterior finishes, ember-resistant vents, and appropriate setbacks from wildland areas.

e. Hailstorm Risk: Hailstorms represent a moderate risk to data centres, particularly in regions exposed to severe convective weather.
The primary exposure relates to damage to the building envelope and core-and-shell elements, most notably roofing systems and cooling infrastructure, which can quickly escalate into operational disruption or outages. Roofing membranes such as EPDM, TPO, and PVC are vulnerable to impact puncture, seam damage, and surface degradation during large hail events.
Repeated hail impacts may compromise waterproofing integrity, insulation layers, and overall thermal performance, increasing the risk of water ingress and secondary equipment damage.
While electrical equipment is generally designed for outdoor installation and limited direct damage is expected, it is recommended that the client verify the hail resistance of exposed components including radiators, louvers, bushings, and control panels, particularly in locations subject to higher hail intensity. Ensuring that outdoor electrical and mechanical equipment is suitably rated or protected for regional hail risk will reduce the likelihood of cascading failures following severe weather events.

References:

1. *Amazon plans to invest $25 billion in Mississippi data centres, create 2,000 jobs*
2. *Gigawatt Scale Power for Next-Gen AI | Fermi America*
3. *Aon names global builders risk lead in push on data centre construction insurance | Insurance Business*
4. *FM raises data centre capacity to $5 billion, citing sector growth | Insurance Business*
5. *The future of US hyperscale data centres | McKinsey.*
6. *JLL Research*
7. *BlackRock, Nvidia-Backed Group Strikes $40 Billion AI Data Centre Deal*
8. *Blackstone Data Centres: A Deep Dive into Their Latest Investments and Global Expansion*
9. *Blackstone (BX) Plans Public Company for AI Data-Centre Buying Spree – Bloomberg*
10. *Brookfield Launches $100 Billion AI Infrastructure Programme | Brookfield Asset Management (BAM)*
11. *Dominion Energy nearly doubles data centre capacity under contract to 40GW – DCD*
12. *AI to drive 165% increase in data centre power demand by 2030 | Goldman Sachs*
13. *DOE Releases New Report Evaluating Increase in Electricity Demand from Data Centres | Department of Energy*
14. *10 Largest Underground Data Centres in the World*
15. *https://www.forbes.com/sites/suwannagauntlett/2025/10/20/china-has-an-underwater-data-centre-the-us-will-build-them-in-space/*
16. *Microsoft finds underwater data centres are reliable, practical and use energy sustainably – Source*
17. *5 Largest Underwater Data Centres in the World*
18. *China Powers AI Boom with Undersea Data Centres | Scientific American*
19. *Immersion cooling safety in a data centre – DCD*
20. *AI and Heat Pumps: How Data Centres are Shaping the Future of Energy, according to the IEA in a New Report – HPT – Heat Pumping Technologies*

Explore more on AI risk and infrastructure

Swiss Re Corporate Solutions America
Holding Corporation
222 West Adams Street Suite 3000,
60606 Chicago,
United States (USA)

For queries, please reach out:

Madhu Latha Palikala
Office: +1 312 821 3814
Mobile: +1 312 838 0023
Email: MadhuLatha_Palikala@swissre.com

Rodrigo Davila
Office: +41 43 285 28 43
Mobile: +41 79 548 76 65
Email: Rodrigo_Davila@swissre.com