HyperNext Data Centers

RESEARCH PAPER

HN-RP-008

The Water Math

Putting data centre water use in context for India

Where data centres actually sit on the water-consumption hierarchy, and why the conversation needs proportionality.

This is what is Next.

Series	HyperNext Research
Paper	HN-RP-008
Issued	27 May 2026
Version	1.0
Classification	Public release
Citation	HyperNext Research, HN-RP-008

The Water Math

This paper is part of the HyperNext Research series. Methodology, assumptions, and source data are stated openly so other operators can reproduce the analysis on their own facilities. Citation as "HyperNext Research, HN-RP-008" is welcome.

Contents

1. Why we have to write this paper

ABSTRACT

The public conversation about data centre water consumption in India has drifted some distance from the published numbers. Hyperscale data centres are routinely described as outsized consumers of water, with comparisons drawn that lack scale calibration. This paper compiles the actual numbers. A 64 MW HyperNext campus with closed-glycol direct-to-chip cooling and dry-cooler heat rejection consumes approximately 70,000 cubic metres of direct on-site water annually (530,000 cubic metres including embedded water in electricity), comparable to a 200-room luxury hotel, less than half what a single 18-hole golf course in arid Indian geography consumes, and about three percent of what a thermal power station of equivalent generating capacity withdraws. We compare data centre water consumption to other urban and industrial uses of comparable economic and infrastructure scale. We do not include agriculture in the comparison because agriculture serves a distinct social function and the comparison would mislead in both directions. The argument the paper makes is for proportionality in the debate, not for an exemption from scrutiny. Data centres should be measured, reported, and held to standards. The standards should also be calibrated to the actual scale of what is being measured.

The water debate around data centres has become a proxy for the broader anxiety about AI, technology, and infrastructure scale. That is understandable but it is not the same as engaging with the numbers.

Where the conversation went

The framing has converged on three claims that have moved from technical reports into mainstream discourse with little examination of their basis. First, that data centres are massive water consumers per facility. Second, that AI specifically has accelerated the per-megawatt water intensity. Third, that the trajectory is unsustainable for water-stressed geographies like India.

The first claim depends on what one is comparing to. Yes, data centres are large water consumers in absolute terms, but so is any major industrial or commercial facility of comparable scale. The honest version of the claim is "data centres consume large amounts of water in absolute terms, like other facilities of comparable scale, and the water consumption should be reported." That version does not generate a headline.

The second claim is partially true and partially false. AI workloads do increase the per-megawatt cooling load relative to traditional enterprise computing, and the water consumption of conventional evaporative cooling scales with cooling load. But the architectures being deployed for AI by serious operators (HyperNext among them) move toward dry cooling primary precisely because the water economics of evaporative cooling at AI density are not viable. The trajectory of AI-era data centres is toward less water per megawatt, not more.

The third claim is the one this paper is least willing to let stand uncorrected. Water stress in India is genuine and acute. The contribution of data centres to that stress is small in absolute terms and shrinking with each generation of architecture. The dominant contributors to Indian water stress are well documented and have been so for thirty years. Treating data centres as one of those dominant contributors confuses the policy conversation and delays the interventions that would actually move the needle.

What this paper is and is not

This paper is a comparative analysis. It collects published water consumption data from other urban, commercial, and industrial uses and places hyperscale data centre consumption alongside them. It quotes the methodology and sources for each comparison so other operators, researchers, and policy stakeholders can run the same analysis.

This paper is not an argument that data centres should not be measured. Every HyperNext facility reports WUE quarterly under ISO/IEC 30134-9. The methodology is public. The numbers are public. This paper is not an argument that the industry should be exempted from water reporting or watershed accounting. The Nagmati Programme described in HN-RP-001 is the opposite position.

This paper does not include agriculture in the comparison set. Indian agriculture consumes approximately 78 percent of all freshwater withdrawal in the country. Comparing data centre consumption to agriculture would technically support the proportionality argument, but it would also be a category error. Agriculture serves a different social function with different policy considerations and including it in the comparison would invite (correctly) the response that food production is non-negotiable. The comparison is more useful when restricted to other industrial and urban water uses that compete with data centres for the same water resources.

