Data Centres

Kindly supported by:

FORESIGHT

With input from:

HOWDEN

STL PARTNERS

Executive briefing report

WHEN DAYS COST MILLIONS: GETTING AHEAD OF DATA CENTRE DELAY RISK

This report examines why data centre delivery delays occur, what they cost developers, where liability sits, and how CFOs can mitigate delay impacts through project operations and financial planning.



Joe Hurman, Principal Consultant | joe.hurman@stlpartners.com | October 2026

Foreword

Methodology

This report presents findings from research conducted between August and October 2026. It combines the Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), with five in-depth interviews with data centre development professionals, input from Foresight, Howden and Munich Re, and a review of company filings, rating agency reports, and legal and insurer analysis. It builds on STL Partners' 2025 study of the cost of data centre delays, produced in collaboration with Foresight Works. Foresight ran the survey between September 25 and October 2, 2026 and shared the results with STL Partners, which carried out the analysis. Survey results reflect practitioners' perceptions rather than measured project outcomes. US figures are shown separately only where the difference from the full sample is large enough for the US subsample of 49 to support it.

Figures used throughout this report are illustrative, and further detail on methodology is available on request. If you have any questions on methodology or findings, please contact:

joe.hurman@stlpartners.com

Interviewee profiles

No.	Organization type	Role	Region
1	Hyperscale developer	Commercial Manager	US
2	Hyperscale & colocation developer	Head of Site Selection	EMEA
3	Hyperscale developer	CFO	US
4	Hyperscale & colocation developer	CCO	Global
5	Hyperscale & colocation developer	VP, Commercial Strategy and Operations	Global
6	Howden (insurance broker)	Commercial Director	Global
7	Munich Re (insurer)	Senior Underwriter, New Technologies	Global
8	Munich Re (insurer)	Senior Underwriter, New Technologies	Global
9	Foresight Works	Atif Ansar, Co-Founder	Global

Editorial independence

This report was commissioned by Foresight Works, which provides delivery analytics of the kind discussed here, and prepared by STL Partners. Howden and Munich Re, which work commercially with Foresight, contributed an insurance industry perspective. Macquarie Capital Venture Capital, which led Foresight's 2026 Series A, co-hosts the launch event. Interviewees were recruited through an independent expert network and are identified by role only.

The analysis and conclusions are STL Partners' own. Companies and products are mentioned to illustrate market developments, not as endorsements.

Glossary

This report uses the following terms and abbreviations:

AI accelerator – Specialised processors, typically graphics processing units (GPUs), that run AI training and inference. Some developers own and operate them inside the facility.

Anchor tenant – The main customer leasing a data centre phase, often a hyperscaler, whose lease typically underpins the project's financing.

Builder's risk – Insurance, also known as construction and erection all risks, that covers physical damage to the works during construction.

Build-to-suit – A development designed and built to one tenant's specification, usually under a lease agreed before construction starts.

CFCI (contractor-furnished, contractor-installed) – An arrangement in which the contractor procures and installs equipment, and so carries the risk of late delivery.

Colocation – A data centre model in which the operator leases space, power and cooling to one or more customers.

Commissioning – Integrated testing of a facility's systems before handover. It cannot start until the on-site substation is energised, and is often where earlier defects are found.

Contractor compensation / liquidated damages (LD) cover – Insurance that covers a contractor's liability for liquidated damages under the construction contract, up to a policy limit.

Delay-in-start-up – An extension to builder's risk that covers lost income when physical damage delays handover. It does not respond where nothing is damaged.

Design freeze – The point after which a facility's design is fixed, so that later changes need formal agreement and usually add cost and time.

Developer LD gap cover – Insurance that may cover an agreed portion of the difference between eligible delay compensation owed or paid by the developer to the tenant and contractor liquidated damages recoverable by the developer, where the delay is attributable to an insured contractor.

Financial close – The point at which a project's financing is committed and funds can be drawn.

Force majeure – A contract provision that excuses or defers a party's obligations when events beyond its control prevent performance.

General contractor (GC) – The main contractor responsible for building the facility and coordinating subcontractors.

Greenfield – A development on a previously undeveloped site, with no existing buildings or grid connection to reuse.

HVAC – Heating, ventilation and air conditioning: the systems that control temperature, humidity and air quality in a facility.

Hyperscaler – One of the largest cloud and AI providers, which lease data centre capacity at very large scale.

Interconnection facilities – Dedicated infrastructure, such as cabling, that connects the utility substation to a single customer's site. The developer can take on its delivery.

Letter of credit – A bank's undertaking to pay a set sum on demand if a party fails to meet its obligations, so the beneficiary does not have to win a dispute to be paid.

Liquid-ready – A facility designed to support liquid cooling, needed for the high-density racks used for AI.

Liquidated damages (LDs) – Pre-agreed sums a contractor pays for each day of delay it causes, typically capped at 10–15% of the construction contract and excluding additional compensation.

Long-lead equipment (LLE) – Items with long manufacturing lead times, such as transformers, switchgear and generators, that must be ordered early in a build.

Master services agreement (MSA) – A framework agreement between developer and tenant that fixes terms, including delay penalties, across multiple projects.

Megawatt (MW)/gigawatt (GW) – Units of power capacity used to size data centres. One gigawatt is 1,000 megawatts.

Network upgrades – Reinforcement works on the wider utility-owned transmission network, outside the developer's control.

OFCI (owner-furnished, contractor-installed) – An arrangement in which the developer procures equipment and the contractor installs it, so the developer carries the risk of late delivery.

Offtaker – The customer that contracts to take a facility's capacity. In this report, usually the tenant.

On-site substation – The facility that transforms grid power to the voltage used on campus. It must be available before data centre commissioning can begin.

Permit cover – Emerging insurance that covers a permit being withdrawn or challenged in court.

Pre-lease – A lease signed before the capacity it covers has been built.

Ready for service (RFS) – The date on which capacity is handed over to the tenant ready for use. Rent and delay compensations are measured against it.

Rent credits – Free rent granted to a tenant as compensation for late delivery, often banked and traded for lower rent on later leases.

Retained payments – A share of each payment due to a contractor that is held back until the work is complete, and can be used to offset compensations.

Self-generation – Power generated on or near the site by the developer, rather than drawn from the grid.

Service level agreement (SLA) compensations – In this report, the compensations a developer owes its tenant for late delivery under the lease or MSA. Often referred to as monthly rental credits (MRCs) or service level credits (SLCs).

Surety and subcontractor default insurance – Cover for the insolvency of a contractor or subcontractor, which contractor liquidated damages cover excludes.

White space/grey space – White space is the floor area that houses IT equipment; grey space houses the supporting power and cooling plant.

