Powering AI: Energy, Land and the Real Bottlenecks in Gulf Data-Centre Investment

The constraints behind the buildout that investors must price: where power, land and connectivity bind, and how to invest around them

Chennakeshav Adya

Independent Researcher

June 2026

ABSTRACT

The demand for AI computing is, for practical purposes, unbounded; the supply is not. The binding constraints on the Gulf data-centre buildout are physical, power above all, then suitable land, cooling and connectivity, and these constraints, not demand, determine which projects succeed and what returns investors earn. This paper sets out the real bottlenecks behind the buildout and how an investor should price and invest around them. Drawing on the literature on infrastructure constraints and lead times, on energy economics, and on real-option and timing value, it advances five propositions concerning the bottlenecks in Gulf data-centre investment. Using a stylised, clearly-labelled framework, it ranks the bottlenecks, contrasts the capacity and lead time of the available power solutions, sets out the cost of power by source, weights the factors of site suitability, and shows how power constraints hit a project. The analysis finds that power, in both availability and speed of connection, is the dominant bottleneck; that the value of a project lies substantially in its time-to-power, so that a grid-ready site commands a premium over a greenfield one; that land is valued chiefly for its proximity to power, fibre and water rather than for itself; and that an investor's returns depend on securing the binding inputs early and pricing the lead times that others underestimate. The paper provides a bottleneck framework, a siting checklist and a glossary, and discusses the limitations and avenues for further research.

JEL Classification: L94, Q40, R52, G31, O18

Keywords: data centres, power, energy, land, infrastructure constraints, lead times, digital infrastructure, GCC, real options

1. INTRODUCTION

The story of the AI data-centre boom is usually told as a demand story: artificial intelligence requires vast computing capacity, the demand is enormous and growing, and so capacity must be built. But the demand, however large, is not the constraint. The constraint is supply, and the supply of data-centre capacity is limited not by capital or by willingness to build but by physical bottlenecks, chiefly power, then suitable land, cooling and connectivity, that take time and effort to overcome. These bottlenecks determine which projects can actually be delivered, how quickly, and at what cost, and therefore which investors earn the returns the demand implies. An investor that understands the demand but not the bottlenecks will misjudge the opportunity, because it is in the bottlenecks, not the demand, that the returns are won and lost.

This paper sets out the real bottlenecks behind the Gulf data-centre buildout and how an investor should price and invest around them. It is written for the infrastructure investor, the developer and the lender who must assess not whether demand exists, which it does, but whether a given project can secure the power, land and connectivity it needs, on the timeline it assumes, at the cost it projects. Its purpose is to move the investor's attention from the demand, which is well understood, to the supply constraints, which are less so and which actually determine outcomes. The central message is that power, in both its availability and the speed of connecting it, is the dominant bottleneck; that the value of a project lies substantially in its time-to-power; and that the investor's returns depend on securing the binding inputs early and pricing the lead times that less careful participants underestimate.

The companion papers in this cluster make the strategic case for the buildout and set out the underwriting of a specific asset; this paper examines the physical constraints that both must contend with. Where the strategic paper identifies power and offtake as the dominant value drivers, this paper develops the power and land constraints in depth, and where the underwriting paper treats power security as a gating item, this paper explains why it gates and how to assess it. The three together give an investor the demand, the asset and the constraints, and this paper supplies the constraints that the others take as given.

The Gulf context is, on the power dimension, unusually favourable, which is part of why the region is building so aggressively: energy is abundant and inexpensive, and increasingly complemented by large-scale solar. But abundance at the national level does not eliminate the project-level bottleneck, which is the speed and certainty of connecting firm power to a specific site on the project's timeline, and which depends on grid capacity, substation availability and the permitting and construction of connections. The Gulf's structural energy advantage eases but does not remove this constraint, and the investor must assess it project by project rather than assuming the national abundance flows automatically to its site.

The paper makes three contributions. First, it ranks the bottlenecks and shows that power dominates. Second, it contrasts the capacity and lead time of the available power solutions and sets out the cost of power by source, showing why time-to-power drives value. Third, it weights the factors of site suitability, showing that land is valued for its proximity to power and connectivity, and shows how power constraints hit a project's return. Throughout, the figures are modelled and clearly labelled; they illustrate the constraint economics rather than forecasting outcomes, and they are not forecasts.

It is worth stating plainly at the outset why a paper on bottlenecks belongs in a series otherwise focused on capital and returns. Capital allocators are trained to think about demand, pricing and risk, and less often about the physical and temporal constraints that determine whether a project can be built at all. Yet in a supply-constrained asset class like data centres, those constraints are precisely where the returns are decided: two investors with identical capital and identical demand views will earn very different returns depending on whether they secured power and a connectable site, and how realistically they priced the lead times. The bottlenecks are therefore not an engineering footnote to the investment case but its core, and treating them as such is the orientation this paper urges on the investor. An allocator who masters the demand and the financing but neglects the bottlenecks has understood the easy part of the problem and missed the part that determines who wins.

The remainder of the paper is structured as follows. Section 2 reviews the literature on infrastructure constraints and lead times, on energy economics, and on real-option and timing value, and develops five propositions. Section 3 describes the data and the stylised methodology.

Section 4 presents and discusses the results, covering the ranking of bottlenecks, the power solutions and their lead times, the cost of power by source, the factors of site suitability, the time-to-power, and the impact of power constraints. Section 5 reports robustness. Section 6 sets out implementation considerations. Section 7 offers concluding comments, including limitations and directions for further research. A statement of disclosures, the references and appendices follow.

2. LITERATURE REVIEW AND PROPOSITIONS

The analysis draws on three strands: the literature on infrastructure constraints and lead times; energy economics; and real-option and timing value. This section reviews each and develops the propositions the analysis examines.

The constraints literature also distinguishes between constraints that capital can relieve and those it cannot, a distinction central to the data-centre case. Some constraints yield to money: a costlier site can be bought, a more expensive power solution procured, additional cooling installed. Others are time-bound and resist capital: a grid upgrade takes its construction time regardless of how much is spent, a permit takes its approval time, new generation takes years to build whatever the budget. The binding constraints in data centres are largely of the second kind, time-bound rather than capital-bound, which is why capital abundance, of which there is plenty chasing the theme, does not relieve them and why the projects that succeed are those that secured the time-bound inputs early rather than those with the most capital. This distinction explains a feature that puzzles capital-focused investors: that in a theme awash with capital, the constraint is not money, and throwing more capital at a power-constrained project does not power it. Recognising which constraints capital can relieve and which it cannot is essential to investing in the asset class, and the framework directs the investor's attention to the time-bound constraints that capital cannot solve.

2.1 Infrastructure Constraints and Lead Times

A literature on infrastructure development establishes that the binding constraints on delivering capacity are often physical and time-bound, the lead times to secure inputs, permits and connections, rather than capital or demand (Flyvbjerg, 2014). The consistent finding is that projects are delayed and value eroded most often by underestimated lead times for critical inputs, and that the input with the longest and least controllable lead time becomes the binding constraint that determines the project's timeline. For data centres, power, and specifically the connection of firm power to the site, is typically that input. The literature establishes that an investor must identify and secure the binding-constraint input early, because it is the long-lead item that determines whether and when the project delivers. This underpins the first and second propositions, on the dominance of power and the value of time-to-power.

The infrastructure-lead-time literature also establishes a behavioural regularity directly relevant to data-centre investment: the systematic underestimation of lead times and the optimism bias in project planning (Flyvbjerg, 2014). Across infrastructure types and geographies, projects are found to take longer and cost more than their planners assumed, not at random but systematically, because optimism and competitive pressure bias the planning toward the favourable. For data centres in a boom, this bias is acute: the pressure to deliver into a hot market and to win deals encourages optimistic timelines for power and connection, exactly the long-lead items most prone to delay. The implication is that an investor should treat planners' timelines, including its own, with deliberate scepticism, building in the contingency that the systematic optimism bias implies, and should regard a project whose plan assumes best-case lead times for power as under-provisioned for delay. This behavioural finding is the basis for the analysis's insistence on realistic, conservative lead-time pricing, and it warns that the natural tendency runs the other way.

2.2 Energy Economics

A second literature examines the economics of energy supply, the cost and availability of power from different sources, and the trade-offs between conventional and renewable generation (IEA, 2024). The findings relevant here are that the levelised cost of power varies by source, that renewable sources such as solar are increasingly cost-competitive, particularly in high-irradiance regions like the Gulf, and that the reliability and timing of supply, not only its cost, determine its value to a continuous, power-intensive load like a data centre. The literature establishes that a data centre needs firm, affordable power, and that the Gulf's combination of cheap conventional energy and abundant solar is a structural cost advantage, subject to the firmness and timing the project requires. This underpins the third proposition, on the cost of power by source.

The energy-economics literature also bears on a tension that data-centre investors increasingly face: the trade-off between the speed of securing power and the sustainability of its source. The fastest firm power may be conventional and carbon-intensive, while the cleaner solar-and-storage solutions that the largest tenants increasingly require can take longer to build and integrate, particularly at the scale and firmness a data centre needs. This tension means the investor cannot always optimise speed and sustainability together, and must weigh the time-to-power premium against the offtake and reputational value of clean power, which the companion strategic paper notes the best tenants demand. The Gulf's abundant solar resource eases this tension by making clean power both cheap and, increasingly, deliverable at scale, but the investor must still sequence and structure its power solution to satisfy both the timeline and the sustainability requirements, which is a more complex optimisation than power-cost alone. This adds a dimension to the power bottleneck that the analysis flags and that will grow in importance as clean-power requirements tighten.