2. What a hyperscale AI data centre actually consumes

Before the comparison, the actual numbers for a HyperNext-class campus. The numbers below are the engineering basis for the HyperNext Hyderabad Phase 1 facility commissioning in December 2026, and the planned Kakinada AI Factory ramping from Q1 2028. They reflect the HyperNext cooling architecture: closed propylene glycol loops carrying heat from direct-to-chip cold plates to external dry coolers, rejecting the heat to ambient air without evaporation. No cooling towers. No water chillers. No evaporative supplement. The only on-site water uses are humidification, sanitation, and facility services.

Phase 1 Hyderabad: 64 MW IT load

Layer	Component	Annual water (m^3)	Notes
1	Cooling water (primary glycol loop → dry coolers)	0	Closed loop, sealed. Topped up only on maintenance event, negligible annually.
1	Humidification and process	30,000	White-space humidity control; substantially lower than evap architectures because there is no cooling-tower drift to compensate for.
1	Sanitation and staff facilities	25,000	Approximately 500 operations staff plus visitor traffic.
1	Other operational	15,000	Cleaning, landscaping (drought-tolerant native species), fire-suppression system tests.
2	Embedded in grid electricity	290,000	15% grid mix at REF target 85%. CEA specific water consumption basis.
2	Embedded in captive renewable electricity	20,000	700 MW Khawada solar plus partnership wind. Mostly panel cleaning, minimal operational consumption.
3	Watershed impact (construction amortised + runoff loss + recharge interruption)	150,000	Modelled per Nagmati methodology in HN-RP-001.
Total gross water footprint		**530,000**	Per HN-RP-001 WUE-3 worksheet, updated for the dry-cooler architecture.
Nagmati restoration credit (2:1)		**-1,060,000**	Watershed projects in Kutch district.

Net water contribution (negative = positive)	**-530,000**	Facility is water-positive on the published target.

On the operationally relevant comparison (gross consumption only), Phase 1 Hyderabad at 64 MW IT load consumes approximately 530,000 cubic metres of water per year across all three layers, of which only 70,000 cubic metres is direct on-site water (Layer 1). That direct on-site number is roughly an order of magnitude lower than a conventional evaporative-cooled facility at the same IT load, which would consume between 700,000 and 1,100,000 cubic metres in cooling water alone.

The 70,000 cubic metre direct on-site footprint corresponds to a WUE well below 0.05 litres per kWh of IT load. The HyperNext published target is WUE below 0.05 L/kWh; the engineered design achieves close to zero on the cooling loop itself, with the small residual driven by humidification, sanitation, and operational maintenance.

Kakinada AI Factory: 1.2 GW IT load (forecast)

The Kakinada facility scales the same architecture. Closed propylene glycol primary and secondary loops, direct-to-chip cold plates on the rack side, dry coolers on the heat rejection side. No evaporative cooling at any ambient temperature. Watershed restoration commitment scaled to consumption per the Nagmati methodology.

Component	Annual water (m^3)	Notes
Layer 1 (direct site water: humidification, sanitation, facility services)	1,300,000	Scales linearly with facility load and headcount. Cooling water remains zero.
Layer 2 (embedded in electricity)	5,800,000	Same 85% REF target, larger absolute load.
Layer 3 (watershed impact)	2,800,000	Larger construction footprint.
Total gross water footprint	9,900,000	0.99 crore m^3 per year.
Nagmati restoration credit (2:1)	-19,800,000	Watershed projects scaled to consumption.
Net water contribution	-9,900,000	Facility is water-positive.

The Kakinada facility at 1.2 GW IT load consumes approximately 9.9 million cubic metres of water gross per year on the HyperNext architecture. For a comparison framework, a conventional evaporative-cooling data centre at the same IT load would consume approximately 45 to 55 million cubic metres per year. The architectural choice (glycol primary loops to dry coolers, no evaporative cooling at any stage) saves roughly 35 to 45 million cubic metres per year on a single 1.2 GW campus, before any watershed restoration commitment.

The dominant arithmetic point: HyperNext architecture at the largest planned scale (Kakinada 1.2 GW) consumes 9.9 million m^3/y gross. That is the engineering baseline for the comparisons in Sections 3 and 4. We treat it as the headline figure throughout the paper.

3. Comparable urban water uses

Indian urban water consumption is large in aggregate and distributed across categories that draw from the same municipal and groundwater sources that data centres do. The comparisons below are facility-by-facility or use-by-use at scales that align with the discussion. Sources are footnoted where the underlying data is from regulatory filings or industry reports rather than HyperNext modelling.

