India Data Centre Energy SLAs 2026: Availability, Curtailment and BESS Risk

By Sudarshan Karweer · 05 September 2026 · growthifye.com

A 2026 practitioner guide to data centre energy SLAs in India: uptime, curtailment, BESS, OA risk, penalties and contract design.

India's data centre market is scaling fast, but energy contracting is still catching up to uptime reality. Most power discussions in the sector still focus on landed tariff, open access savings, renewable share, and whether battery storage can shave peak demand or support outage ride-through. Those are necessary questions, but they are not sufficient for a facility that sells uptime.

For a hyperscale or colocation data centre, the real commercial issue is not just what power costs per kWh. It is what the power service actually guarantees, under which failure conditions, across which node of the system, with what penalties, and who bears the cost when the grid, open access supply, battery system, or plant-side controls do not perform. In 2026, that means energy SLAs have moved to the centre of project design, lender diligence, and utility negotiation.

This article looks at how Indian data centre operators, developers, utilities, renewable suppliers and lenders should structure energy service-level agreements for high-availability facilities. The focus is on practical clauses, typical thresholds, Indian regulatory realities, and the commercial treatment of curtailment, outage duration, BESS availability, and scheduling risk.

Why energy SLAs matter more in data centres than in other C&I loads

A conventional industrial consumer can often tolerate process interruptions, reschedule production, or absorb some variability if the delivered energy cost is low enough. A data centre is different. Its economic model depends on service continuity, rack utilisation, contractual uptime to tenants, and reputation. A few minutes of energy service failure can trigger losses far above the monthly savings achieved through cheaper open access procurement.

That changes procurement priorities in three ways:

- The buyer must define reliability at multiple layers, not just at the delivery meter.
- Contracted renewable or RTC supply must be evaluated against actual hourly critical load coverage, not annual energy volume.
- Battery systems, EMS platforms, and grid redundancy must be contracted with measurable performance guarantees.

In India, this challenge is amplified by state-level variations in distribution network quality, substation congestion, outage restoration practices, open access approval timelines, and curtailment behaviour. A 33 kV connected facility in one state may see stable supply with monthly interruption duration under 30 minutes, while another site with a similar sanctioned load may face feeder trips, bay constraints, or switching-related outages that materially change backup design economics.

For that reason, a bankable data centre energy SLA should not be treated as a generic annexure to a PPA or utility supply agreement. It should be built as a structured risk document alongside [Load & reliability engineering](/coreservices/datacentersenergymanagement/loadandreliabilityengineering) and overall power architecture planning.

The five reliability layers that should appear in every SLA

Many contracts describe availability too loosely. They refer to “best efforts”, “continuous supply”, or “firm power” without defining where the obligation begins and ends. For data centre applications, the SLA should separate at least five reliability layers.

- Grid connection availability: whether the utility or licensed distributor maintains energisation up to the agreed delivery point, including bay, line, transformer and switching infrastructure.
- Open access delivery availability: whether contracted third-party power can physically and commercially flow through the approved network path.
- Renewable/hybrid generation availability: whether the supply portfolio produces contracted power according to the schedule or band.
- BESS availability and performance: whether the battery is online, charged, dispatchable, and able to meet duration and response-time guarantees.
- Site-level power continuity: whether the combined architecture supports critical IT and mechanical loads through transfer events, disturbances, and outages.

Each layer should have its own KPI. Otherwise, suppliers tend to aggregate responsibilities and push non-performance into force majeure or “grid event” buckets.

A practical KPI stack for a 25 MW to 100 MW data centre campus in 2026 could include:

- Utility delivery point availability: 99.90% to 99.95% monthly
- Feeder restoration time for non-force-majeure faults: event-wise thresholds such as 30 minutes, 2 hours and 6 hours linked to compensation slabs
- OA schedule fulfilment ratio: 90% to 98% depending on product type and balancing support
- RTC supply hourly coverage band: for example 85% to 95% of nominated hourly volume excluding agreed outage windows
- BESS technical availability: 97% to 99%
- BESS guaranteed discharge response time: sub-second to 2 seconds depending on UPS and switchgear interface
- BESS usable round-trip efficiency: often 85% to 90% AC block level for lithium-ion systems at rated operating conditions
- BESS guaranteed end-of-year usable capacity: commonly linked to degradation curve, for example 85% to 90% of day-one contractual capacity by year 10 depending on cycling regime

These KPIs should be tied to test methods, telemetry source, event classification, and settlement formulas.

Curtailment is the most misunderstood risk in RTC and 24/7 CFE structures

In India’s 2026 market, curtailment remains one of the least well-allocated risks in data centre renewable supply deals. The issue is not only whether wind-solar-hybrid generation is curtailed. The issue is who absorbs the consequences when curtailment disrupts an hourly power shape expected to support a critical load.