2.3 Real Options and Timing Value

A third strand concerns the value of timing and optionality in investment, the premium attached to being able to deliver sooner, and the value of a secured option on a scarce input (Dixit and Pindyck, 1994). For data centres, a site with secured, ready power can deliver capacity into a supply-constrained market sooner than a site that must build its power, and that earlier delivery is worth a premium because it captures demand and revenue that the delayed project misses. The literature establishes that in a supply-constrained market, time-to-delivery has substantial value, and that securing a scarce input ahead of need is itself valuable as an option. This underpins the second and fifth propositions, on time-to-power as a value driver and on securing inputs early.

2.4 Synthesis and Propositions

Taken together, these literatures imply that the binding constraint is the long-lead physical input, power, that time-to-power has substantial value in a supply-constrained market, and that securing scarce inputs early is valuable. From this the paper develops five propositions:

Proposition 1. Power, in both availability and the speed of connection, is the dominant bottleneck in the Gulf data-centre buildout, ahead of land, cooling and connectivity.

Proposition 2. The value of a data-centre project lies substantially in its time-to-power; a grid-ready site commands a premium over one that must build its power, because it can deliver capacity sooner.

Proposition 3. Power cost varies by source, and the Gulf's cheap conventional energy and abundant solar are a structural cost advantage, subject to securing firm, affordable supply on the project's timeline.

Proposition 4. Land is valued chiefly for its proximity to power, fibre and water, not for itself; a site's suitability is determined by its access to the binding inputs.

Proposition 5. An investor's returns depend on securing the binding inputs early and pricing the lead times that less careful participants underestimate; the bottlenecks, not the demand, determine who wins.

3. DATA AND METHODOLOGY

This study is analytical, using a transparent, stylised framework whose assumptions are stated explicitly rather than estimated from a proprietary dataset, which suits a question about how constraints shape outcomes. The figures are modelled and clearly labelled; they illustrate the constraint economics an investor can apply, and they are not forecasts.

3.1 The Bottleneck Framework

The framework has three components: a ranking of the bottlenecks by how often they constrain projects; a power analysis contrasting the capacity, lead time and cost of the available power solutions; and a land-and-siting analysis weighting the factors of site suitability and showing how power constraints hit a project. Together these identify the binding constraints and how an investor should invest around them.

3.1a What the Framework Does and Does Not Do

It is important to be clear about the framework's scope. It is a tool for reasoning about the physical constraints on data-centre delivery, the ranking of bottlenecks, the trade-offs among power solutions, the determinants of site suitability, and the value of time-to-power; it is not an engineering or grid-planning model, and it does not predict the lead time or cost of any specific connection. The figures should be read as illustrations of the constraint economics, the way power dominates, the way capacity trades against lead time, the way time-to-power carries value, rather than as numerical predictions. The detailed engineering, grid and permitting assessment of a specific site requires specialist expertise beyond this framework; what the framework offers is the structure within which to commission and interpret that expertise, so that the investor knows which constraints to investigate most rigorously and how they bear on the investment return.

The framework is deliberately oriented to the investor's decision rather than to the engineer's, because the common failure is not a failure of engineering but a failure of investment judgement about engineering realities, committing capital on optimistic timelines, paying for land without securing power, treating a power constraint as a priced risk rather than a gating one. A precise grid model in the hands of an investor who has not internalised that power gates the project will not prevent these errors; an investor who has internalised the framework's priorities will commission the right engineering and act on it correctly. The framework's contribution is thus to align the investor's judgement with the physical realities, which is the prerequisite for using engineering expertise well rather than a substitute for it.

3.2 Assumptions and Their Justification

The stylised inputs, the bottleneck frequencies, the lead times and costs by power solution, and the site-suitability weights, are calibrations informed by the infrastructure and energy literature (Flyvbjerg, 2014; IEA, 2024) and by the general characteristics of data-centre development; they are indicative rather than precise estimates and are recorded in Appendix A. The lead-time and power-cost assumptions are treated as the central uncertain inputs and are tested in Section 5. The base-case assumptions are set out in Table 1.

Table 1. Base-Case Assumptions for the Bottleneck Framework

Parameter	Assumption
Dominant bottleneck	Power (availability and connection speed)
Grid connection lead time	~12-18 months (illustrative)
New generation lead time	~36-48 months (illustrative)
Solar PPA cost vs grid	~70% of conventional grid (index)
Primary site-suitability factor	Proximity to power
Time-to-power premium	Grid-ready site commands a premium
Power	Abundant nationally; project-level connection binds
Cooling	Water and efficiency a secondary constraint

Indicative, forward-looking assumptions used to illustrate constraint economics; not forecasts. See Appendix A.

4. RESULTS AND DISCUSSION

This section presents the framework in the order of the propositions: the ranking of bottlenecks (Proposition 1); the power solutions, lead times and costs (Propositions 2 and 3); the site-suitability factors (Proposition 4); the time-to-power; and the impact of power constraints (Proposition 5).

A useful way to internalise the inverted structure is to ask, of any data-centre project, what would actually stop it. It is rarely a lack of customers, given the demand; it is far more often that power cannot be secured, that the site cannot be connected, that the permits do not come, or that the timeline stretches until the opportunity passes. These are the things that stop projects, and they are all supply constraints. An investor that habitually asks what would stop this project, rather than will there be demand for it, is led directly to the bottlenecks, and to the diligence and securing actions that address them. This question, what would actually stop it, is a simple but powerful discipline for keeping the investor's attention on the constraints that determine outcomes, and it naturally produces the supply-focused analysis the framework recommends, because the honest answers are almost all about supply.

4.0a Why Demand Is Not the Constraint

Before ranking the bottlenecks, it is worth establishing the premise that distinguishes this analysis: that demand, for once, is not the binding constraint. In most investment, demand uncertainty dominates, the central question is whether customers will want the product, and supply is assumed to follow demand. The AI data-centre boom inverts this. The demand for computing capacity is, for the foreseeable horizon, expected to exceed the capacity the industry can build, so the question of whether there will be demand for a well-located, well-powered data centre is, within reason, settled in the affirmative. What is not settled is whether a given project can secure the power, land and connectivity to be built and to deliver capacity, and when. This inversion, demand abundant, supply constrained, is why the analysis can take demand largely as given and concentrate on the supply bottlenecks, and it is a genuinely unusual feature that investors accustomed to demand-driven analysis must consciously adopt. The investor that keeps worrying about demand while neglecting supply has not adjusted to the inverted structure of this particular opportunity.

This premise should be held with appropriate qualification, as the companion strategic paper notes: demand could disappoint if AI adoption stalls or computing efficiency leaps, and the breadth of demand drivers, not only AI, cushions but does not eliminate that risk. But for the purpose of analysing the bottlenecks, the relevant point is that for any plausible demand path short of a broad collapse, the supply constraints bind before the demand does, so the marginal project's fate is decided by whether it can be powered and built, not by whether its capacity will be wanted. The investor should therefore size its overall exposure to the theme with the demand risk in mind, as the strategic paper counsels, but should analyse each specific project through the supply bottlenecks, because that is where the individual project succeeds or fails. Holding both lenses, demand risk at the portfolio level, supply constraints at the project level, is the complete posture, and this paper supplies the project-level supply lens.

4.1 Ranking the Bottlenecks

The analysis begins by identifying what actually constrains projects. Figure 1 ranks the bottlenecks.

Figure 1. The Real Bottlenecks in Gulf Data-Centre Buildout

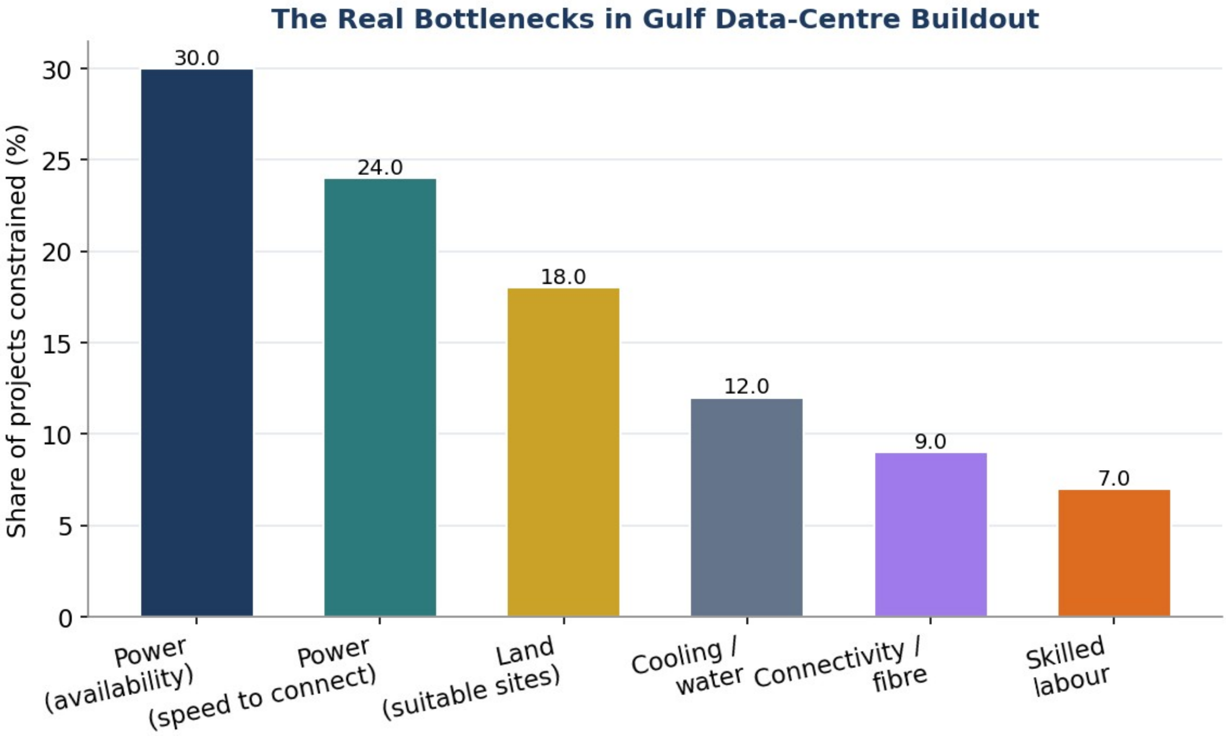

Indicative; power, in availability and connection speed, constrains the most projects, ahead of land, cooling and connectivity.