Per facility comparison

Facility type	Typical Indian example scale	Annual water (m^3/y)	Source basis
HyperNext Hyderabad Phase 1 (data centre)	64 MW IT load	1,044,000	HyperNext WUE-3 worksheet (HN-RP-001)
HyperNext Kakinada (data centre)	1.2 GW IT load	16,750,000	HyperNext forecast
Conventional evaporative-cooled DC, same scale	1.2 GW IT load	45,000,000 to 55,000,000	Industry average WUE 1.8 L/kWh
18-hole golf course (semi-arid Indian geography)	60 hectares	150,000 to 250,000 per course	Indian Golf Union water audit data
5-star hotel (200 rooms, India tier 1 city)	200 rooms + amenities	120,000 to 280,000	Indian Hotels Association sustainability reports
Large urban shopping mall (1 million sq ft)	Footfall 50,000/day	180,000 to 320,000	Commercial real estate sustainability filings
500-bed multi-speciality hospital	Tier 1 city, full ICU complement	200,000 to 450,000	National Accreditation Board for Hospitals data
Office complex (1 million sq ft, 8,000 workers)	Tier 1 city	140,000 to 220,000	Building services industry benchmarks
Indoor water park	Large attraction, multiple slides	500,000 to 1,200,000	Indoor amusement industry data
Olympic swimming complex (training facility)	50m pool + 25m pools + dive	30,000 to 50,000 (recurring evap and exchange)	Federation Internationale de Natation standards
University campus (residential, 30,000 students)	Public university tier	1,200,000 to 1,800,000	UGC infrastructure audits

Power station (thermal, 1.2 GW)	Coal-fired, modern	24,000,000 to 38,000,000	Central Electricity Authority

Read the table

A 1.2 GW HyperNext data centre at 16.7 million cubic metres per year sits below a coal-fired 1.2 GW thermal power plant of equivalent generating capacity (24 to 38 million cubic metres per year). The thermal plant is consuming more water to produce the same electricity that a renewable PPA portfolio would produce at a fraction of the embedded water. The water that data centres displace by procuring renewable electricity is part of the water saving, not part of the water cost.

A 64 MW Phase 1 HyperNext data centre at 1.04 million cubic metres per year sits between a large university campus (1.2 to 1.8 million) and a large urban shopping mall (0.18 to 0.32 million). It consumes more water than three to four 18-hole golf courses combined, less water than one indoor water park, and roughly the same as a moderately-sized residential university.

A 1.2 GW HyperNext data centre at 16.7 million cubic metres per year is roughly equivalent to 10 to 14 large university campuses, or 50 to 90 18-hole golf courses, or 60 to 90 multi-speciality hospitals. It serves a customer base of (depending on workload mix) tens to hundreds of millions of users.

The user-served basis

On a per-user-served basis the comparison reframes substantially. A 1.2 GW HyperNext data centre supports inference for approximately 50 to 200 million concurrent users on currently anticipated workload patterns. At the upper end of that range, the water consumption per user-year is approximately 84 litres. The Indian per-capita domestic water consumption average is approximately 49 cubic metres per person per year, or 49,000 litres. The data centre water cost of being a heavy AI user is approximately 0.17 percent of one\'s own domestic water consumption.

That is, if every Indian were a heavy AI user served by HyperNext-class capacity, the total water consumption for that AI infrastructure would be approximately 84 litres per person per year. The total Indian domestic water consumption is approximately 49,000 litres per person per year. AI inference is 0.17 percent.

4. Other industrial water users

Industrial water consumption in India is substantially larger than commercial and urban water consumption. The categories below are sized at the scale relevant for comparison with hyperscale data centres. Indian industrial water is dominated by power generation, with refineries, chemical processing, steel, paper, and textile manufacturing as significant categories.

Industrial comparators at hyperscale

Industry	Typical Indian facility scale	Annual water (million m^3/y)	Notes
Coal-fired thermal power plant	1.2 GW, modern subcritical	24 to 38	Per CEA specific water consumption data, 3.0 to 4.5 $m^3/MWh \times 8760 \times CF$
Coal-fired thermal power plant	1.2 GW, older subcritical (pre-2010)	32 to 45	Pre-modernisation water intensities
Refinery (large)	15 MMTPA crude throughput	20 to 30	Indian refinery sustainability filings
Steel plant (integrated, mid-size)	5 MTPA crude steel	18 to 35	Indian Steel Authority data, varies with technology vintage
Paper mill (kraft, integrated)	500,000 tonnes/year	15 to 30	Indian paper industry data
Textile mill (vertically integrated)	50,000 tonnes/year fabric	6 to 15	Indian textile industry water audits
Bottling plant (large beverage)	100 million litres/year output	1.5 to 3.5	Beverage industry water audits
Brewery (medium)	200 million litres/year beer	1.0 to 2.0	Brewing industry standard ratios
Dairy (large processing)	500,000 litres/day milk processed	0.4 to 0.8	Indian dairy industry data
Pharma manufacturing (mid-size)	API plus formulations	0.3 to 0.9	Indian pharmaceutical industry filings
Cement plant (large)	5 MTPA	0.6 to 1.2	Indian cement industry data, dry-process modern

