Features

Structuring financeable data centres

Data centres sit at the convergence of real estate, energy infrastructure and capital markets. By **JOHN BAIN**, **JOHN FORBUSH**, **DANIELLE GARBIEN**, **ERIC POGUE**, **DALE SMITH** and **WES SMITH**, **WILLKIE FARR & GALLAGHER LLP.**

Data center, server room. Fantastic supercomputer. © Funtap P | Dreamstime.com

Multi-hundred megawatt campuses, hyperscale build-to-suits and networked co-location platforms are competing for power, scarce interconnection capacity and long-lead equipment, complicated further by evolving regulatory frameworks. Whether these projects reach notice to proceed – and how they are financed – centre around the ability of developers to understand and appropriately structure the interface between the data centre and the power and utility assets required for their operation.

This article is organised around four substantive areas: (a) division of real estate and site access and use rights; (b) contracting for power; (c) intercreditor and collateral structuring considerations; and (d) energy regulatory considerations.

Site access and use rights

The foundational elements of a data centre campus consist of a large site and a significant power allocation. Once a site is identified and a power allocation is obtained, the developer will work with design professionals, engineers, architects and the prospective tenants to prepare the initial campus layout and, ultimately, a general site plan. The campus design phase and site plan preparation establish the major components of the campus and their respective locations.

While each campus is unique, data centre campuses are generally composed of three main aspects: individual data centre buildings; electrical infrastructure; and non-electrical shared infrastructure such as roads, water and sewer lines, stormwater facilities, fibre lines and connections, and natural gas pipelines and related improvements. The site plan serves as a conceptual guide for determining the approximate location of each of these major campus elements. The site plan is also a key document in pursuing zoning and permitting approvals required to move the campus through the pre-development process.

- *Parcelisation, the whys and the hows* – Once a site plan is in place, the overall site must be subdivided into individual parcels. Subdividing the site is a common practice that accommodates a phased construction process, financing, and where applicable, satisfies the leasing requirements of hyperscale tenants. The enormous capital required to fund the construction of a single building, which depending on the power allocation, can exceed US$1bn, drives the need for dedicated debt and capital stacks for each such building.

Accordingly, each building may be owned and funded through an individual joint venture and financed by an individual construction loan. This requires each such building to sit on a distinct, separate parcel, allowing for individual fee simple ownership, and the granting of a mortgage in favour of a construction lender which encumbers only the parcel that the subject building is constructed on.

Likewise, in cases where the intended user(s)/tenant(s) for such building require a purchase right, generally a right of first refusal or right of first offer, under their lease, situating the building on a single parcel allows for those rights to be granted solely with respect to the parcel containing the building (and if co-located, the individual data hall) that is covered by their lease.

Harmonising the above requirements with the reality that each building must receive the benefit of shared campus infrastructure is addressed through the subdivision process, together with preparing and implementing a set of agreements that set forth each individual data centre owner's rights and responsibilities with respect to such shared infrastructure. There are multiple common approaches to achieve this outcome: condominium structures, master declarations of covenants, conditions and restrictions, joint facilities agreements, and tenancy in common structures, among others.

Ultimately, each structure solves the same basic set of issues – ensuring that each parcel/building owner has the access and rights to shared infrastructure necessary to operate its building, establishing management, maintenance and governance architecture, and allocating costs among the owners. Larger shared improvements, such as private substations, capacitor banks, drainage ponds etc, are generally located on separate parcels that are owned by an association or similar entity or fractionally owned in a tenancy in common structure. For other shared improvements, such as electrical lines, pipelines, fibre lines and roads, the core agreements – ie, master declarations or condominium declaration – will establish easements that crisscross the site.

• *Easements and the campus nervous system* – Every modern data centre campus depends on an intricate lattice of easements that allows power, fuel, gas laterals, fibre, water intake and discharge lines, thermal loops, stormwater facilities, and private roads to move efficiently across what will eventually become separately owned parcels. The agreements setting forth these rights must be established early, recorded in the real property records in a form that survives foreclosure or sale, and drafted with enough flexibility to accommodate future phases while protecting lenders and operators from interruption.

A common approach is to incorporate detailed easement exhibits directly into the master condominium declaration or the declaration of covenants, conditions and restrictions, supplemented by a standalone master utility easement agreement or a shared facilities agreement when multiple non-affiliated owners are contemplated from the outset. Easements are granted as appurtenant interests that run with the land and expressly bind successors, lenders, and assigns.