The figure supports Proposition 1. Power, taken as both the availability of capacity and the speed of connecting it, constrains the largest share of projects, ahead of suitable land, cooling and water, connectivity, and skilled labour. The dominance of power reflects its nature: a data centre is, in essence, a facility for converting large amounts of electricity into computing, so power is its defining input, and the quantity required, rising as AI workloads grow denser, strains the grids and the connection processes that must deliver it. The other bottlenecks matter, and can bind on particular projects, but power is the one most likely to be the binding constraint, and it is therefore the one on which the investor should concentrate its attention. The practical lesson is that an investor assessing a project should ask first and most rigorously about power, because that is where the project is most likely to be constrained, and a project with a clear, secured power solution has overcome the bottleneck most likely to defeat it.

A practical implication of the capacity-versus-lead-time trade-off is that many projects will adopt a phased or hybrid power strategy rather than a single solution. A project might secure an initial allocation of grid capacity to deliver a first phase quickly, while building or contracting additional capacity, a grid upgrade, new generation, or behind-the-meter solar-and-storage, to support later phases, so that it delivers early capacity into the supply-constrained market while scaling its power over time. This phasing aligns the power solution with the demand ramp and the capacity-lead-time trade-off, capturing the time-to-power premium on the first phase while building the larger capacity the project will ultimately need. The investor should therefore think of the power solution not as a single binary choice but as a sequenced plan matched to the project's phasing, and should assess whether the project's power roadmap credibly delivers each phase's capacity on its timeline. A well-phased power plan is often the practical resolution of the capacity-versus-lead-time tension, and the investor should look for it rather than expecting a single solution to deliver both speed and scale.

4.2 Power Solutions and Their Lead Times

Solving the power bottleneck means choosing among solutions with different lead times. Figure 2 contrasts them.

Figure 2. Power Solutions: Capacity versus Lead Time

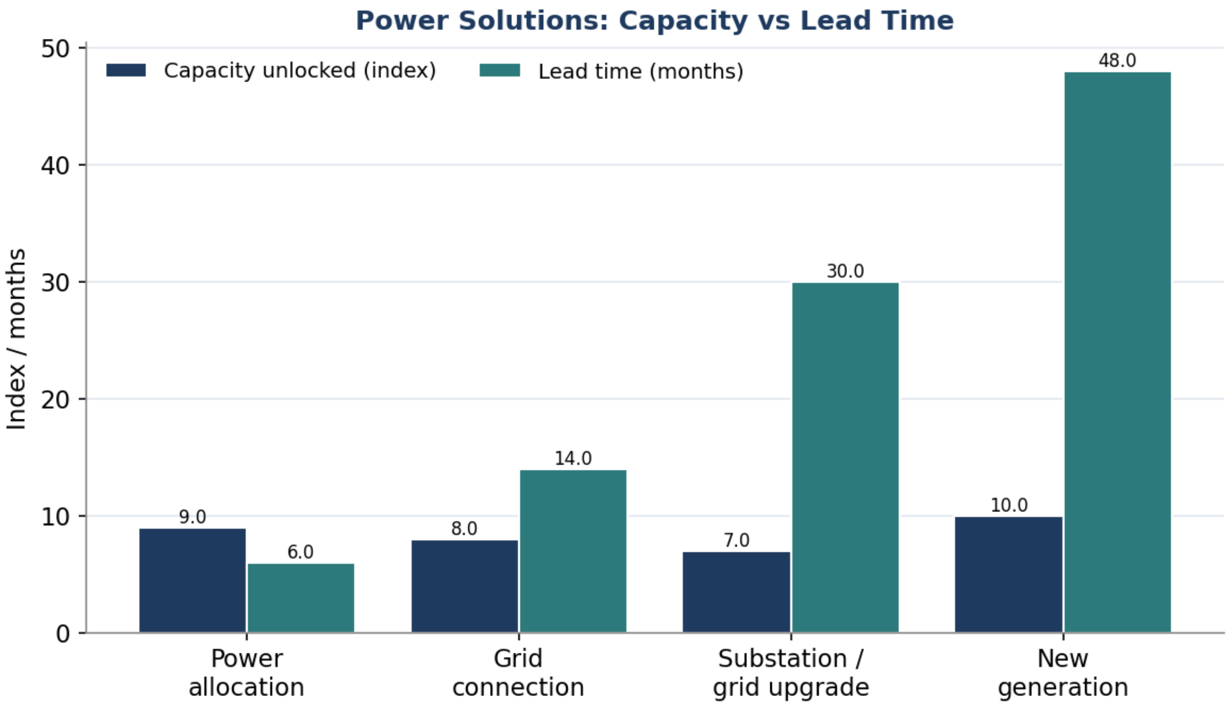

Indicative; solutions that unlock more capacity, such as new generation, take far longer to deliver.

The figure supports Proposition 2. The available power solutions trade capacity against lead time. Securing an allocation of existing grid capacity is fastest but limited by what the grid can spare; connecting to the grid takes longer; upgrading the grid or building a substation takes longer still; and building new generation, while it can unlock the most capacity, takes the longest, often several years. The crucial implication is that the power solution chosen determines the project's timeline, and a project relying on a long-lead solution, new generation or a major grid upgrade, cannot deliver capacity for years, however strong its demand. An investor must therefore match the power solution to the timeline it needs, recognising that the fastest solutions are constrained in capacity and the highest-capacity solutions are slow. The interplay of capacity and lead time is the heart of the power bottleneck, and a project's viability often turns on whether it can secure enough power quickly enough, which is a question of which solution it relies on.

It is worth connecting the power-cost analysis to the sustainability requirements that increasingly shape the offtake, because the two interact in the Gulf to the region's advantage. The largest data-centre tenants increasingly require that their capacity be powered substantially by renewable energy, both for their own commitments and for the expectations of their customers, so the ability to supply clean power is becoming a condition of securing the strongest offtakes, not merely a cost or reputational matter. The Gulf's exceptional solar resource means it can meet this requirement more cheaply than most regions, supplying clean power at low cost and at scale, which positions Gulf assets to attract exactly the premium tenants whose offtakes drive value. This links the power bottleneck to the offtake question the strategic and underwriting companions emphasise: securing not just firm, affordable power but firm, affordable, clean power strengthens the project's ability to secure the strong, long offtake on which its value and financeability depend. The investor should therefore weigh the renewable profile of its power solution as part of the offtake strategy, and the Gulf's solar advantage as a contributor to both cost and offtake quality.

4.3 The Cost of Power by Source

Power cost, as well as availability, varies by source. Figure 3 shows the cost stack.

Figure 3. Illustrative Levelised Power Cost by Source (Index)

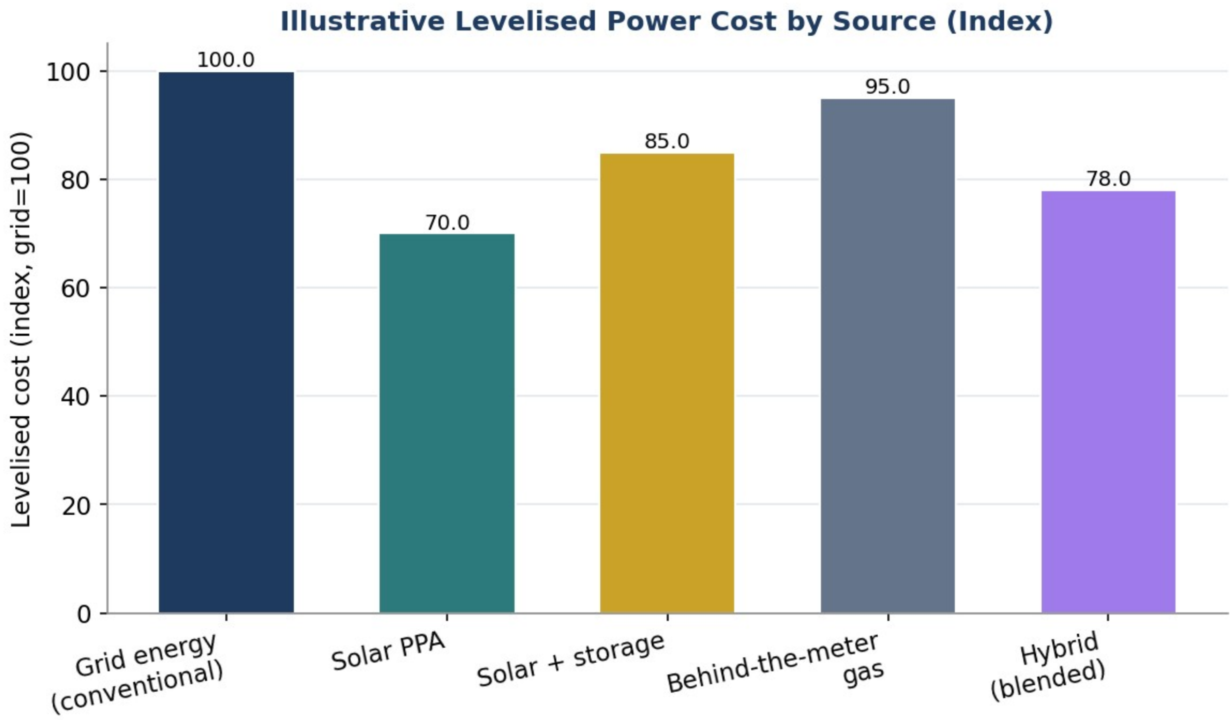

Indicative; solar and blended hybrid sources can undercut conventional grid power in the high-irradiance Gulf.