Table of Contents

Foreword.....2
Methodology.....2
Editorial independence.....2
Glossary.....3
Table of Figures.....6
Executive summary: The delay gap.....7
What are the stakes? Escalating SLA compensations in data centre developments.....9
What causes delays?.....13
Who owns the liability, and how do developers protect themselves today?.....16
Are SLA compensations actually enforced? Early signals that times are changing.....20
Closing the gap: Insurance and the CFO agenda.....21
Recommendations for CFOs.....24

Table of Figures

Figure 1: For a major project, what does one month of delay cost as a share of construction cost?…9
Figure 2: Across your portfolio, how material is each source of delay? (top four shown).....13
Figure 3: How often does SLA exposure exceed recoverable contractor liquidated damages?.....16
Figure 4: Developer delay liability stems from a variety of sources.....18
Figure 5: When SLA exposure exceeds liquidated damages, who carries the difference?.....19
Figure 6: What share of your delay exposure is covered by contract or insurance?.....22

Executive summary: The delay gap

93%

of developers owe tenants more for late delivery than they can recover from contractors (US 96%)

79%

cover less than half of their delay exposure through contracts and insurance (US 86%)

$86M

owed to the tenant vs $6-12M recoverable from the contractor by day 120 (illustrative 60 MW phase)

Source: Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), September 2026

Data centre developers are racing to deliver capacity that is already sold. In North America's primary markets, 80% of the 7.5 gigawatts (GW) under construction is preleased and vacancy is 1.4%; worldwide, more than half of 2025 projects slipped by three months or more.[1] STL Partners' 2025 research with Foresight put the cost of a one-month delay on a 60 megawatt (MW) facility at **USD14.2mn**. Practitioners now put it higher: 56% of the 99 data centre finance, risk and development leaders in the Foresight survey say a month's delay costs 3% or more of construction cost. On a 60 MW phase costing about USD816mn, that is at least **USD24mn**.

This report examines what contracts add to that cost: the gap between what developers owe tenants for late delivery and what they can recover from contractors and insurers, who carries the remaining exposure, and whether it is ever collected. This is increasingly an issue which has "gone up to the CEO level," according to one executive at a US-based hyperscale developer. The report finds that:

> These [delay exposure] conversations have gone up to the CEO level.
>
> **Commercial Manager, Hyperscale Developer, US**

- **Tenant compensations add a second bill**, on terms often fixed in master agreements years before handover. For an illustrative 60 MW phase, two days' free rent for each day late (European practice) comes to about USD21.6mn a month and USD86mn at the 120-day cap. US terms of up to 1% of contract value a day would reach USD130mn by day 100, reading contract value as one year's rent (the interviewee did not define it); on the value of the whole lease, the exposure would be many times higher.

- **Delay shows up in commissioning but usually starts earlier**, in installation, equipment and power. Much of it is the developer's own: it typically buys the long-lead equipment (LLE), usually owns the on-site substation and may build the dedicated grid connection, and 52% of respondents (US 41%) rate owner decisions, changes and approvals a high or very high source of delay.

- **Little of the bill can be passed on.** Nearly all respondents (93%) say what they owe tenants exceeds what they can recover from contractors at least occasionally, and 46% say so on half or

1 Sources: CBRE, North America Data Center Trends H1 2026; JLL, global data centre outlook, January 6, 2026

more of their projects. Construction contracts typically waive the developer's claim for lost rent and leave liquidated damages, a fixed daily sum agreed in advance, as its only remedy for delay. Those sums are set far below the developer's loss and cover only the days the contractor causes, so at rates cited in Americas disputes a developer could recover USD6–12mn of the USD86mn it can owe by day 120.

- **In practice, little is collected in cash, for now.** Good contractors are scarce and needed for the next phase, so developers often waive their compensations or settle at cost. Tenants rarely take cash either: "Most of the time, all of those compensations never get paid," one European site selection executive told us. Instead they trade credits for lower rent on later leases or move future capacity elsewhere, so the developer still pays, later and out of sight. That restraint reflects today's shortage of capacity and should not be relied on: in a downturn, tenants seeking liquidity could claim in cash.

- **Insurance closes only part of what is left.** Delay-in-start-up cover pays only when physical damage causes the delay, which most delays do not involve, and only 39% of respondents say they bought it on their last project. Four in five (79%) estimate that contracts and insurance together cover less than half of their delay exposure. Developer liquidated damages (LD) gap cover may pay part of the difference between eligible tenant delay compensation and recoverable contractor liquidated damages, but only where the contractor is liable and subject to the policy limit, deductible, co-insurance and other policy terms.

Most of the gap is set before construction starts. We recommend that chief financial officers (CFOs):

1. **Make delay a core sensitivity in the project finance model**, cause by cause, counting only recoveries that are secured.

2. **Draft contractor and supplier contracts to make recovery collectable**, with compensations that follow the lease formula and are backed by letters of credit or retained payments.

3. **Engage and evaluate the insurance market actively**: involve brokers early and have them set out in writing what current and new cover would pay, then decide what each phase insures, retains or funds.

4. **Invest to avoid delay** rather than delivering at minimum viable cost: reserve factory capacity for LLE and staff commissioning to plan.

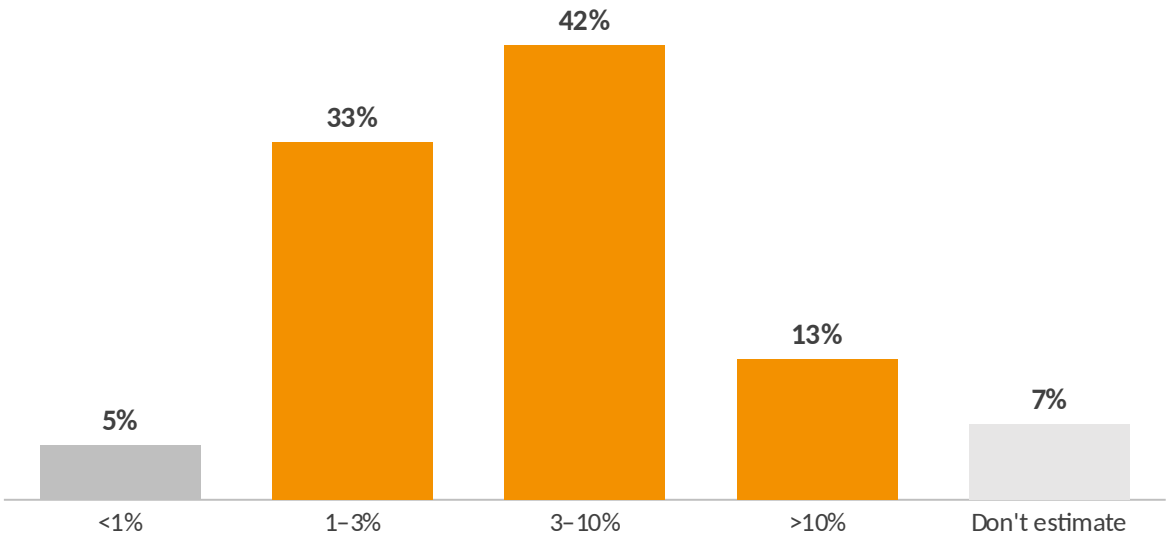
5. **Keep every party working from the same facts**, sharing schedule data with the tenant, lender and insurers so that individual delays do not escalate.