HyperNext Kakinada (data centre)	1.2 GW IT load	16.7	HyperNext architectural forecast
Conventional DC of same scale	1.2 GW IT load	45 to 55	Industry average WUE 1.8 L/kWh

The picture this paints

On the per-facility basis, hyperscale data centres of the HyperNext architecture consume less water than refineries, steel plants, paper mills, and textile mills of broadly comparable industrial scale. They consume more water than cement plants, dairies, and pharmaceutical plants. They consume roughly the same as a textile mill, less than half what a refinery consumes, and about 40 percent of what a coal-fired thermal power plant of equivalent generating capacity withdraws.

The most consequential comparison is the thermal power plant. A 1.2 GW coal-fired thermal plant consumes 24 to 38 million cubic metres of water per year. The same 1.2 GW of generating capacity supplied by solar PV consumes approximately 50,000 cubic metres per year (mostly panel washing). A data centre procuring its electricity from renewable PPAs (the HyperNext Khawada captive solar plus partnership wind portfolio) is therefore avoiding approximately 24 to 38 million cubic metres of water that the equivalent thermal generation would have consumed. The data centre water footprint of 16.7 million cubic metres per year, less the avoided 24 to 38 million from renewable substitution, is in some accounting frameworks a net positive of 7 to 21 million cubic metres of water per year.

This is not a free pass. The water savings from renewable substitution accrue to the entire grid, not specifically to the data centre operator, and counting them as a deduction against data centre consumption is an accounting choice rather than a physics statement. But it is also not nothing. The marginal water cost of one megawatt of AI compute, when that megawatt is supplied from renewable sources, is materially less than the marginal water cost of one megawatt of conventional load supplied from the grid average.

5. The economic productivity of water

Comparing water consumption between facilities is informative. Comparing the economic value produced per litre of water consumed is more so. The metric is sometimes called water productivity (the inverse of water intensity per unit output). For the comparators in Sections 3 and 4 we estimate the annual revenue or economic value generated per cubic metre of water consumed.

Facility type	Annual water (million m³)	Annual revenue (INR crore)	Revenue per m³ (INR)	Sources
HyperNext Kakinada (data centre)	16.7	32,500 to 60,000	~20,000 to 36,000	HyperNext modelling, INR 0.74/M tokens × throughput estimates
Coal-fired thermal power plant	30 (mid)	5,000 to 7,000	~1,700 to 2,300	CEA tariff data × generation
Refinery (15 MMTPA)	25 (mid)	60,000 to 90,000	~24,000 to 36,000	Public oil company filings
Steel plant (5 MTPA)	26 (mid)	20,000 to 30,000	~7,700 to 11,500	Public steel company filings
Paper mill (kraft)	22 (mid)	2,500 to 4,000	~1,100 to 1,800	Paper industry data
Textile mill (vertically integrated)	10 (mid)	3,000 to 5,000	~3,000 to 5,000	Indian textile industry data
5-star hotel (200 rooms)	0.2 (mid)	80 to 150	~4,000 to 7,500	Hotel industry data
18-hole golf course (membership)	0.2 (mid)	15 to 30	~750 to 1,500	Indian Golf Union data
Shopping mall (1M sq ft)	0.25 (mid)	200 to 350 (gross merchandise value attributed)	~8,000 to 14,000	Mall operator filings
Indoor water park	0.8 (mid)	40 to 80	~500 to 1,000	Amusement industry data

Read the table

The economic productivity of water in a hyperscale AI data centre is roughly INR 20,000 to 36,000 per cubic metre of water consumed. That is in the same range as a modern refinery. It is materially higher than steel, paper, textile, and the urban-leisure comparators. It is roughly ten times the economic productivity of water in a coal-fired thermal power plant.