The figure supports Proposition 3. The levelised cost of power varies by source, and in the Gulf the picture is favourable: conventional grid energy is already inexpensive, and solar, in a region of exceptional irradiance, can undercut it, while solar with storage and behind-the-meter gas occupy intermediate positions and a blended hybrid can optimise cost and firmness together. The Gulf's structural advantage is that it can power data centres more cheaply than most regions, through cheap conventional energy complemented by some of the world's lowest-cost solar, which lowers the largest operating cost of a data centre and improves its economics. The qualification, central to the analysis, is that cost is secondary to firmness and timing: a cheap power source that cannot deliver firm capacity on the project's timeline is of little use, so the investor must secure power that is both affordable and firm and timely, with the Gulf's cost advantage a benefit on top of, not a substitute for, the availability the project requires. The cost advantage makes Gulf projects attractive once the availability is secured; it does not overcome an availability bottleneck.

A practical consequence of power dominating the cost stack is that the choice of power source materially affects the asset's competitiveness over its life, not only at development. Because energy is the largest ongoing operating cost of a data centre, an asset powered cheaply, through low-cost Gulf solar and conventional energy, has a durable operating-cost advantage over one powered expensively, and that advantage compounds over the asset's life into a meaningful difference in profitability and competitiveness for tenants. An investor should therefore weigh not only whether power can be secured and how quickly, but how cheaply it can be secured over the long term, because the power-cost position is a lasting source of advantage or disadvantage. The Gulf's structural energy-cost advantage is, in this light, not merely a development convenience but a durable competitive edge for assets located there, provided the firm, affordable supply is actually secured, and it is part of why the region is attractive for data-centre investment beyond the immediate boom.

4.4 What Makes a Site Suitable

Land is the second bottleneck, but it is valued for its access to the first. Figure 4 weights the factors.

Figure 4. What Makes a Site Suitable for a Data Centre

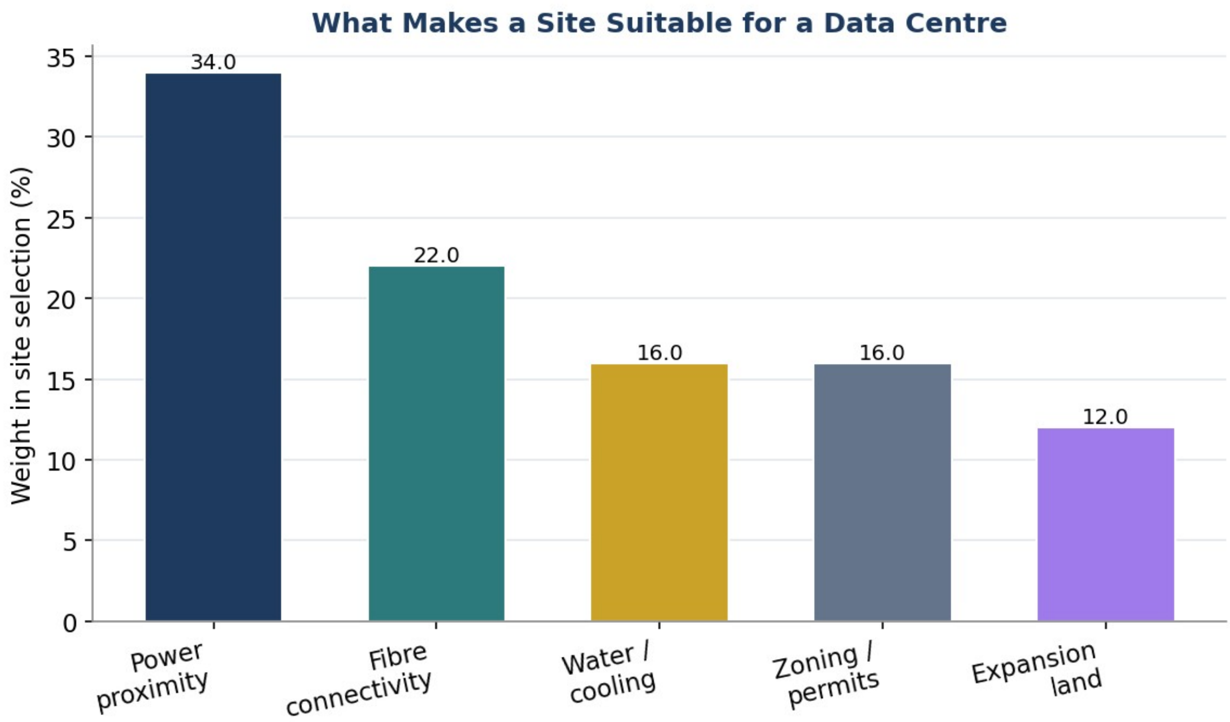

Indicative; proximity to power dominates site selection, ahead of fibre, water, permits and expansion land.

The figure supports Proposition 4. The factors that make a site suitable for a data centre are dominated by proximity to power, the ease and speed of connecting firm capacity, followed by fibre connectivity, water and cooling, zoning and permits, and room to expand. The decisive point is that land is valued not for itself, its area or its price, but for its access to the binding inputs, above all power. A cheap, large site far from power and fibre is less valuable for a data centre than a smaller, costlier site adjacent to a substation and a fibre route, because the former cannot readily be powered and connected while the latter can. This inverts the usual real-estate logic, in which location is valued for its market access; for a data centre, location is valued for its infrastructure access, and the investor selecting a site should weight power and connectivity proximity far above land cost or size. The land bottleneck is, in large part, a power-and-connectivity bottleneck expressed through land, which is why the analysis treats power as dominant even when the apparent constraint is finding a site.

It is worth noting that the time-to-power premium accrues not only to the developer but, ultimately, to whoever controls the scarce ready power, which shapes how value is distributed along the chain. A landowner with a powered, connectable site, a utility with spare allocatable capacity, or a developer holding a secured power position each owns a piece of the scarcity, and the premium for time-to-power is shared among them according to their bargaining positions. An investor entering the chain must understand where in it the scarcity value sits and how much of the premium it can capture versus pay away to those who control the inputs upstream. This is why securing the scarce inputs directly, rather than buying a site whose power must still be obtained from others, is so valuable: it moves the investor up the chain to where the scarcity premium is captured rather than paid. The distribution of the time-to-power premium along the chain is a strategic consideration the investor should map, positioning itself to capture rather than surrender the value that the supply constraint creates.

4.5 Time-to-Power as a Value Driver

The premium on speed is best seen in the time-to-power. Figure 5 contrasts the routes.

Figure 5. Time to Power: Grid-Ready versus Greenfield Build

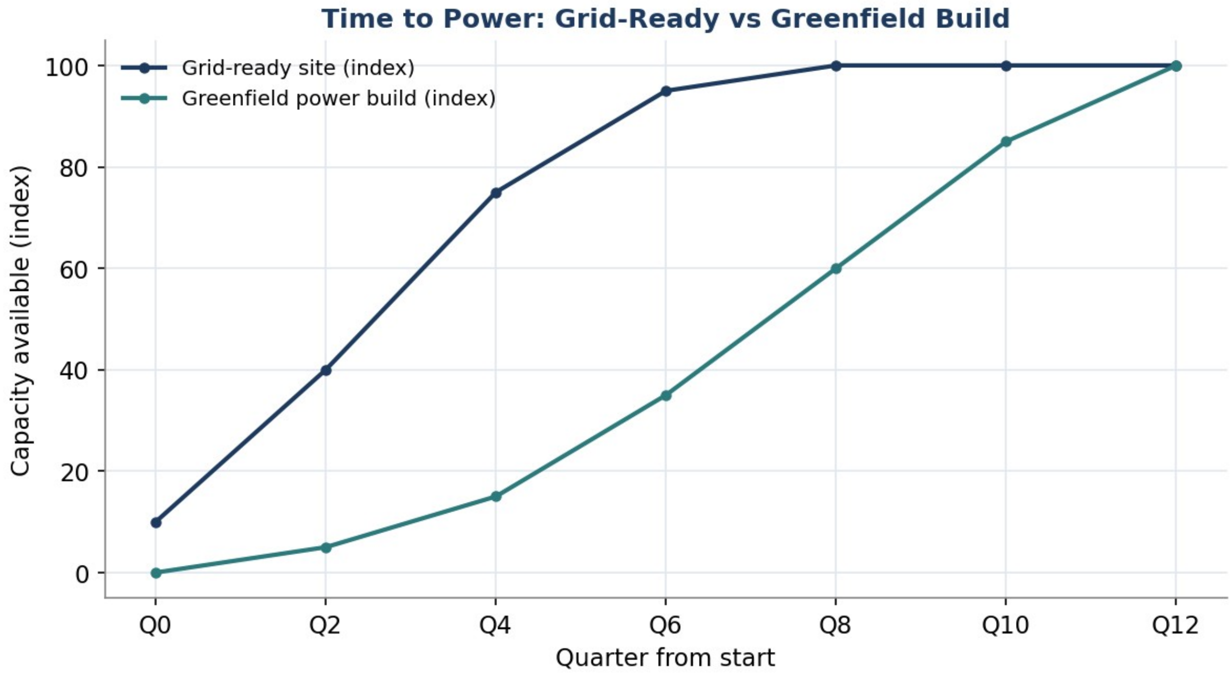

Indicative; a grid-ready site delivers capacity far sooner than one that must build its power.