What are the stakes? Escalating SLA compensations in data centre developments

Delays in delivery of data centre facilities are no longer a noteworthy event. In fact, delays are increasingly becoming the norm as the significant expansion in demand for development services, combined with growing pressure on rapid construction velocity, puts pressure on delivery from both sides. With multiple parties managing significant cross-dependencies, when something gives, the resulting delay and its financial impact can be hard to rein in. As one European site selection executive told us, "almost every data centre I've ever been involved in has been delayed to some extent."

These delays can escalate to millions of dollars in financial impact to the developer each week beyond the target handover date. Most respondents (88%) put the financial impact above 1% of construction cost per month, and with liquid-ready greenfield colocation costing approximately USD13.6mn per MW to develop in Northern Virginia (a benchmark mid-cost market), it is clear to see the financial impact running into the tens of millions of dollars.

Figure 1: For a major project, what does one month of delay cost as a share of construction cost?



Source: Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), September 2026

STL Partners' 2025 research with Foresight put the cost of a one-month delay on a 60 MW facility at USD14.2mn, roughly 1.7% of the construction cost of a phase costing about USD816mn. That estimate assumed rent of USD180 per kilowatt (kW) a month, the rate used throughout this report. Practitioners now put it higher: 56% of respondents say a month's delay costs 3% or more of construction cost, which on the same phase is at least USD24mn. The survey captures practitioners' perceptions and is not a re-estimate of the 2025 figure, but the higher figure is likely symptomatic of the growing scale and complexity of projects resulting in ever more prevalent delays, and greater financial impacts of these, even on a per MW basis.

Revenue loss and SLA compensations

Revenue loss and service level agreement (SLA) compensations a data centre developer incurs from its anchor tenant are the most pressing concerns when it comes to the financial impact of delays. Interviewees described two common structures to calculate delay SLA compensations used by the largest offtakers, such as hyperscalers.

1. The most commonly cited practice involves an SLA compensation calculated as free rent for a multiple of the delay period. Specifically, one US developer CFO described pushing for compensations that ramp up with the length of the delay: 1x monthly rent for the first 30 days, 1.25x for 30 to 90 days and 1.5x beyond that, capped at 12 months. On the indicative 60 MW project with USD129.6mn of annual rent (USD10.8mn monthly), a full year's delay ramping up to 1.5x as outlined above would still cost around USD184mn, 23% of the USD816mn build. In Europe, one site selection executive described the standard practice for delay compensation as two days' free rent for each day the ready for service (RFS) date extends past the schedule, applied specifically to the portion of capacity which is contracted with that specific RFS date. In this example, credits stop building up at 120 days late. For an illustrative 60 MW phase renting at USD180 per kW a month (about USD10.8mn a month), that is about USD21.6mn of free rent for each month late, reaching USD86mn if the full 120-day limit is reached. At this point, contract terms state that the tenant has cause to terminate the contract, although such is the scarcity of available capacity that this rarely happens in practice.

2. In the US, another interviewee described compensations of up to 1% of contract value a day or 5% a week, capped at up to 100%, so the maximum is reached in three to five months depending on the methodology used. The interviewee did not say what contract value refers to, so this report reads it as one year's rent. On the same illustrative phase, on that reading (about USD130mn), 1% a day reaches the cap, USD130mn, by day 100. If it instead means the value of the whole lease, the exposure would be many times larger, so developers should check the definition in their own agreements.

In many instances these delay compensations are independent of specific projects, and fixed in the master services agreement (MSA) years before handover, leaving only the amount and schedule to be negotiated for each project. And these MSAs, and the ongoing multi-project relationship between data centre developer and customer which the MSAs formalise, are often what hamper the ability of the customers to enforce such SLAs, a topic which is covered later in this report.

STL Partners dives into the topic of data centre delays, including causes, impacts and remedies, in more detail in our 2025 report "Preventing multimillion dollar data centre losses through reporting". In addition to revenue loss and SLA compensations, four additional direct financial impacts, and two indirect financial impacts, are covered in that report, and a summary of each is given below:

Other direct financial impacts

Labour cost overruns

Contractors, project managers and engineers stay on beyond schedule. Contractors may bear the cost, but only if the delay is proven to be specifically their fault, and minimum labour commitments

mean paying even when work stalls. Operations staff hired for the planned start date must also be paid.

Delays can also cause churn among the individuals working on the development. Many skilled workers move between short-term contracts, and if a delay runs past their availability window, they can leave for other jobs. This can cause shortages and a loss of expertise, often at a critical phase of the build, extending the delay and increasing the financial impact.

Beyond this sits a long tail of further costs that grow as a delay runs on:

Cost of capital

Large projects are typically financed with 20–40% equity and 60–80% debt[2], with repayments aligned to revenue dates. Delay brings unplanned interest and can force refinancing on tighter terms. On USD570mn of debt for a 60 MW facility (70% debt at USD816mn development cost), and at 6% interest, one extra month of interest before repayment starts costs about USD2.9mn. Equipment purchases timed to favourable exchange rates can also be disrupted, which is especially impactful for data centre developers who also own and operate the AI accelerators (typically graphics processing units, or GPUs) housed within the facility.

Equipment storage

Sensitive components such as the aforementioned AI accelerators, as well as switches, servers and other equipment across white and grey space, need controlled temperature, humidity and air quality. This prevents installation until heating, ventilation and air conditioning (HVAC) and related works are complete, and, notably, until the substation has been delivered and commissioned. Delay can force developers to store them in specialist conditions on short notice, eroding their value against depreciation. In extreme cases a new model is released before installation, changing customer specifications and resulting in stranded duplicate equipment.

Construction equipment

Cranes, scaffolding and heavy machinery are often rented on fixed terms, so delay pushes their cost beyond budget. Unlike labour, rental fees accrue whether or not the equipment is used.

Indirect financial impacts

Reputational damage: The long-term trust deficit

Delays erode trust among customers, partners and investors. Hyperscalers score developers on metrics such as schedule adherence and cost efficiency, so a single late handover can count against future bids.