There are at least four different curtailment types that should be separately defined:

- Grid security curtailment at generator evacuation level
- State load despatch or scheduling curtailment affecting OA flow

- Distribution-level restrictions or outage events at receiving end
- Commercial curtailment through scheduling under-declaration, deviation control, or banking restrictions

A typical mistake is to write a single sentence saying curtailment by transmission utility or SLDC is a pass-through force majeure event. That may be acceptable for a standard C&I renewable PPA, but it is weak risk allocation for a data centre using the contract as part of a wider uptime strategy.

Instead, the contract should specify:

- Whether curtailment hours count as supplier non-performance for SLA purposes
- Whether substitute power must be arranged, and at what price cap
- Whether curtailment relief is limited if alternative evacuation routes or balancing assets were reasonably available
- Whether BESS held by the supplier must be dispatched first before declaring shortfall
- Whether repeated curtailment beyond an annual threshold triggers portfolio rebalancing or pricing reset

For example, an RTC supply agreement for a 40 MW average data centre load may allow up to 3% annual curtailment without liquidated damages if such curtailment is certified by the relevant load despatch authority and if the seller uses commercially reasonable efforts to mitigate through alternate scheduling and storage dispatch. Beyond that threshold, the seller may be required to compensate the buyer for replacement power cost above contract tariff, subject to a cap.

In 2026, replacement power cost during evening peak periods in several states can easily rise above Rs 8-11/kWh depending on the source, while contracted hybrid RTC tariffs may sit in the Rs 5.2-6.8/kWh range ex-delivery point before local network charges. That spread makes curtailment treatment a material credit issue.

How to contract BESS for reliability, not just arbitrage

Battery storage in Indian data centre discussions is often framed around diesel displacement, peak shaving, or demand-response value. Those use cases are relevant, but for SLA design, BESS should be contracted as a reliability asset with measurable obligations.

The key mistake buyers make is to procure a battery based on MWh nameplate and broad cycle assumptions without locking in delivered-service metrics. For critical loads, what matters is not only installed capacity but dispatch certainty under the site’s actual duty cycle.

A robust BESS schedule should define:

- Contracted usable energy in MWh at AC output
- Continuous discharge power in MW
- Duration at guaranteed power
- State-of-charge reserve policy for contingency use
- Recharge obligation after dispatch event
- Auxiliary load treatment
- Ambient and HVAC operating envelope
- Degradation warranty under specified cycle profile
- Black-start or islanding support obligations, where applicable
- EMS integration and telemetry availability

Suppose a 30 MW IT-plus-cooling critical block requires 10 minutes of bridge support for transfer stability and an additional 20 MW of managed support for one hour during feeder restoration uncertainty. The battery SLA cannot simply state “30 MW/30 MWh BESS installed”. It must state the operational hierarchy: what SoC is reserved for emergency use, what share is available for tariff arbitrage, how quickly it can re-arm after an event, and what happens if the system is partially unavailable.

Typical 2026 battery commercial terms in India for this class of application increasingly include:

- Monthly technical availability guarantees above 98%
- Performance LDs for failure to meet contracted MW or MWh on dispatch instruction
- Separate exclusions for utility outage duration exceeding the design basis if recharge was impossible
- Cybersecurity and EMS uptime clauses because dispatch visibility now affects SLA compliance
- Degradation true-up formulas when actual throughput materially differs from the base case

This is where [On-site generation & BESS](/coreservices/datacentersenergymanagement/onsitegenerationandbess) and Energy management systems need to be considered together. A battery that is technically healthy but poorly integrated with site EMS, transfer logic, or open access schedule management can still fail the commercial objective.

Utility, OA and backup interface: where most disputes actually occur

Most severe disputes do not arise because a single generation asset underperforms. They arise because contract boundaries between utility supply, open access supply, backup power systems and internal distribution are badly drafted.

Consider three common 2026 scenarios:

- The utility feeder trips, the site transfers to backup architecture, but the seller argues its OA supply was still available at injection and therefore no energy SLA breach occurred.
- The renewable supplier under-schedules due to forecast error, the battery partially covers the gap, and the buyer claims a critical reserve was consumed, increasing site risk for the rest of the day.
- A planned utility shutdown is notified late, forcing diesel or market purchase, but the distribution licensee considers the outage compliant with maintenance norms.

To avoid these outcomes, the SLA and supporting agreements should answer six operational questions explicitly:

- What is the contractual delivery point for each source of power?
- Which events are measured at generator bus, interconnection point, utility meter, data centre incomer, and critical load bus?
- Who has dispatch priority among utility supply, OA renewable, battery and backup generation?
- What is the outage notification protocol and minimum notice period?
- What telemetry is the single source of truth for settlement?
- How are overlapping failures allocated?