The figure supports Proposition 2. A grid-ready site, one with secured, ready power, can deliver capacity into the market quickly, ramping to full capacity within a year or so; a greenfield site that must build its power, a substation, a grid upgrade, or new generation, delivers capacity only after years. In a supply-constrained market, this difference in time-to-power has substantial value: the grid-ready site captures demand and revenue during the years the greenfield site is still building, and it does so when capacity is scarce and pricing strong. This is why a grid-ready site commands a premium, often a large one, over a comparable site without ready power: the premium is the value of the time-to-power, the ability to deliver sooner into a market that wants capacity now. For the investor, this means that securing ready power, or a site with it, is not merely a matter of avoiding delay but a source of value in itself, and that the premium paid for a powered site can be justified by the earlier, scarcer-market revenue it captures. Time-to-power is the value driver the supply-constrained market rewards, and the investor that secures it captures a premium the demand alone would not provide.

It is worth distinguishing the two ways a power constraint manifests, because they call for different responses. The first is a hard constraint: power simply cannot be secured for the site in the required quantity, in which case the project is not viable as conceived and the investor must reduce its scale, change its site, or abandon it. The second is a soft constraint: power can be secured but only on a longer timeline or at a higher cost than planned, in which case the project remains viable but its returns are eroded by the delay and cost, and the investor must reprice or restructure accordingly. The investor's diligence should establish which kind of constraint a project faces, because a hard constraint is a reason to decline while a soft one is a matter of repricing and patience. Mistaking a hard constraint for a soft one, proceeding in the hope that power will eventually be found, is the costliest error, because it commits capital to a project that cannot be completed; mistaking a soft constraint for a hard one needlessly forgoes a viable project. Distinguishing the two accurately is part of the power diligence the framework places first.

It is worth adding that cooling, though ranked below power, is closely bound to it and deserves specific attention in the Gulf's climate. The region's high ambient temperatures make cooling more demanding and more energy-intensive than in temperate climates, so a Gulf data centre must devote a larger share of its power to cooling, which both raises its energy requirement and links the cooling constraint directly to the power one. Water-based cooling can ease the energy burden but introduces a water constraint in a water-scarce region, while air and other cooling approaches consume more power; the choice involves a trade-off between the power and water constraints that the investor must weigh for its specific site and design. This climate-driven cooling burden is a Gulf-specific refinement of the bottleneck analysis: it raises the effective power requirement, ties cooling to both power and water, and rewards efficient cooling designs that reduce the energy and water the asset consumes. The investor should assess cooling not as a separate technical matter but as part of the power-and-resource constraint, because in the Gulf's climate the cooling and power bottlenecks are tightly coupled.

4.6 How Power Constraints Hit a Project

The cost of failing to secure power is severe. Figure 6 shows the impact.

Figure 6. How Power Constraints Hit a Project

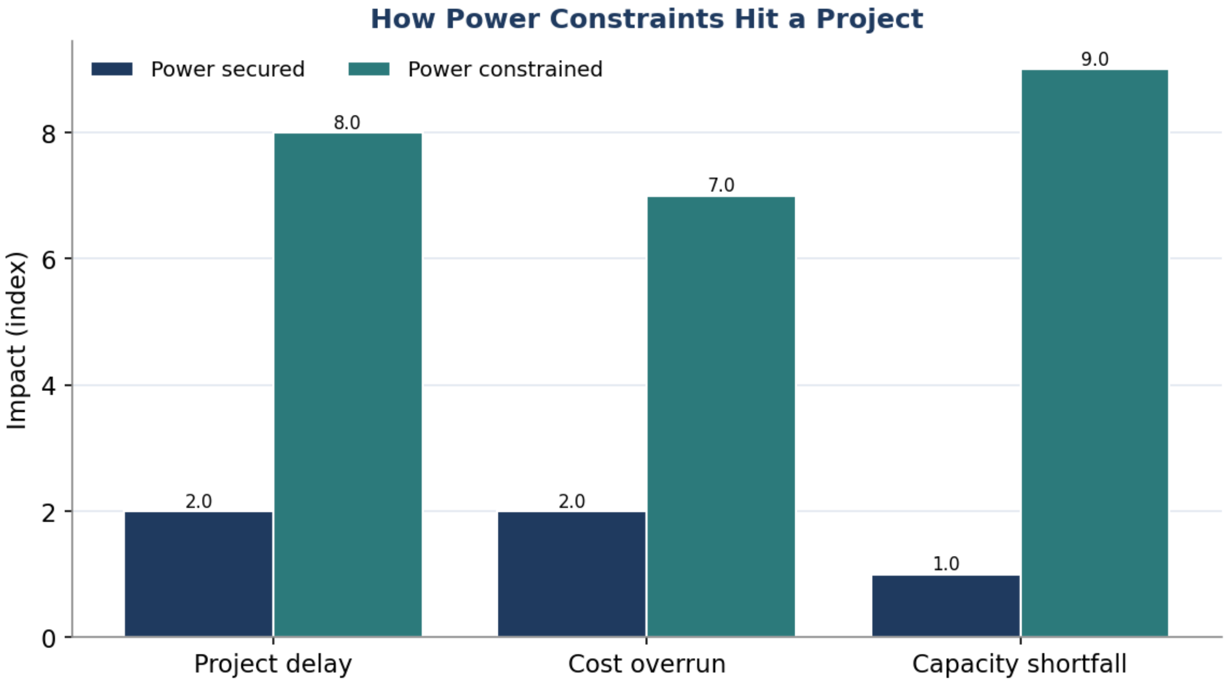

Indicative; a power-constrained project suffers delay, cost overrun and capacity shortfall together.

The figure supports Proposition 5. A project that has secured firm power on its timeline proceeds with little power-related delay, cost overrun or capacity shortfall; a power-constrained project suffers all three, a delayed timeline as it waits for power, cost overruns as it pursues alternatives or carries an idle asset, and a capacity shortfall if it cannot power its full design. These impacts compound: a delay both defers revenue and raises cost, and a capacity shortfall permanently reduces the asset's income. The severity of the impact is why power is a gating item rather than a priced risk, as the underwriting companion paper argues: a power failure does not merely reduce the return, it can undermine the whole project. The practical lesson is that the investor must secure firm power before committing, treating an unsecured power position as a reason to wait or decline rather than a risk to accept for a higher return, because no return compensates for an asset that cannot be powered. The asymmetry, secured power proceeds smoothly while constrained power suffers compounding harm, is what makes power the dominant consideration in the analysis.

Figure 7 places the sensitivity of the project return to the bottleneck variables in a single view.

Figure 7. Sensitivity of Project Return to the Bottleneck Variables

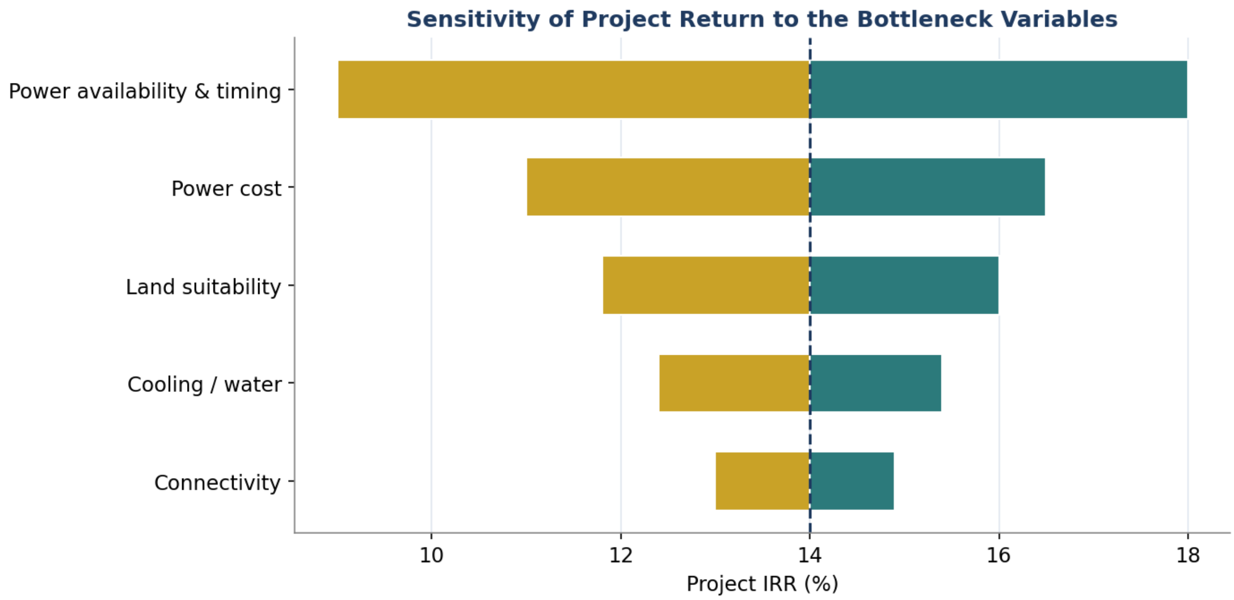

Indicative; power availability and timing dominate, ahead of power cost, land and connectivity.