Competitive positioning: Losing market share to faster competitors

Developers that deliver on time gain a clear advantage while supply is constrained. In new regions, delay can cost first-mover advantage for years, especially where power is scarce and the first to secure grid access and permits can lock out competitors.[3]

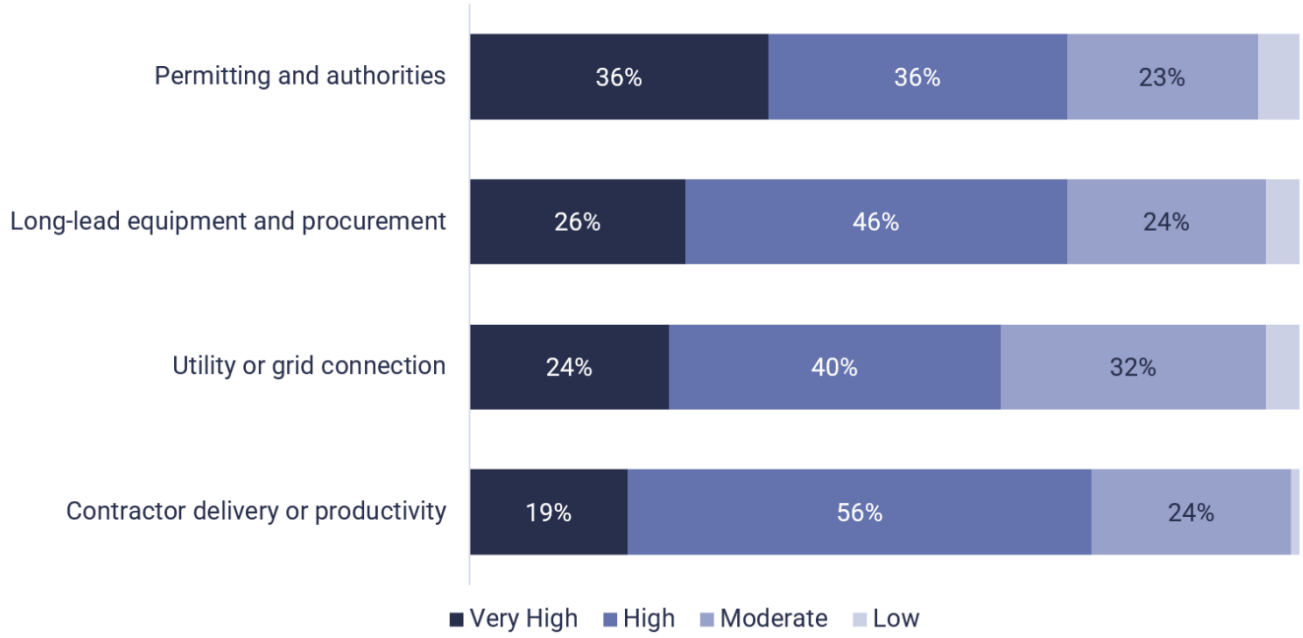
2 Source: Norton Rose Fulbright
3 Source: DatacenterDynamics

There is an additional tangible benefit to being able to repurpose capabilities and funding towards new projects which is compromised if there is a delay on an existing development. This can lead to knock-on impacts across a portfolio of data centre developments. This is particularly prevalent when considering multiple parallel data centre builds on a single campus.

What causes delays?

The first step to understanding how data centre developers can build resilience into the economics of their data centre builds is to understand what can go wrong. While external factors such as access to power and permitting often dominate the trade media when it comes to data centre delays, developers can find just as great a threat to project timelines, and to the opportunities that follow, by looking inward.

Figure 2: Across your portfolio, how material is each source of delay? (top four shown)



Source: Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), September 2026

In the Foresight survey, excluding permitting as this is typically secured prior to construction commencing, respondents rate three key sources of delay risk above the rest:

- Contractor delivery or productivity – arises when work is installed poorly and the fault is found only once it is tested. Poor fit-out by the general contractor (GC) often goes unnoticed until commissioning, when integrated testing exposes it and the work has to be stripped out and redone. This happens at the point in the build with the least slack left in the schedule, and while it is often associated with the commissioning process as this is when the fault is found and remedied, the seeds are sown for this delay much earlier in delivery. Skills shortages make this worse. STL Partners' 2025 research into data centre delays found that a GC can need 10 to 15 projects before it fully understands a hyperscaler's standards[4], with the explosion in data centre builds meaning more and more projects are being delivered by those inexperienced in delivering such builds, both at a personal and organisational level. It also found that less experienced site teams cause more rework and failed quality checks, and report progress less accurately, so defects are both more likely and more likely to stay hidden until testing. Poor workmanship is the

4 Source: STL Partners, Preventing multimillion dollar data centre losses through reporting, March 2025

contractor's responsibility, so in principle the delay is covered by the contractor's compensations. As later sections show, however, these compensations are capped and rarely collected in full.

- Long-lead equipment and procurement – lead times for LLE, such as the transformers described below, were already long and have since stretched further. Supply chain disruption during COVID nearly doubled them, and average equipment lead times remain about 50% above pre-2020 levels, at around 33 weeks worldwide. Any design hold-up that delays ordering this equipment falls on whoever procures it. Under owner-furnished, contractor-installed (OFCI) arrangements, common for LLE, the delay is the developer's own; under contractor-furnished, contractor-installed (CFCI) arrangements, it sits with the contractor.

- Utility or grid connection – encompassing any delays in getting power generated, transmitted and transformed into the right power profile for usage within the data centre. Risks to timelines within this domain can be categorised into three areas:

  - Generation delays – usually known early, before a pre-lease is signed, so they weaken the business case of an early-stage development rather than trigger tenant compensations. If generation delays are off-campus and grid-connected, any delay is typically unrecoverable financially and unavoidable in terms of its impact on project timelines, with utilities typically offering no more than a token time-bound compensation and alternative options few and far between. Self-generation is an option in theory, but a three-year equipment lead time plus permitting means it rarely rescues a build if it is already under way.

  - Transmission delays – relate to delivery of fit-for-purpose infrastructure between the location of generation and the data centre campus, typically terminating at a developer-owned substation. This can occur before or after a pre-lease has been signed depending on the project, and the financial impacts depend on the specifics of the issue. Network upgrades relate to reinforcement works across the wider utility-owned transmission network, firmly outside the data centre developer's remit. However, interconnection facilities built to serve a single offtake customer can be shifted from the utility to the data centre developer's responsibility, by taking ownership of the delivery of dedicated cabling connecting the utility substation to the developer on-site substation (or direct to facility).

  - Substation delays – a significant majority of today's data centre builds have an on-site substation, and completion of this facility is a key milestone in any data centre build, with sign-off on the substation required to commence commissioning of the data centre itself. Such is the complexity of the substation build that this is typically delivered by a specialist outside the GC's remit, creating a new layer of complexity in interdependence between parties. In addition, lead times for some high-voltage transformers have risen to several years, from around a year in 2020–21, and large substation transformers are now passing three years, so despite being the first item ordered, they can end up being a blocker to expediting RFS dates, unless these lead times can be worked around.