Developers now routinely include both fixed metes-and-bounds corridors for known infrastructure and floating or criteria-based expansion easements that trigger only when a new phase meets pre-agreed conditions. Relocation rights are granted to the owner of the servient parcel, but are always subject to a “no material impairment” standard and an obligation to bear all costs.

Fibre conduits receive particular attention: hyperscalers increasingly demand exclusive conduit zones and irrevocable dark-fibre rights of use embedded in the master declaration to prevent a future owner from trenching through an existing carrier’s path. With the rise of district energy and heat-reuse mandates in US jurisdictions such as Loudoun County, Virginia, and the Arizona Department of Water Resources service territory, thermal discharge and intake easements are reserved at the master-planning stage – their value is beginning to rival that of the megawatts themselves. To ensure financeability, master declarations often include non-disturbance and lender recognition provisions so that foreclosure of a servient parcel cannot interrupt power, cooling, or connectivity to the benefitted data centre buildings and parcels.

• *Cost sharing and the economics of shared infrastructure* – The allocation of ongoing operating and capital costs across separately owned parcels is one of the most heavily negotiated areas of the entire campus-governance package. A poorly drafted cost-sharing regime creates perpetual disputes, erodes financeability, and can render an otherwise excellent site effectively uninvestable. The starting point is a comprehensive cost-waterfall schedule incorporated into, or that supplements, the shared facilities agreement, CCR or condominium declaration. Common areas – roads, stormwater ponds, security perimeters, and landscaping – are typically allocated pro rata by platted acreage or percentage interest in the condominium or other applicable regime.

Shared electrical infrastructure – the private substation, switchyard, and medium-voltage backbone – is allocated based on reserved capacity. New phases that trigger upgrades to shared facilities bear the incremental cost under latecomer provisions, preserving the economics for earlier parcels. Water and cooling infrastructure costs are increasingly allocated by peak thermal load or gallons per day to reflect the growing complexity of heat-reuse and district-energy programs.

Property taxes, once any abatement period expires, flow through by assessed value per parcel. Reserve accounts corresponding to an individual borrower/owner’s share of these costs have become a common requirement for lenders. Major decisions – budget approvals and large capital expenditures – are governed by weighted voting thresholds coupled with drag-along rights to prevent minority holdouts from blocking necessary work.

• *Special case, on-site or adjacent hybrid power plants* – When a campus incorporates material on-site or adjacent generation – which is becoming increasingly common for new hyperscale projects, whether gas-fired peakers, reciprocating engines, fuel cells, battery storage, solar-plus-storage, or emerging nuclear small modular reactor interfaces – the power plant must be physically and legally ringfenced from the data centre parcels from the very beginning. The generation assets are placed on one or more separately platted parcels or condominium units, often with independent access roads and security perimeters.

As discussed in more detail below, this separation is non-negotiable for several reasons – including financeability considerations.

Contracting for power

Data centre campuses require large amounts of electrical power. In cases where a contemplated campus development requires more grid-sourced power than is then allocated to the proposed site, the process for acquiring additional power starts with submitting a large-load power allocation request to the local utility. This request initiates a coordinated process between the data centre campus developer and the applicable local utility, where the local utility analyses the size of the request and the characteristics of the proposed use, together with the present capacity of the local grid to accommodate the request.

As part of this process, the utility identifies grid upgrades required to support the proposed campus-owned grid-connected assets, and, where applicable, accommodate on-site generation and storage. These facilities and the agreements governing their ownership, construction and operation must be aligned to ensure grid performance, a coordinated construction and energisation schedule, and reliability of power supply across the campus.

Based on the location of the campus, developers and tenants should have a complete understanding of the structure of the utility provider – ie, a member-owned utility cooperative, investor owned utility etc – local market rules, and tariff requirements, all of which will bear on the economic and operational realities of

Features

executing on a large load request in accordance with the agreed ramp schedule. In markets where power delivery timelines can be unpredictable or where recent power delivery failures have occurred, developers who want to ensure timely performance under their leases should arrange in advance for bridge power solutions that satisfy tenant requirements.

For grid-connected data centre campuses, approval of a large load request often requires developers and/or data centre customers to enter into facility upgrade agreements, line extension agreements, and similar arrangements providing for line extensions to the campus site where applicable, and grid upgrades required for the delivery of the grid power allocated to the campus.