The sensitivity analysis confirms the hierarchy. Power availability and timing have by far the largest effect on the project return, ahead of power cost, land suitability, cooling and connectivity. This ordering is the central practical message: the project's return is made or lost chiefly on whether it secures firm power on its timeline, and an investor should concentrate its diligence and its securing effort there before all else. The cost of power and the other factors matter, but they are second-order to the availability and timing of power, which is why the analysis treats time-to-power as the master variable.

Figure 8 compares three approaches to solving the power-and-land problem.

Figure 8. Three Approaches to Solving the Power-and-Land Problem

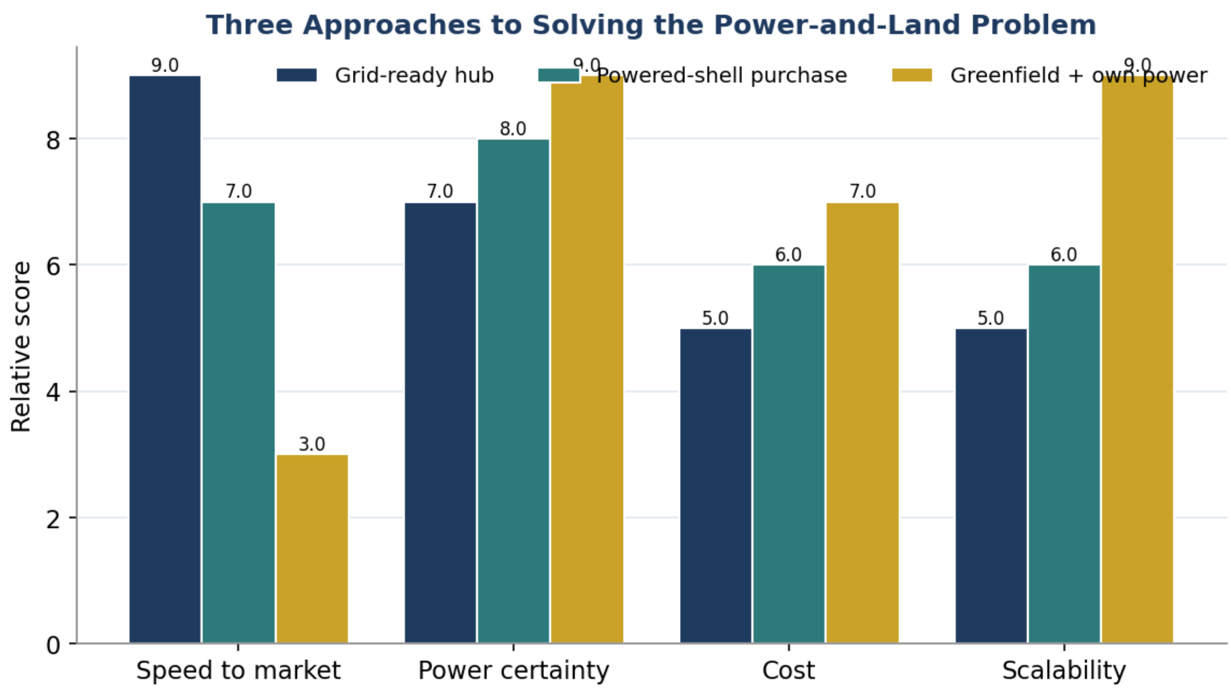

Indicative; each approach trades speed, power certainty, cost and scalability differently.

The comparison sets out the investor's strategic options. Locating in a grid-ready hub offers the fastest route to market and good power certainty but at higher land cost and with limited room to scale. Purchasing a powered shell, a built structure with power provisioned, offers strong power certainty and reasonable speed. Developing greenfield with its own power offers the most scalability and control and potentially the lowest long-run cost, but the slowest delivery and the greatest power-execution risk. The right approach depends on the investor's priorities, speed versus scale, certainty versus cost, and on its capability to execute a power build. For many investors prioritising speed to market in a supply-constrained window, a grid-ready hub or a powered shell captures the time-to-power premium at acceptable cost, while greenfield-with-own-power suits large, capable investors building for scale over a longer horizon. The framework's contribution is to make the trade-offs explicit so the investor chooses deliberately, matching its approach to its priorities and capability.

A second illustration shows the optionality strategy in action. Consider an investor that, rather than developing a single project, assembles a pipeline of secured power allocations and connectable land positions across several Gulf locations ahead of committing to build any one of them. Each secured position is an option: the investor can develop it when demand, offtake and capital align, and can let it lapse or trade it if not. By holding these options on the scarce inputs, the investor can move quickly when a strong offtake appears, developing a powered, connectable site into capacity faster than a competitor who must start securing inputs only after winning the tenant. The pipeline of secured inputs is thus both a source of optionality and a competitive moat, because the scarce inputs the investor has secured are exactly what its competitors lack. This strategy, securing the binding inputs ahead of need and developing into demand, is how the most capable infrastructure investors turn the bottlenecks from an obstacle into an advantage, and it is the fullest expression of the analysis's insight that in a supply-constrained market, owning the scarce inputs is owning the opportunity.

4.7 A Worked Illustration of Investing Around the Bottlenecks

To make the framework concrete, consider an investor choosing between two Gulf sites for a data-centre development. Site A is large and inexpensive but distant from spare grid capacity and fibre, requiring a substantial grid upgrade with a multi-year lead time. Site B is smaller and costlier but adjacent to a substation with spare capacity and on a fibre route, connectable within a year. A naive comparison on land cost favours Site A; the bottleneck framework favours Site B decisively, because its time-to-power is far shorter, letting it deliver capacity into the supply-constrained market years sooner, capturing scarce-market revenue while Site A is still building its power. The investor applying the framework pays the premium for Site B, secures the power allocation and connection agreement early, prices the connection lead time conservatively, and proceeds, while recognising that Site A might suit a different, longer-horizon, larger-scale strategy if its power could be secured. The decision turns not on land cost but on time-to-power, which is the value the framework places at the centre.

The illustration shows the framework overturning the intuitive answer. An investor reasoning from land cost, the figure most visible on a term sheet, would choose Site A and discover, years later, that its power had not arrived and its capital sat idle; an investor reasoning from the bottlenecks chooses Site B, pays more for the land, and delivers revenue while the first investor waits. This is the practical payoff of putting the bottlenecks, not the headline costs, at the centre of the analysis: it directs the investor to the decisions that actually determine the return, which are about power and time, not about the land price that dominates a superficial comparison. The investor that internalises this reads every site and project through the lens of the binding constraints, and so makes the choices that the demand alone would not reveal.

4.8 Bringing the Bottleneck Analysis Together

The elements of the framework, the ranking of bottlenecks, the power solutions and their lead times, the cost of power, the determinants of site suitability, and the value of time-to-power, form a single connected analysis centred on one insight: that the binding physical constraint, almost always power, determines the project, and that securing it early and quickly is where value is created. The ranking tells the investor where to look; the power-solution and site analyses tell it how the constraint can be overcome and at what speed and cost; and the time-to-power and constraint-impact analyses tell it why speed matters and what failure costs. Applied together, they shift the investor's attention decisively from the demand, which it shares with every competitor, to the supply constraints, where it can gain or lose advantage. The recurring failure in this asset class is to invest from the demand, assuming the supply will follow, and the framework's connected structure is designed to prevent exactly that by making the supply constraints the centre of the analysis. An investor that holds power and time-to-power at the centre will consistently make better siting, timing and securing decisions than one drawn by demand and land cost.

5. ROBUSTNESS AND SCENARIO ANALYSIS

Because the framework is stylised, its value lies in whether its conclusions hold under different assumptions. This section varies the central inputs.

5.0a Reading the Scenarios as a Set

The scenarios that follow, longer power lead times, shifting renewable economics and changing land constraints, stress different aspects of the supply picture, but they share a unifying lesson worth stating before they are taken in turn. In every scenario, the value accrues to the investor that secured the binding inputs early and priced the lead times realistically, and the harm falls on the investor that assumed the supply would follow the demand. Whether lead times lengthen, renewable economics shift, or land tightens, the investor with secured power and a connectable site is protected, while the one still seeking its inputs is exposed. This invariance is the practical heart of the analysis: across the plausible futures, the discipline that protects the investor is the same, secure the scarce inputs early, value time-to-power, price lead times conservatively, and the scenarios differ only in how severely they punish its absence. An investor that adopts this discipline is robust across the scenarios; one that does not is exposed in all of them, which is why the analysis treats the discipline rather than any single forecast as the key.