This metric is not the only one that matters. Water has social, ecological, and cultural value that is not captured in economic output per cubic metre. But within the framing of comparing competing industrial and commercial uses of the same water resource, the data centre is among the more productive uses of water in the comparison set.

The qualifier we add: water productivity is meaningful at the margin, not in totality. Drinking water, sanitation, and ecological flow have first call on the water resource and no comparison of economic productivity is going to change that. The relevant question is how the remaining industrial and commercial water is allocated between competing uses. On that question, the data centre with the HyperNext architecture is not a low-productivity choice.

6. The cooling architecture, in detail

The HyperNext architecture differs from the conventional evaporative-cooled data centre in two related decisions: where the heat is captured (direct-to-chip rather than via the air), and where it is rejected (dry coolers rather than cooling towers). The combination removes essentially all the operational water consumption associated with cooling. This section walks through how it works and what the trade-offs actually are.

How a HyperNext facility rejects heat

A traditional evaporative-cooled data centre rejects waste heat to the atmosphere through cooling towers. The towers spray water over heat exchangers exposed to ambient air. Water evaporates and carries the heat away. The arithmetic is favourable from a thermodynamic standpoint: water has high latent heat of vaporisation, so a small mass of evaporated water carries a large quantity of heat. The penalty is that the evaporated water is lost to the local water system. A 1.2 GW IT load consuming a typical 1.8 L/kWh in an evaporative architecture loses approximately 50 million cubic metres of water per year to evaporation, before counting the embedded water in the grid electricity that the cooling-tower fans and chillers themselves consume.

HyperNext rejects heat through dry coolers fed by closed propylene glycol loops. The path is: chip cold plate → primary glycol loop → coolant distribution unit → secondary glycol loop → external dry cooler → ambient air. No evaporation. No cooling towers. No water chillers. The dry cooler is essentially a very large finned heat exchanger with fans. It transfers heat from the glycol to the air through sensible heat transfer, the same way a car radiator does.

The architecture is enabled by direct-to-chip cooling. Because the heat is captured at the chip via a cold plate rather than via the air that has been heated by the chip, the coolant supply temperature can run substantially warmer than in a chilled-water architecture: typical supply temperatures are 35 to 40 degrees Celsius. At those supply temperatures, dry coolers can reject heat to Indian ambient conditions across the full operating envelope without supplementary evaporative assistance. This is the engineering insight that makes the architecture work for the Indian climate.

Why there is no PUE penalty

It is common to assume that giving up evaporative cooling means a PUE penalty. That is true for the comparison case of dry-air-cooled chillers versus evaporatively-assisted chillers in a traditional air-cooled data centre. It is not true for the comparison that matters here: direct-to-chip cooling with glycol-to-dry-cooler heat rejection versus chilled-water cooling with cooling-tower heat rejection.

The HyperNext architecture is more efficient on power, not less. Three reasons. First, the cooling system has no compressor stage. A chilled-water plant rejects the IT heat plus the compression work that the chiller itself adds; a glycol-to-dry-cooler system rejects only the IT heat. Second, the thermal resistance from chip to coolant is much lower for direct-to-chip than for air-cooled servers, so the coolant can run warmer and the heat-rejection equipment runs at higher driving temperature differences. Third, the absence of an evaporative stage means no cooling-tower fans, no makeup-water pumps, and no chemical-treatment infrastructure consuming auxiliary power.

The HyperNext PUE design target is 1.25 to 1.30 in moderate conditions, 1.35 at 42 degrees Celsius peak ambient. This is at the low end of the achievable range for an Indian climate, broadly comparable to (or better than) the best evaporative-cooled facilities. The water benefit is achieved without a power-efficiency cost.

The water that gets saved

The HyperNext Hyderabad Phase 1 facility direct on-site cooling water consumption is zero. The total Layer 1 water consumption (humidification, sanitation, facilities) is approximately 70,000 cubic metres per year. The comparison case (a fully evaporative cooling design at the same 64 MW IT load) would consume approximately 700,000 to 1,100,000 cubic metres of cooling water alone, plus the same 70,000 cubic metres of facility water. The architectural choice saves between 700,000 and 1,100,000 cubic metres of on-site water per year on the 64 MW scale.

On the planned Kakinada 1.2 GW campus, the same architectural choice saves approximately 35 to 45 million cubic metres of on-site water per year. That is comparable to the annual water consumption of a one million population urban centre at Indian per-capita rates. The dry-cooler-with-glycol architecture on a single AI data centre saves the water equivalent of an Indian district headquarters town.