These agreements govern the payment responsibilities between the parties for the upgrade costs, construction timelines and milestones, payment security requirements and change-order protocols. Grid upgrade requirements often include existing substation modifications – such as increasing the primary and secondary transformer windings – and/or newly constructed substations that will ultimately be owned and operated by the utility, electrical switchyards, and capacitor banks to manage the highly variable load profiles associated with AI applications. For upgrades that will be utility owned, the developer is often tasked with acquiring the land upon which these facilities will be located, and conveying or leasing that land to the utility prior to the commencement of construction. On the "customer side" of the substation, revenue-grade utility meters should be installed at the point of interconnection and submeters for tenant pass-through with SCADA, outage notifications and curtailment signals integrated for tariff and regulatory compliance, as applicable.

In addition to securing the upgrades to serve the campus load via a physical path to grid power, developers and hyperscalers routinely enter into long-term electrical supply agreements with the local utility and/or power purchase agreements directly with a generation resource often located behind-the-meter or adjacent to the campus that lock in price, term, curtailment rights, and increasingly carbon-free attributes for a material portion of the campus load.

These agreements often include take-or-pay provisions under which the customer commits to minimum quantity of energy or capacity-reservation charges in exchange for capped energy rates and phased energisation certainty. Regardless of structure, given the need for data centre campuses to have a continuous stable power source, having a power supply contract in place is typically a condition to lease commencement under certain tenant leases, and a precondition to funding under data centre construction loans.

- *Special case, on-site or adjacent hybrid power plants* – Data centre developers are increasingly looking to on-site power options. There is a wide range of possible development models, with true, co-located behind-the-meter generation at one extreme and then a spectrum of hybrid approaches, including adjacent power plants that are interconnected to the grid, have a power purchase agreement with the data centre (virtual or otherwise) and share certain common infrastructure with the data centre campus.

For on-site generation facilities, the power supply contracts are typically between the data centre buyer and the on-site generation or storage owner, often a separate SPV or third-party independent power producer. Depending on the type of on-site resource, the commercial terms for the power and capacity supply arrangements include: availability guarantees, capacity payments, variable energy payments typically fuel-indexed depending on the resource, dispatch rights and responsibilities, rights to curtail the power supply under certain circumstances and the economics associated with such curtailment, liquidated damages for failure to perform, change-in-law and compliance with reliability standards.

If renewables and storage on-site facilities are contemplated, state-of-charge management, round-trip efficiency standards and degradation mechanics and reserves should be properly documented. A private switchyard and medium voltage network are integral parts of balance of plant. Revenue-grade meters are required at the point of delivery to properly measure and document the transfer of energy from the on-site generation owner to the data centre customer.

For hybrid generation facilities, utilities and data centre customers often enter into an interconnection agreement to connect an onsite generation project located on the data centre's campus with the utility's grid. The interconnection agreement typically includes technical specifications for the distributed energy system to ensure it can operate safely and reliably in parallel with the utility's grid. An engineering, procurement and construction contract between the project owner (often a separate SPV or third-party independent power producer) for the development of the generation facility is executed.

The EPC contract typically includes development milestones, design, procurement and commissioning of the project as well as performance testing for electrical balance-of-plant, validating grid and BTM dispatch modes. Separately, the on-site generation owner enters into power and capacity supply contracts with the data centre customer and any ancillary dispatch or control agreements that define dispatch priority, curtailment restrictions and compliance with applicable local tariff and permit requirements. The metering relied upon for the onsite generation resource must be able to distinguish between energy derived from the BTM resource and the grid.

Intercreditor and collateral structuring

Data centre developments weave together a complex set of ownership models including with respect to real property rights, as discussed above, and commercial contracts, including for power. In the context of structuring equity and financing transactions, this creates a unique set of challenges. Construction lenders, equipment lenders and permanent lenders generally overlap and have competing positions and need to understand the bundle of rights required to operate the asset and enterprise into which they are extending capital. The collateral framework must address lien priority, step-in and other enforcement rights if certain defaults occur, pre and post-completion, and proceeds waterfalls.

Investors face a similar challenge in sorting through what rights are shared with other owners and owner groups in order to structure their investment so as to not assume liability or obligations (such as one hundred percent of maintenance costs for shared facilities) with respect to assets outside of its joint venture.