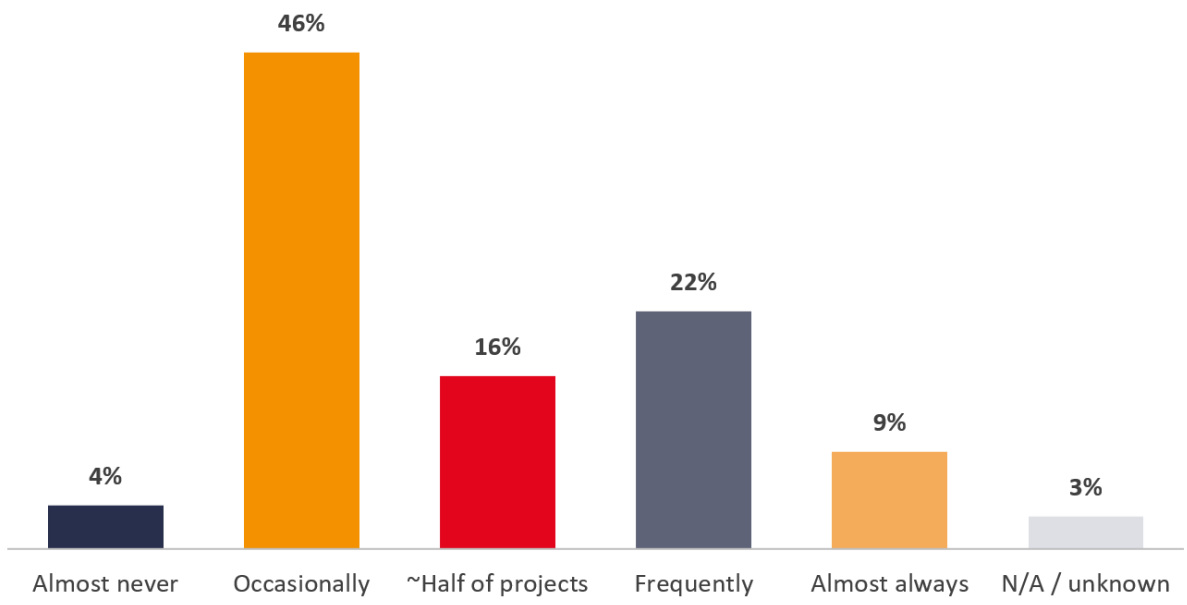
The other delay risks routinely referenced relate to the management of the development itself as opposed to one of the individual delivery workstreams, specifically related to project management and

coordination, as well as changes in requirements from either the developer or their customer. STL Partners' 2025 research corroborates the first of these two areas, identifying that a key challenge in data centre construction was the lag between an issue arising and mitigations being implemented, caused by weak reporting, inconsistent data and siloed communication. Decisions are therefore taken late and on imperfect information, and minor slippage compounds into missed deadlines. Where the developer's own coordination is at fault, the delay is its own and cannot be passed down the chain.

Who owns the liability, and how do developers protect themselves today?

While responsibility for on-time delivery is spread across a growing web of contractors, subcontractors and equipment vendors, the liability for timely delivery sits squarely with the data centre developer. Contractual provisions to protect the data centre developer from delays driven by dependencies outside of its organisation do exist, across contractual interfaces with contractors, equipment vendors, and power generation and transmission; however, these rarely come close to protecting the data centre developer from the financial impact of delays. Nearly all survey respondents (93%) say SLA exposure has exceeded recoverable contractor liquidated damages at least occasionally.

Figure 3: How often does SLA exposure exceed recoverable contractor liquidated damages?



Source: Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), September 2026

A variety of these delay causes relate to contractor delivery, and these can be claimed through liquidated damages provisions included in contracts between a data centre developer and its GC or specialist contractors. These include the following:

- Poor fit-out and rework – defects found at commissioning must be stripped out and redone. Workmanship is the contractor's responsibility, making this the clearest case for liquidated damages.
- CFCI equipment – late delivery of such equipment sits with the contractor. Under OFCI arrangements, common for LLE, the delay is the developer's own.

- Transmission: interconnection – where the developer takes on the dedicated cabling to its site, it can pass delivery, and the delay risk, to a contractor. Wider network upgrades stay with the utility.

- On-site substation – usually delivered by a specialist outside the GC's remit under a separate contract, and a key blocker to timely site delivery as commissioning cannot start until it is energised.

The contractual provisions related to delayed delivery by the GC do not match up to the financial impact of delayed RFS which a developer is exposed to. Standard US construction contracts waive the developer's claims for lost rent and income (American Institute of Architects form AIA A201, §15.1.7), leaving liquidated damages as its only remedy for delay: a fixed daily sum agreed in advance in place of its actual losses, lost rent included. Those sums are set at modest daily rates and capped at a share of the construction contract, so they cover only a fraction of the tenant compensations a developer faces. Contractor delay compensations cited in Americas data centre disputes run at about USD50,000–100,000 a day,[5] or USD6–12mn over 120 days, against the USD86mn the illustrative 60 MW phase could owe its tenant over the same period.

In reality there is a litany of challenges in place which stop data centre developers being able to claim even this reduced amount, including, but not limited to:

- Scarcity of experienced data centre contractors – developers need the same few hyperscale-capable contractors for later phases and subsequent projects, so pursuing compensations is rarely worth the risk of fracturing an existing relationship with an experienced contractor. As one executive at a US-based hyperscale developer put it, "this is right now a GC market." The same restraint extends to equipment vendors: the same interviewee described a USD50mn equipment failure on which compensations were waived because the vendor also supplied the client in other regions.

- Owner-side dependencies – late owner-furnished equipment or mid-build redesigns let contractors argue that the delay is the developer's own. Such interdependencies and complexities are near-inevitable, such is the complexity and time pressure of data centre developments.

- Disputes – contractors pay only for the days they cause, so every claim turns on proving causation with records kept at the time. The dispute process itself takes effort and time from senior staff which would otherwise be used to push ahead with delivery. Pursuing a claim mid-build can also make the delay worse, as working relationships break down on the phases still to be delivered. As one European site selection executive put it, "if we then go into dispute… the rest of the build is just never going to deliver on time either, because they're spending all their time fighting with you on those liquidated damages. It's just enhancing the problem."

Figure 4 summarises the main causes of delay, where liability for each one sits, and to what extent this is recoverable for a data centre developer. The sections that follow show how much of that liability developers end up carrying.

5 Source: Global Arbitration Review, The Arbitration Review of the Americas 2026

Figure 4: Developer delay liability stems from a variety of sources

Area	Delay source	Summary	Responsible *for timely delivery*	Recoverable?
Contractor delivery or productivity	Poor fit-out and rework	Inexperienced GCs and site teams; defects hidden until testing	General contractor	Partial, from contractor
Long-lead equipment and procurement	OFCI equipment	Design hold-ups that delay orders; lead times ~50% above pre-2020	Developer	No
	CFCI equipment	As above, with the contractor placing the orders	Contractor	Partial, from contractor
Utility or grid connection	Generation	Off-campus grid supply; self-generation takes ~3 years plus permits	Utility / generator	*n/a, likely occur before lease agreement*
	Transmission: network upgrades	Reinforcement of the wider utility-owned network	Utility	No
	Transmission: interconnection	Dedicated cabling to site, where the developer takes delivery over	Developer and cable contractor	Partial, from contractor
	On-site substation	Specialist build outside GC remit; 3–4 year transformer lead times	Specialist contractor	Partial, from contractor
Development management	Coordination and reporting	Slow issue detection from weak reporting, inconsistent data and silos	Developer	No
	Changes in requirements	Revised specifications from the developer or its customer	Developer or its customer	Varies, specified in lease