5.1 Longer Power Lead Times

If power lead times prove longer than assumed, as grids strain under rising demand, the value of secured, ready power and the premium on time-to-power both rise, and the disadvantage of relying on a long-lead power solution worsens. The conclusion is reinforced: the scarcer and slower power becomes, the more valuable it is to have secured it, and the more an investor should pay for a grid-ready site or a secured power position. A scenario of lengthening lead times strengthens rather than weakens the analysis's central emphasis on securing power early and valuing time-to-power, because it widens the gap between projects that have power and those still waiting for it.

A connected consideration is the rising power density of AI workloads, which compounds the power bottleneck over time. AI computing, particularly training, packs more power into each rack than traditional workloads, so a data centre serving AI demand draws more power per unit of floor space and strains both the power supply and the cooling that must remove the resulting heat. This rising density means the power required for a given data-centre footprint is growing, intensifying the bottleneck the analysis identifies, and it raises the importance of cooling, which the ranking placed below power but which becomes more binding as density rises. An investor should therefore assess not only the power a project needs today but the trajectory of its power and cooling requirements as workloads densify, favouring sites and designs that can accommodate rising density. The densification of AI computing is a structural force pushing in the direction of tighter power and cooling constraints, which strengthens the analysis's central emphasis on securing ample, firm power and adequate cooling, and warns against sites and designs sized only for today's densities.

5.2 Cheaper or Constrained Renewables

If solar and storage become cheaper and more available, the Gulf's power-cost advantage widens and behind-the-meter renewable solutions become more attractive, easing both the cost and, where they can be built quickly, the availability of power. If, conversely, renewable build is constrained by land, grid-integration or supply-chain limits, the reliance on conventional and grid power persists. Either way, the conclusion holds: the investor must secure firm, affordable power on its timeline, and the mix of sources that achieves this will evolve. The robustness point is that the analysis's emphasis on firmness and timing is source-agnostic; whether the firm power comes from grid, gas, solar-with-storage or a hybrid, what matters is that it is secured, and the cost advantage of the Gulf's renewables is a benefit on top of that requirement.

A final scenario worth considering is the maturation of the market, in which grid capacity expands, connection processes streamline, and the power bottleneck eases over time as utilities and regulators respond to the demand. This is the favourable long-run scenario, and it would gradually reduce the time-to-power premium and the scarcity value of secured power as supply catches up with demand. The investor should welcome this prospect but not assume it in current decisions, because the easing, if it comes, will take years, and the premium on secured power is real in the meantime. The implication mirrors the analysis throughout: the scarcity-driven advantages, the time-to-power premium, the value of secured inputs, are largest now, in the constrained phase, and will erode as the market matures, which is an argument for capturing them through early, well-secured projects rather than waiting for a maturity that will compete the advantages away. The early, disciplined investor captures the scarcity premium; the one who waits for the bottlenecks to ease captures a smaller opportunity in a more competitive market.

5.3 Easing or Worsening Land Constraints

If suitable, well-located land becomes scarcer, the value of sites with power and connectivity proximity rises, and early securing of such sites becomes more valuable; if land constraints ease, the power bottleneck dominates more clearly still, because land ceases to bind. Under either scenario, the analysis's ranking holds, with power the dominant constraint and land valued for its access to power and connectivity. The robustness conclusion is that the investor should secure well-located, power-proximate land early where it is scarce, and concentrate on power where land is ample, in both cases keeping power at the centre of its assessment, because land matters chiefly as the route to power and connectivity rather than as a constraint in its own right.

6. IMPLEMENTATION CONSIDERATIONS

Translating the framework into practice involves securing power first, selecting sites for infrastructure access, pricing lead times realistically, and building optionality on scarce inputs.

A foundational sequencing point, parallel to the other papers in this cluster, is that the power question must be settled before, not after, the capital is committed and the site acquired, because the investor's flexibility is greatest before it has sunk costs into a particular project. Before committing, the investor can choose among sites and power solutions, walk away from a project whose power cannot be secured, and negotiate from a position of choice; after committing to a site and incurring costs, it is anchored to that project and its power problem, with far less flexibility to respond if the power proves hard to secure. The discipline is therefore to resolve the power question, to secure firm power or a clear, evidenced path to it, as a condition of committing, so that the investor never finds itself holding a site it cannot power. This sequencing, power before commitment, is the single most important procedural discipline the framework implies, because it preserves the investor's ability to avoid the worst outcome, a committed, unpowered project, which is the outcome the bottleneck analysis most warns against.

6.1 Securing Power First

The foundational discipline is to secure firm power before committing to a project, treating an unsecured power position as a reason to wait rather than a risk to accept. This means obtaining a contracted, firm power supply on the project's timeline, including for planned expansion, and understanding the lead time and certainty of the connection. An investor that secures power first has overcome the dominant bottleneck and can proceed with confidence; one that commits before securing power exposes itself to the compounding delay, cost and shortfall that power constraints inflict. Power should be the first question and the first secured item, not a matter resolved along the way.

A practical refinement on site selection is to assess not only a site's current infrastructure access but its trajectory, because grid and fibre networks evolve. A site that is currently power-proximate may face rising competition for the same grid capacity as other developments cluster around it, while a site currently less connected may sit on a planned grid expansion or fibre route that will improve its access. The investor should therefore assess the forward path of the infrastructure around a site, the planned grid investments, the committed fibre builds, the other developments competing for the same capacity, rather than only its present state, because the power and connectivity available when the project is built and operating matter more than those available today. This forward assessment requires engagement with the utilities and planners who shape the infrastructure's trajectory, reinforcing the value of the early relationships the analysis recommends, and it guards against the error of selecting a site for a current access that competition or congestion may erode by the time the project needs it.

6.2 Selecting Sites for Infrastructure Access

Because land is valued for its access to power, fibre and water, the investor should select sites primarily for their infrastructure proximity rather than their cost or size. A site adjacent to spare grid capacity and a fibre route, with water for cooling and clear permitting, is worth more for a data centre than a cheaper, larger site lacking these, because it can be powered and connected quickly. The investor should assess each candidate site for its access to the binding inputs and weight that access far above land price, recognising that the apparent saving on a poorly-connected site is illusory if the site cannot be powered on the timeline the project needs.

A practical aid to realistic lead-time pricing is to benchmark a project's assumed timelines against the actual delivered timelines of comparable recent projects, rather than against the project sponsor's plan. Sponsors, subject to the optimism bias the literature documents, tend to present favourable timelines, and an investor relying on them inherits the bias; an investor that instead asks how long comparable connections and builds have actually taken in the relevant jurisdiction grounds its timeline in evidence rather than hope. This reference-class approach, comparing the project to the realised experience of its peers, is a well-established corrective to planning optimism, and it is especially valuable for the power and connection timelines that dominate data-centre delivery. An investor that prices its lead times from the realised experience of comparable projects, with contingency on top, will plan and commit on a sounder basis than one that accepts the sponsor's plan, and it will avoid the disappointment of timelines that systematically run long. Grounding lead-time estimates in realised peer experience is the practical discipline that operationalises the conservatism the analysis recommends.

6.3 Pricing Lead Times Realistically

The investor should price the lead times of its power and connection solutions realistically, and conservatively, because underestimated lead times are the most common cause of delay and value erosion. This means basing the project timeline on the actual, evidenced lead time for the chosen power solution, not on an optimistic assumption, and building contingency for the delays that complex connections often incur. An investor that prices lead times realistically avoids the disappointment of a project that takes years longer than planned; one that assumes optimistic timelines will find its returns eroded by delays it should have anticipated. Realistic lead-time pricing is the discipline that aligns the project's plan with the physical reality of the bottlenecks.

A further implementation discipline is to engage early with the entities that control the scarce inputs, the utilities, grid operators, regulators and landholders, because access to power and connectable land is often allocated through relationships and processes that reward early, credible engagement. An investor that builds relationships with the power authorities and understands their allocation processes and timelines can secure power positions ahead of less-engaged competitors, while one that approaches them late, after committing to a site, takes its place in a queue. In the Gulf, where much power and land is controlled by government and government-linked entities, this early, credible engagement with the controllers of the scarce inputs is particularly valuable, and it is part of the local capability the cluster emphasises. The investor should treat the relationships that unlock power and land as strategic assets to be cultivated before they are needed, not as transactions to be conducted once a project is otherwise set, because the timing and certainty of the scarce inputs often depend on them.

A final implementation point concerns the management of a portfolio of secured but undeveloped positions, because holding options on scarce inputs carries its own discipline. Secured power allocations and land positions typically carry holding costs, option fees, land carry, and sometimes use-it-or-lose-it conditions that require development within a period, so an investor holding a pipeline of positions must manage them actively, developing those that align with demand and capital, and releasing or trading those that do not before their costs or conditions erode their value. The optionality is valuable but not free, and an investor that accumulates secured positions without a plan to develop or release them in time can find the holding costs and lapse conditions eroding the value it secured. Managing the pipeline as a portfolio of options with costs and expiries, developing and releasing deliberately, is the discipline that turns secured inputs into realised value rather than a stranded inventory of carrying costs, and it is the operational counterpart to the strategic insight that owning the scarce inputs is owning the opportunity.