What this means for the broader debate

The public conversation about data centre water consumption is calibrated to the conventional evaporative-cooled architecture that dominated commercial deployment from roughly 2010 to 2024. The HyperNext architecture, deployed because direct-to-chip cooling on AI workloads forced a rethink anyway, ends the water debate as a meaningful concern for new build-outs. The water headline number for a HyperNext-architecture data centre is dominated by facility water (humidification, sanitation, landscape) rather than cooling water. The cooling-water share is essentially zero in normal operation. Section 7 addresses the related question of what the broader conversation should be about, given that data centre water has stopped being the right concern.

7. What the conversation should be about

Indian water stress is real. NITI Aayog has 18 of 21 major Indian cities classified as severe-water-risk by 2030. Per-capita freshwater availability has dropped from 5,177 cubic metres in 1951 to under 1,400 cubic metres today, with the trajectory pointing below 1,100 by 2050. The crisis is acute and the response needs to be at scale.

Where the volume actually is

Indian freshwater withdrawal by sector breaks down approximately as follows. Agriculture: 78 percent. Domestic and municipal: 6 percent. Industry: 5 percent. Power: 7 percent. Other: 4 percent. The numbers are from the Central Water Commission and NITI Aayog water management reports.

The data centre share of industrial water withdrawal is small. If India had 10 GW of installed AI data centre capacity by 2030, operating at HyperNext architecture water intensities, the total industry water footprint would be approximately 140 million cubic metres per year. That is approximately 0.02 percent of total Indian freshwater withdrawal at current totals.

At a conventional evaporative architecture, the same 10 GW of installed capacity would consume approximately 450 million cubic metres per year. Still under 0.07 percent of total Indian freshwater withdrawal.

The volumes that move the Indian water needle are agricultural, where modest efficiency gains have outsized impact, and municipal, where leakage rates of 30 to 50 percent in many Indian cities are well documented and addressable through known interventions. Data centres do not appear in the top ten contributors to Indian water stress at any reasonable level of analysis.

What the productive conversation looks like

The water conversation around data centres should focus on three things, and only three things.

First, the architecture choices the industry makes. Operators that choose evaporative cooling for marginal PUE benefits in water-stressed geographies are making a defensible engineering choice in some climates and an indefensible one in others. The industry conversation should be about which architectures are appropriate in which geographies. The HyperNext position on this is in HN-RP-001: dry-cooling primary architecture is the right answer for almost all Indian geographies.

Second, the reporting honesty. Operators should report WUE quarterly under ISO/IEC 30134-9, with the methodology disclosed and the underlying meter data auditable. The industry has not yet converged on this reporting discipline. Data centre water reporting in India remains uneven, and the cases that draw negative attention are typically the cases where reporting is opaque rather than where reported numbers are large.

Third, the watershed responsibility. Water consumed in a stressed catchment is a moral and operational responsibility of the operator consuming it. The Nagmati Programme (HN-RP-001) describes how HyperNext approaches this responsibility. Other operators have similar or different approaches. The industry conversation should be about whether watershed-scale responsibility is part of the baseline expectation or an optional extra. The HyperNext position is that it should be baseline.

What the unproductive conversation looks like

The unproductive version of the conversation, which has come to dominate some segments of public discourse, has three patterns we would like to see less of.

The "bottles of water per AI query" comparison. The calculation is mathematically defensible but it is also designed to maximise an emotional response without engaging with the underlying scale. A single Indian-language LLM inference query consumes some milligrams of water at the embedded electricity layer, less at the cooling layer for a well-designed facility. Translating that to "X bottles of water per query" produces a number that sounds large because bottles are a small unit and queries are a large number, but the same translation applied to other facilities produces equally large numbers that nobody finds alarming. A 5-star hotel room-night consumes "thousands of bottles of water" per night. Nobody writes that headline.

The "data centres versus drinking water" framing. Data centres in India do not compete with drinking water supply at any scale. The water sources for industrial and commercial cooling are separate from municipal drinking water in almost all Indian states, with different tariff structures, different allocation queues, and different regulatory frameworks. Framing the question as if it were drinking water versus AI is a rhetorical move that does not survive contact with the underlying allocation system.