Source: STL Partners

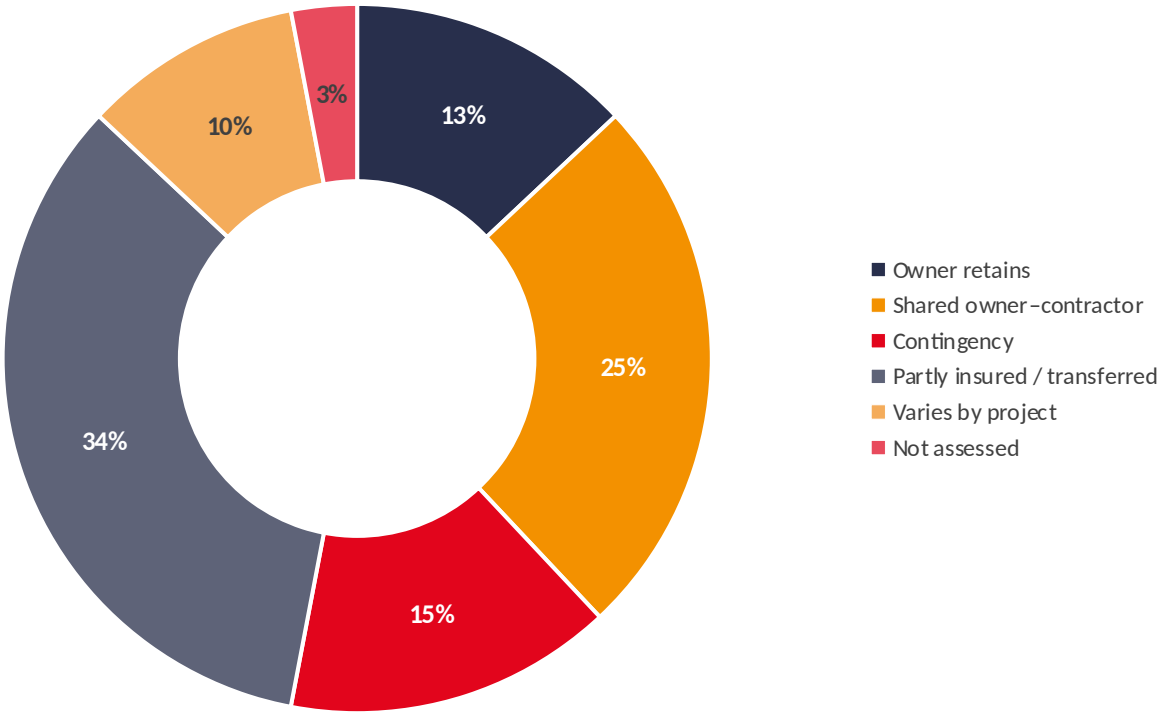
Liquidated damages recoverable from contractors offset only a small part of the exposure. Developers deal with the remaining liability gap in one of three ways.

- Absorb it. Developers hold contingency for the compensation period, manage exposure across all the phases on a campus rather than one at a time, and accept that tenants often take compensations as rent credits, traded for lower prices on later leases. Some do so deliberately, relying on a long-term partnership with the tenant rather than on recovery or insurance. As one US developer CFO put it, they are "happy to carry [the risk] on the balance sheet… and trust that [the tenant] is going to operate as a partner during the build-out phase, not as a judge and jury that's going to execute you."

- Pass some of it on. Builder's risk insurance with delay-in-start-up cover pays only when physical damage causes the delay. Contractor liquidated damages, and insurance covering eligible contractor LD liability, may respond without physical damage. Developer LD gap cover may transfer an additional part of the remaining tenant-compensation gap where the contractor is liable. Leases can also shift some risk back to the tenant, by excusing utility delays, moving the handover date when the tenant changes the design or, in some build-to-suit deals with investment-grade tenants, starting rent on a fixed date.

- Do both. Most developers combine the two, project by project, depending on the tenant and lease terms, their balance sheet and financing, how big the phase is relative to the insurance available, and what lenders will accept. Four in five respondents (79%) say contracts and insurance cover less than half of their delay exposure, suggesting most developers still carry the bulk of it themselves.

Figure 5: When SLA exposure exceeds liquidated damages, who carries the difference?



Source: Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), September 2026

Figure 5 shows that there is no industry standard for managing this liability gap. Each developer's approach depends on factors such as its tenant and lease terms, the strength of its balance sheet and its financing structure. The insurance market is a further factor: products are still maturing to keep pace with the rapid growth in the scale and complexity of developments, so the cover available to transfer this risk varies widely from project to project.

Are SLA compensations actually enforced? Early signals that times are changing

Compensations on paper do not always translate into payments in practice. As one US developer CFO put it, "the SLAs are a compensation in theory only… the biggest compensation is really the missed rent." The same market forces that discourage developers from enforcing liquidated damages against their contractors, set out above, also shape the relationship between developers and their anchor tenants. Tenants rarely enforce a cash payout by their developers, with the paper value of the SLA compensations most commonly traded in for favourable lease rates or other contract terms on either the project in question or future projects. Termination is rare once a tenant has committed to fibre, fit-out and staff, such is the sunk cost into the facility. One European site selection executive told us that, short of a year's delay or more, they did not think anyone would ever use the right to terminate. Of course, this does not negate the financial impact associated with lost revenue, but the SLA compensations charged as a multiple of this are often paid in kind.

The non-monetary impacts of a protracted delay compensation negotiation can also have significant costs for the developer, principally in the cost of losing out on future development projects from the tenant in question, as well as the opportunity cost of time and effort from the management team of the developer in such a negotiation. Such a negotiation also incurs a significant spend on professional services in sourcing a third party to represent the developer in the negotiation.

Developers should not treat this reluctance to enforce SLA compensations as permanent, however. It reflects today's shortage of capacity, and in an AI downturn, AI-exposed tenants seeking liquidity could enforce compensations in cash, so prudent risk management should plan for that exposure rather than assume it away.

Oracle's Project Jupiter campus in New Mexico shows the shift under way. On September 24, 2026, Oracle sent a force majeure notice to Stack Infrastructure, the Blue Owl Capital unit developing the campus, citing potential delays in securing power. The notice would allow Oracle to defer payments if the campus misses its 2028 target to come online. Oracle is not exiting the lease and maintains that the project "remains on our planned schedule," while Blue Owl has said the notice does not change its financial commitments to the project. This is not a tenant collecting compensations: under the contract, securing power for the site is Oracle's own responsibility.[6] But it shows that delay is now contested through formal, public contract notices, and that the cost can land on the developer: here, the lower development-stage rent would run for longer.[7] Prudent CFOs will invest where possible in delivery capabilities to minimise the risk of such delays.

6 Source: Reuters, September 24, 2026 (via BNN Bloomberg)
7 Source: Reuters, September 24, 2026 (via BNN Bloomberg)

Closing the gap: Insurance and the CFO agenda

Today's insurance market covers physical damage well, but cover for delays has remained narrow until the last couple of years. Builder's risk policies are bought on most data centre projects, and only 39% of respondents say they bought the delay-in-start-up extensions sold alongside them on their last project, but both pay out only after physical damage. Delays caused by the owner, such as late owner-furnished equipment, slow grid energisation or the owner's own scope changes, are not covered, yet these are the delay causes most often quoted within the industry. The main policies are set out below, each labelled by how commonly it is held today: standard (bought on most projects), specialist (available, but underwritten selectively) or developing (products repositioned for emerging risks).