6.4 Building Optionality on Scarce Inputs

Where power and well-located land are scarce, securing options on them ahead of need, a power allocation, a connection agreement, a land position, is itself valuable, because it gives the investor the ability to deliver when others cannot. An investor that builds a pipeline of secured power and land positions holds options on the scarce inputs that the demand will require, and can develop into a supply-constrained market while less prepared competitors are still seeking their inputs. This optionality, securing the binding inputs ahead of need, is a strategy in itself, and it is how the most capable investors turn the bottlenecks from an obstacle into a source of advantage, owning the scarce inputs that the boom makes valuable.

7. CONCLUDING COMMENTS

The Gulf data-centre boom is a demand story with a supply problem, and it is the supply problem, the physical bottlenecks of power, land, cooling and connectivity, that determines which projects succeed and what returns investors earn. This paper has set out the real bottlenecks, shown that power dominates, and explained why time-to-power drives value and how an investor should invest around the constraints.

The evidence and analysis support five conclusions. First, power, in availability and connection speed, is the dominant bottleneck. Second, the value of a project lies substantially in its time-to-power, so a grid-ready site commands a premium. Third, power cost varies by source, and the Gulf's cheap conventional energy and abundant solar are a structural advantage, subject to securing firm, timely supply. Fourth, land is valued for its proximity to power, fibre and water, not for itself. Fifth, the investor's returns depend on securing the binding inputs early and pricing the lead times others underestimate.

Before turning to the limitations, it is worth placing this paper within its cluster. The strategic companion paper establishes that the Gulf data-centre opportunity is large and that power and offtake are its dominant value drivers; the underwriting companion treats power security as a gating item in lending against a specific asset; and this paper explains why power gates, how the constraint behaves, and how an investor should invest around it. Together they take an investor from the strategic case, through the underwriting of an asset, to the physical constraints that determine whether any asset can be built. This paper's contribution is to make the supply constraints, which the others take as given, the explicit object of analysis, so that the investor neither assumes the supply will follow the demand nor underestimates the lead times that ration the opportunity to those who solve them. The bottlenecks are the part of the data-centre story least discussed and most decisive, and this paper aims to give them the central place the returns show they deserve.

These conclusions are subject to limitations that also mark out the avenues for further work. The framework is stylised and illustrative, and actual bottlenecks, lead times and costs vary by location and over time; the value is in the structure, not the specific numbers. The analysis treats the Gulf at a regional level, whereas the power, land and grid situation differs sharply across jurisdictions and sites, which further work could examine in detail. The lead times and costs are calibrations to illustrate the constraint economics rather than current market readings, and an investor should populate them with live, project-specific data. And the analysis treats the physical bottlenecks; the financing, offtake and strategic dimensions are addressed in the companion papers. These are extensions rather than corrections: the central message, that the bottlenecks, not the demand, determine outcomes, that power dominates, and that securing the binding inputs early and valuing time-to-power are the keys to returns, is robust. An investor that secures power first, selects sites for infrastructure access, prices lead times realistically and builds optionality on scarce inputs can invest around the bottlenecks that defeat less careful participants, and so capture the returns the demand implies but the supply constraints ration to those who solve them.

DISCLOSURES

This paper is provided for general information and educational purposes only. It does not constitute investment, financial, legal, tax or other professional advice, nor an offer or solicitation to buy or sell any security or to pursue any investment strategy. The figures and scenarios are stylised, forward-looking illustrations used to convey constraint economics; they are not forecasts, do not represent the past or expected performance of any actual asset or project, and should not be relied upon as such. Data-centre development carries power, construction, land and market risks that can result in loss. The author is an independent researcher; the views are his own and do not represent those of any institution. Readers should obtain professional advice tailored to their circumstances before making any decision.

REFERENCES

References

[1] Dixit, A. K. and Pindyck, R. S. (1994). Investment under Uncertainty. Princeton, NJ: Princeton University Press.

[2] Flyvbjerg, B. (2014). What You Should Know about Megaprojects and Why: An Overview. Project Management Journal, 45(2), 6-19.

[3] International Energy Agency (2024). Electricity 2024: Analysis and Forecast to 2026. Paris: IEA.

[4] International Energy Agency (2025). Energy and AI: Electricity Demand from Data Centres. Paris: IEA.

[5] Koomey, J. and Masanet, E. (2021). Does Not Compute: Avoiding Pitfalls in Assessing the Internet's Energy and Carbon Impacts. Joule, 5(7), 1625-1628.

[6] Masanet, E., Shehabi, A., Lei, N., Smith, S. and Koomey, J. (2020). Recalibrating Global Data Center Energy-Use Estimates. Science, 367(6481), 984-986.

[7] Trigeorgis, L. (1996). Real Options: Managerial Flexibility and Strategy in Resource Allocation. Cambridge, MA: MIT Press.

ABOUT THE AUTHOR

About the Author

Chennakeshav Adya is an independent researcher and corporate-finance practitioner with more than 25 years of industry experience as both an operator and an investment banker. He holds a

Bachelor of Engineering from Visvesvaraya Technological University (VTU), an MBA from London Business School, and is currently pursuing an LLM (Master of Laws) at University College London (UCL). His research focuses on capital formation, private markets and cross-border investment across the Gulf, the United Kingdom, Europe and the United States, with particular attention to real estate, private credit, digital infrastructure and alternative assets. He writes to bridge institutional investment practice and the practical questions facing allocators and capital seekers.

APPENDIX A. BASE-CASE ASSUMPTIONS

The table below records the stylised assumptions used in the framework. They are indicative, forward-looking calibrations chosen to illustrate constraint economics, informed by the infrastructure and energy literature; they are not forecasts and not estimates from a proprietary dataset.

Table A1. Full Base-Case Assumptions

Parameter	Base case	Rationale
Dominant bottleneck	Power	Defining input of a data centre
Grid connection lead time	~12-18 months	Illustrative
Grid upgrade / substation	~24-30 months	Longer-lead solution
New generation	~36-48 months	Longest-lead, highest-capacity
Solar PPA vs grid	~70% (index)	Gulf irradiance advantage
Top site factor	Power proximity	Land valued for input access
Time-to-power	Premium for grid-ready	Value in supply-constrained market
Cooling / water	Secondary constraint	Matters but rarely binds first

Indicative assumptions; see Section 3.

APPENDIX B. POWER-AND-SITE DILIGENCE CHECKLIST

The following checklist distils the framework into the questions an investor should answer when assessing the bottlenecks for a Gulf data-centre project.

1. Have we identified power as the likely binding constraint and assessed it first and most rigorously?Is firm power contracted, or clearly securable, on the project's timeline, including for planned expansion?What power solution does the project rely on, and what is its realistic, evidenced lead time?Have we matched the power solution to the timeline we need, recognising the capacity-versus-lead-time trade-off?Is the power affordable, and have we taken advantage of the Gulf's solar and conventional cost options where firm?Have we selected the site primarily for proximity to power, fibre and water rather than for cost or size? Does the site have clear zoning, permits, water for cooling, and room to expand?Have we priced the lead times realistically and conservatively, with contingency for connection delays? Have we quantified the value of time-to-power and the premium a grid-ready position commands?Have we secured options on scarce power and land inputs ahead of need where we can?Have we matched our siting approach (grid-ready hub, powered shell, or greenfield)

to our priorities and capability?Have we treated an unsecured power position as a reason to wait or decline rather than a risk to accept?

APPENDIX C. GLOSSARY OF KEY TERMS

The following glossary defines the principal technical terms used in this paper for readers who may be less familiar with data-centre infrastructure constraints.

Bottleneck. A binding constraint that limits the delivery of capacity; for data centres, chiefly power, then land, cooling and connectivity.

Firm power. A contracted, reliable supply of electricity available on demand, as opposed to intermittent or anticipated power.

Grid connection. The physical and contractual link between a site and the electricity grid, often a long-lead item.

Power allocation. A commitment of existing grid capacity to a site, the fastest but most capacity-limited power solution.

Substation. Electrical infrastructure that steps and distributes power; building or upgrading one is a medium-lead solution.

New generation. Building new power-generation capacity to serve a load, the highest-capacity but longest-lead solution.

Lead time. The time required to secure and deliver an input such as power; the binding constraint is the longest-lead critical input.

Time-to-power. The time from project start until firm power is available; a key determinant of when capacity can be delivered.

Grid-ready site. A site with secured, ready power that can deliver capacity quickly, commanding a premium in a supply-constrained market.

Greenfield. A site developed from scratch, including building its power, offering scale but the slowest delivery.

Powered shell. A built structure with power and connectivity provisioned, offering power certainty and reasonable speed.

Levelised cost of power. The average cost of power per unit over the life of a source, used to compare sources.

PPA. A power purchase agreement, a contract to buy power, often from a renewable source such as solar, at agreed terms.

Behind-the-meter. Power generated on-site, behind the grid connection, such as on-site gas or solar, serving the load directly.

Irradiance. The intensity of solar energy at a location; the Gulf's high irradiance makes solar especially cost-effective.

Cooling. The systems that remove the heat a data centre generates; a constraint where water or efficiency is limited.

Connectivity. The fibre and network links a data centre requires; proximity to fibre routes is a site-suitability factor.

Site suitability. The fitness of a site for a data centre, determined chiefly by access to power, fibre and water.

Real option. The value of the flexibility to act, such as the option to deliver sooner or to develop a secured input.

Optionality. The value of holding secured rights to scarce inputs ahead of need, enabling delivery when others cannot.

Supply-constrained market. A market where capacity cannot meet demand, so the ability to deliver sooner carries a premium.

Binding constraint. The single factor that most limits a project, which determines its timeline and viability.