In a grid-connected model, analogous to any industrial end-user of electricity, eg, large manufacturing plants, data centres procure power from local load serving entities. In this model, which remains by far the most common structure for data centre development, the intercreditor and collateral structuring considerations arise between the co-owners of portions of the data centre campus. As discussed elsewhere in this article, discipline with respect to the range of development matters – beginning with the site plan and through and including commercial contract rights (eg sharing interconnection access) – is critical to maintain flexibility for different financing parties and structures.

- *Special Case, on-site or adjacent hybrid power plants* – As discussed above, data centre developers are increasingly looking to on-site and hybrid power options. In an on-site or adjacent hybrid structure, investors and financing parties expect the power plant and related assets, real property and contracts, to be physically and legally ringfenced from the data centre. Developers are best positioned if this structural separation is put in place from the outset for a number of reasons including:

i) Project finance lenders and tax-equity investors, on the power side where applicable, require clear title segregation to structure and protect their collateral including separateness/bankruptcy remoteness considerations. In the context of assets eligible for tax credits, including solar, storage, nuclear and fuel cells, extra scrutiny occurs from the perspective of the tax equity investors, including tax credit buyers, related to tax structuring matters – including credit eligibility and, in the case of investment tax credits, commencement of construction and recapture considerations.

ii) Commercial and regulatory considerations related to the ownership and operation of power projects and the sale of power related thereto may preclude data centre owners from owning the generation assets or create unnecessary regulatory burdens that effectively lead to the same result.

iii) Separating the respective assets allows each technology to be financed with the capital stack best suited to its risk profile and credit life. The funding for developing the generation assets generally consists of dedicated debt and equity solely related to the generation project – including in the case of solar, storage, fuel cells and nuclear, a tax equity investment component. This portion of the capital stack and the market participants related thereto, in many instances is completely separate from the funding for the data centre campus and individual data centre buildings.

Our experience over the last two years has shown that the failure to separate the capital stack for funding the power generation assets at the outset is the single most common reason hybrid campuses fail to attract investment, miss energisation schedules, or become unfinanceable, including with respect to tax credit eligibility matters.

Energy regulatory and commodity issues

Navigating the complex energy regulatory considerations, whether it's grid interconnection, power procurement, or deciding between grid interconnection or co-located power generation, is integral to the data centre financing and development process. Over the last several years, state and federal regulators have devoted increasing attention to the impact data centres and other large load users can have on overall grid stability and reliability.

In the US, at the state level, it is clear that with the large amount of power that data centres require, state utility regulators are overhauling regulatory rules and utility tariffs governing large load interconnection. A key focus of those changes has been cost allocation, requiring data centre customers to bear the cost that large use of power has on the overall grid structure. This focus on grid modernisation will continue as utilities and regulators work to strike a balance between existing customers' use of the existing grid and the improvements necessary to ensure reliable use by data centres.

At the federal level, the US Federal Energy Regulatory Commission has oversight over reliability of the US bulk electric system. However, while FERC has not historically regulated end-use consumers outside of areas like demand response, FERC is tasked by statute to ensure open, non-discriminatory access to the power grid. Now, because of the impact data centres have on overall grid reliability, FERC will inevitably seek to shape regulatory policy over data centre grid usage and interconnection.

This policy is likely to play out in either of two ways: one, with FERC shaping the rules governing use of the regional transmission systems operated by grid operators such as PJM in the Mid-Atlantic or CAISO in California, and second, by attempting to directly regulate grid interconnection by large load users, including data centres, and the cost allocations of grid modernisation efforts that transmission owners are required to expend.

Ultimately, how this continues to play out will be one of the most closely watched areas among industry stakeholders. In terms of data centre financing, there is hope that establishing a better framework of regulations will provide a greater level of consistency. As power and energy usage continues to be the driving issue in data centre development, there is no doubt that energy regulators at both the state and federal level will continue to play an important role. It is imperative that anyone seeking to develop or finance a data centre have their collective ear to the ground to understand how regulatory policy and changes will impact outcomes.

Outlook

Data centre campuses are difficult projects to develop – there are many stakeholders with competing interests and key resources (power) are constrained. For developers to create successful projects that attract capital and top-of-the-market users, it is imperative that they understand and appropriately structure projects from the outset – with rigor and discipline with respect to real property planning beginning with a tailored design phase and site plan, commercial arrangements for shared resources and assets, and power matters.

* This article neither contains legal advice nor establishes an attorney-client relationship in any form. The opinions expressed herein are attributable to the author(s) alone, and they do not reflect the views, positions or opinions of Willkie Farr & Gallagher LLP or other attorneys at the firm. ■