Standard policies

- Builder's risk (construction and erection all risks) – Covers physical damage to the works during construction.

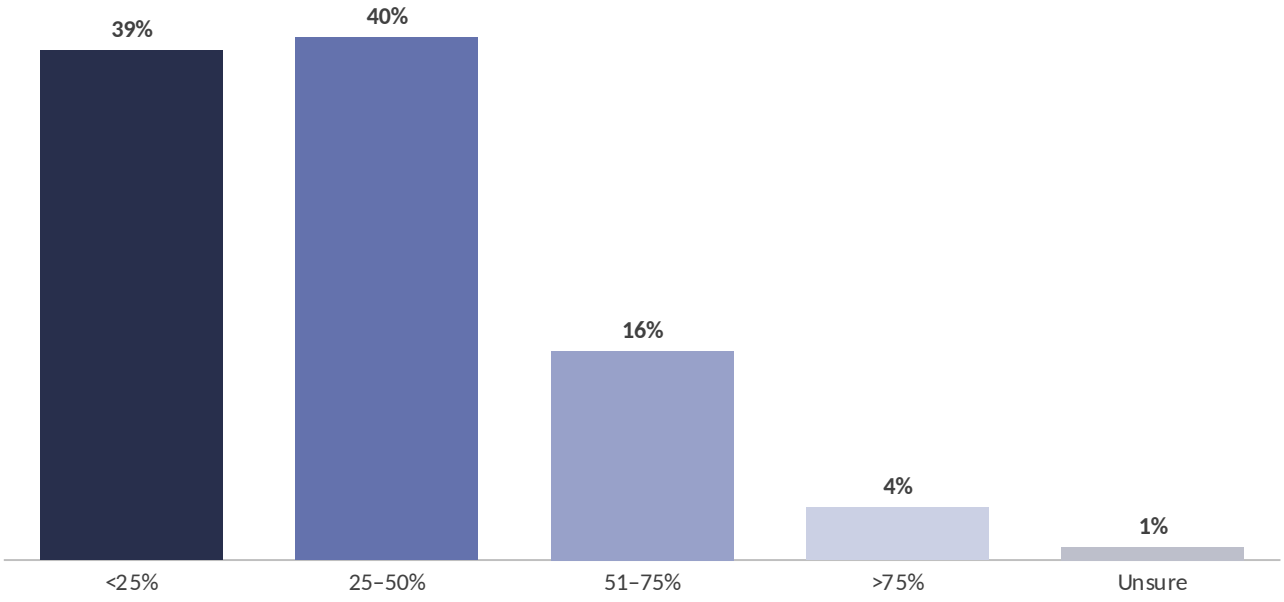
- Delay-in-start-up – An extension to builder's risk that covers lost income when physical damage delays handover. Munich Re notes that it "has been available for years," and only 39% of respondents say they bought it on their last project, but it does not respond where nothing is damaged.

Specialist policies

- Contractor compensation - liquidated damages cover – Covers a contractor's liability for liquidated damages under the construction contract, excluding permitting, material damage and contractor insolvency. These have traditionally maxed out at around USD25–50mn per project, to support the contractor for its liquidated damages payments. Because the cover can pay no more than the contractor's own capped LD liability, it does not by itself reduce the developer's exposure beyond what the construction contract allows.

Together, these policies leave most of a developer's delay exposure uninsured, for three reasons. First, builder's risk and delay-in-start-up respond only to physical damage, whereas most delays arise from late equipment, slow energisation, rework found at commissioning or design changes, none of which damages the works. Second, contractor liquidated damages cover pays only where the contractor is liable; delays caused by the owner, including late owner-furnished equipment and grid connection, fall outside it entirely. Third, where a claim is valid, that cover pays no more than the contractor's own liquidated damages, which are set at modest daily rates and capped under the construction contract: at the daily rates cited in Americas disputes, about USD6–12mn by day 120, against tenant compensations that can reach USD86–130mn within four months on a single 60 MW phase. The result is consistent with the survey: 79% of respondents report that contracts and insurance together cover less than half of their delay exposure, as shown in Figure 6, leaving the remainder on the developer's balance sheet. Interviewees also said the standard construction policies have not kept pace with the scale of today's projects. As one interviewee put it best, "they're treating it the same at 50 megawatts as at two gigawatts… These products need to evolve with the industry."

Figure 6: What share of your delay exposure is covered by contract or insurance?



Source: Foresight survey of 99 data centre finance, risk and development leaders (49 in the US), September 2026

A range of insurance products now seeks to reduce part of the eligible gap and offer developers greater certainty over their future financial returns. Rather than waiting for physical damage, these products respond to the delay itself or to the specific events that cause it: the gap between delay compensation owed to the tenant and contractor liquidated damages recoverable by the developer, or the loss or legal challenge of a permit. Alongside these, surety-backed guarantees can ease the liquidity strain of funding equity up front, giving sponsors more control over when they make their equity contribution. While these policies are not new, there has been a refreshed approach in packaging them for the data centre, responding to the dramatic uplift in number, scale and complexity of data centre projects. Because these risks are harder to price than physical damage, insurers underwrite them selectively. They review the construction contract, the tenant lease terms and the contractor's delivery and compensation history before offering cover, and increasingly draw on delivery data from the project itself to judge how likely a delay is. Such products reduce rather than remove a developer's exposure, but they mark the first meaningful move by the insurance market beyond damage-triggered cover. The main developing policies are set out below.

Developing policies

- Developer LD gap cover – Bought by developers, owners and project companies, it may cover an agreed portion of the difference between eligible delay compensation owed or paid to the tenant and contractor liquidated damages recoverable by the developer. It responds only where the delay is attributable to an insured contractor, and pays out net of a deductible and co-insurance. It is distinct from contractor liquidated damages cover: that policy protects the contractor against its liability to the developer; this one protects the developer against part of the gap between tenant compensation and contractor recovery, subject to the contractor-liability trigger.