For high-availability campuses, a power responsibility matrix is often more useful than prose alone. It should map every asset and event class to the responsible party, response expectation, and settlement consequence.

For instance:

- Utility responsible for bay and feeder outage restoration up to delivery meter
- OA supplier responsible for scheduled energy shape and replacement support under defined shortfall classes
- BESS operator responsible for emergency reserve maintenance and dispatch execution
- Facility operator responsible for internal switching, UPS coordination and critical load shedding logic if defined

This level of drafting materially improves lender confidence because it reduces ambiguity in downside cases.

What lenders and investors now look for in 2026 energy diligence

Lenders financing data centre platforms, captive plants, hybrid RTC portfolios or storage-backed supply structures are no longer satisfied with generic assumptions on “reliable power availability”. They increasingly test energy SLAs as part of downside resilience.

The most common diligence questions now include:

- Is the sanctioned load and ultimate build-out load aligned with the contracted energy architecture?
- Are utility connection milestones and bay readiness realistic relative to COD?
- Are open access approvals state-feasible within the project timeline?
- Do SLA penalties cover actual replacement cost exposure, or are they symbolic?
- Is BESS augmentation capex built into long-term availability assumptions?
- Are force majeure definitions too broad for curtailment-heavy states or congested corridors?
- Is there a clear step-in framework if a supplier repeatedly fails hourly delivery thresholds?

For data centre tenants and colocation customers, downstream pass-through of energy SLA quality is also becoming more important. Operators seeking premium enterprise customers will increasingly need auditable evidence that their energy supply chain has service guarantees consistent with the uptime commitments they market.

This is one reason [24/7 clean power contracting](/coreservices/datacentersenergymanagement/247cleanpowercontracting) is moving beyond annual REC-style narratives into hourly operational risk management. The commercial market is rewarding contracts that can explain exactly how a site remains supplied during a bad monsoon week, an evening peak shortage, a bay outage, or a battery derating event.

A practical drafting checklist for Indian data centre energy SLAs

Below is a compact checklist for sponsors, operators, utilities, developers and advisors drafting or reviewing 2026-era data centre energy service agreements in India.

- Define service objective separately for cost, carbon and reliability
- State delivery point and metering point for every source
- Break out availability metrics by grid, OA, generation and BESS layers
- Define curtailment classes and mitigation obligations
- Cap force majeure relief where alternative balancing options exist
- Specify replacement power methodology and compensation formula
- Include hourly or 15-minute telemetry hierarchy for settlement

- Define outage notice requirements for planned and unplanned events
- Set BESS reserve policy and dispatch priority logic
- Link EMS data retention and cybersecurity requirements to SLA validity
- Provide repeated-failure cure rights and supplier step-in provisions
- Align SLA penalties with actual business interruption exposure where feasible
- Ensure consistency across PPA, connection agreement, O&M contract and financing documents

As a rule of thumb, if the SLA can be read without understanding the facility's single-line diagram, feeder configuration, backup architecture and dispatch logic, it is probably too generic for a serious data centre application.

The commercial takeaway for 2026

The Indian data centre market has now reached the point where energy reliability cannot be outsourced to assumptions. Tariff arbitrage, renewable sourcing and storage deployment only create value when the contractual structure reflects how the site actually runs under stress.

That means the best energy strategy in 2026 is not the one with the lowest quoted blended tariff. It is the one that converts grid dependency, renewable variability, curtailment exposure and battery limitations into a measurable, financeable and enforceable SLA framework.

For developers, this improves bankability. For operators, it protects uptime economics. For lenders, it clarifies downside risk. For utilities and policymakers, it provides a more realistic template for serving one of India's fastest-growing critical infrastructure sectors.

If your team is planning a new data centre campus, expanding grid-connected capacity, or restructuring renewable and storage supply around uptime-grade service levels, contact Growthifye's advisory desk. We help clients design bankable energy architectures, commercial risk allocation and execution-ready contracts for India's evolving data centre power market.

Explore Growthifye's related capabilities

This analysis connects directly to our advisory practice: [Load & reliability engineering](/coreservices/datacentersenergymanagement/loadandreliabilityengineering) · [24/7 clean power contracting](/coreservices/datacentersenergymanagement/247cleanpowercontracting) · [Grid connectivity & redundancy](/coreservices/datacentersenergymanagement/gridconnectivityandredundancy) · [On-site generation & BESS](/coreservices/datacentersenergymanagement/onsitegenerationandbess).

ABOUT THE AUTHOR

Sudarshan Karweer — Chief Executive Officer, Growthifye

Over 23 years in management consulting — concept to scale — building and scaling new-age digital and energy businesses. An EY alumnus, Sudarshan leads Growthifye's renewable energy & storage advisory, EPC, transmission engineering and green-financing practices for developers, utilities and C&I consumers across India.
Contact: info@growthifye.com · +91 84510 99371 · growthifye.com