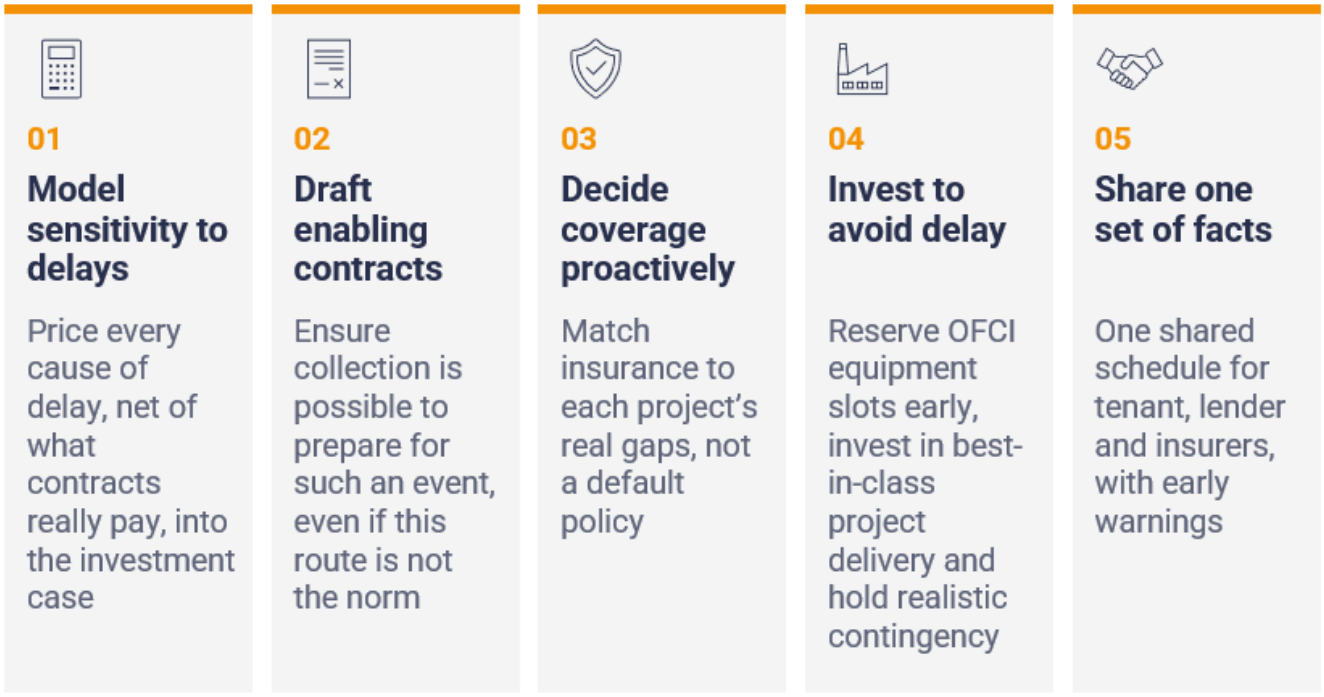
- Deferred Equity Guarantee (DEG) – When financing large infrastructure assets such as data centres or renewable energy assets, equity sponsors will often utilise project finance structures which require them to contribute "equity in first" or proportionately with the debt. These often large cash injections can create liquidity challenges for sponsors who either fund themselves through asset divestments or third-party credit sources. To help manage these challenges, sureties can issue a DEG enabling equity sponsors to have greater control over the timing of their equity contribution.

- Permit cover – Covers losses from the revocation, annulment or material variation of a permit. It can help satisfy a condition precedent to financial close, bring in lenders and investors, and protect against wasted development costs and project delays.

It remains to be seen how widely these policies will be adopted across the industry. Given the frequency and financial impact of delays, however, CFOs should make a deliberate and informed decision on whether and how to use insurance to protect cash flow, rather than leave the exposure unmanaged by default.

Recommendations for CFOs

STL Partners presents five key recommendations to CFOs evaluating the financial impact of delays:



1. **Make delay a core sensitivity in the project finance model.** Only 52% of respondents say they are highly confident of the completion date at financial close, and four in five estimate they cover less than half of their delay exposure. A single delay scenario used to determine downside risk for a project does not provide sufficient information. The model should cover each potential cause of delay (contractor, owner-furnished equipment, utility and permits, tenant design change and physical damage), each with its own likelihood, duration and cost, including tenant compensations, interest carry, lost rent and price concessions on later deals. It should then subtract only what existing contracts and policies would actually pay, counting contractor compensations only where they are backed by a letter of credit or retained payment, as rating agencies do. The output should show the uncovered cost at a cautious completion date and the days remaining before termination and loan deadlines, so that delay is priced into the investment decision rather than discovered after it.

2. **Draft contractor and supplier contracts to maximise recovery, and to make it collectable.** On paper, a developer can recover USD6–12mn of the USD86mn it may owe a tenant by day 120, and interviewees report that even this is rarely collected. Key actions that can be taken here include:

    - Calculate the contractor's compensation with the same formula as the compensations developers face with their tenants, starting with a later handover date and escalating in line with tenant compensations

- Negotiate a higher cap on contractor liquidated damages, towards 25–35% of the construction contract where the market allows. Contractor LD cover, a specialist product where available, can make a higher cap more acceptable to contractors, but it is subject to underwriting and limited capacity, and contractors may price the extra liability into their bids.
- Be as explicit as possible about how responsibility for overlapping delays is handled
- Ensure delay notice periods from contractors and suppliers are kept reasonable, so any delay information can be passed on to the tenant in a timely manner and certainly before the expiration of the developer's own notice deadline
- Enforce compensations via letters of credit or retained payments, so that recovery does not depend on winning a dispute, and tie supplier payments to factory-test and shipment milestones, with deposits protected by bonds

3. **Evaluate the insurance market actively and engage specialty brokers early to choose the right model for each project.** Over half of respondents (55%) say they use insurance to manage the gap, yet only 39% say they bought delay-in-start-up cover, which pays only after physical damage, excludes contractual compensations[8] and carries a 30–90-day waiting period[9]. Developer LD gap cover, the newer non-damage delay cover, may cover an agreed portion of the gap above the contractor liquidated damages recoverable by the developer, but only where the delay is attributable to an insured contractor. Before deciding, CFOs should have brokers set out in writing what current and proposed cover would pay in a variety of realistic delay scenarios. Armed with this data, CFOs can proactively decide for each phase what to insure, what to retain on the balance sheet and what to fund from contingency.
4. **Invest to avoid delay rather than delivering at minimum viable cost.** Half of respondents rate their own decisions as a major source of delay. Developers should:
    - Reserve factory capacity for transformers, switchgear and generators before financial close, as buyers without a reservation wait two to four years
    - Budget for temporary power during commissioning
    - Hold realistic schedule and cost contingency
    - Overprovision commissioning staff if there is previous experience of delays due to commissioning staffing
    - Use delivery analytics and benchmarks from comparable projects to improve time estimates, and track leading indicators, such as daily commissioning headcount against plan, early enough to act
5. **Keep every party working from the same facts, so that individual delays do not escalate.** When delays surface late, one interviewee described three months of forensic analysis on each side

8 Source: Gallagher, Navigating Delay in Start-Up (DSU) insurance for construction projects
9 Source: MDD Forensic Accountants, Stemming the spiralling cost of delay: Delay in Start Up (DSU) or Advance Loss of Profit (ALOP) insurance cover

before a claim could be argued; another, CEO-level decisions on whether to "burn a bridge." Developers should agree test plans and design-freeze points with the tenant early and give the tenant, lender and insurers access to the same schedule data, flagging likely delays early through an agreed escalation route. As tenants and lenders turn to formal notices and termination rights, a shared view of the facts is what prevents one late transformer or substation from becoming a lease or loan default.

Delay compensations are owed in days and recovered, if at all, over years. The developers best placed for the next wave of AI capacity will know their gap before they sign.

STL PARTNERS



Research



Consulting



Events

www.stlpartners.com