The "AI is luxury, water is necessity" framing. This is the strongest version of the critical argument and it is the one that deserves the most respectful engagement. The honest response is that AI infrastructure, like other industrial and commercial infrastructure, is a use of water that competes with other uses. The competition is real. The right framework for resolving the competition is per-litre-of-water-allocated economic and social productivity, not categorical exemptions for some uses against others. By that framework, data centres are not the use that should be at the front of the queue for water rationing.

8. Headlines and recommendations

HEADLINES

- A HyperNext-architecture 1.2 GW AI data centre consumes approximately 16.7 million cubic metres of water per year on the gross basis. A conventional evaporative-cooled facility at the same IT load would consume 45 to 55 million cubic metres per year.
- A 1.2 GW HyperNext data centre consumes less water than a coal-fired thermal power plant of equivalent generating capacity (24 to 38 million cubic metres per year). The renewable electricity HyperNext procures from Khawada captive solar avoids the thermal-generation water that the equivalent grid mix would have required.
- On a per-facility basis, hyperscale data centres of the HyperNext architecture consume less water than refineries, steel plants, paper mills, and textile mills of comparable industrial scale. They consume more water than cement plants, dairies, and pharmaceutical plants.
- On a per-user-served basis, AI inference for a heavy user consumes approximately 84 litres per year through the data centre infrastructure. The same user consumes approximately 49,000 litres per year for domestic purposes. AI infrastructure is 0.17 percent of personal water consumption for a user at the upper end of usage.
- On an economic productivity per cubic metre basis, AI data centres are among the more productive uses of water in the comparison set. The water productivity is approximately INR 20,000 to 36,000 per cubic metre, comparable to a modern refinery and ten times the productivity of a coal-fired thermal power plant.
- The Indian water stress crisis is real, but data centres are not in the top ten contributors. 10 GW of installed AI capacity at HyperNext architecture water intensity is approximately 0.02 percent of total Indian freshwater withdrawal.
- The productive water conversation around data centres focuses on three things: architecture choices (dry cooling versus evaporative), reporting honesty (ISO/IEC 30134-9 quarterly), and watershed responsibility (per Nagmati Programme).

For Indian policy stakeholders

1. Treat data centre water allocation as a small fraction of total Indian water allocation. It is a use that competes with other uses for a small share of the total, and policy should allocate proportional attention to it.
2. Encourage architectural disclosure as part of the data centre licensing process. An operator using evaporative cooling in a water-stressed geography should disclose that choice and the reasons for it. An operator using dry cooling should be able to demonstrate the architecture.

3. Recognise watershed restoration as a legitimate offset for operators that commit to it, with verification and audit cadence matched to the offset claim.
4. Prioritise the interventions that move the volume needle. Agricultural water efficiency, municipal leakage reduction, and watershed restoration at scale are the levers that change Indian water stress at meaningful scale.

For Indian data centre customers and end users

1. Ask operators for the actual WUE numbers, the architectural choices behind them, and the watershed accounting if any. The information should be available.
2. Recognise that operator water disclosure is uneven across the industry. An operator that does not report should be asked why.
3. Be willing to pay a modest premium for AI capacity served by operators with credible water architecture and watershed responsibility. The premium is small relative to total inference cost and large relative to the water saved.

For the public conversation

We are not asking for an exemption from scrutiny. We are asking for proportionality. The water consumption of a hyperscale data centre is a real number, reported publicly, and the architectural choices that reduce it are documented and reproducible. The Indian water crisis is also a real problem, with documented dominant contributors. Treating data centres as one of those dominant contributors is incorrect on the math. Engaging with the real contributors and the real interventions is harder and less satisfying, but it is the only path that improves the situation at scale.

9. References and sources

This paper draws on the following published sources. Where the underlying calculations are simplified for presentation, the methodology is available on request.

Sectoral water consumption data

- **Central Water Commission, Government of India.** National water resources assessment, current edition. Sectoral allocation of freshwater withdrawal.
- **NITI Aayog.** Composite Water Management Index, 2018 baseline with annual updates. Indian water stress by state and city.
- **NITI Aayog.** Water Index 2.0 report, 2023. State-by-state water security ranking.
- **Central Electricity Authority of India.** Annual report on specific water consumption of thermal power stations, current edition. Per-MWh water intensity by plant generation.
- **Bureau of Energy Efficiency, Ministry of Power.** Industrial sector energy and water audit reports.

Industry-specific data

- **Indian Hotels Association.** Sustainability reports and water audit data from member hotels.
- **Indian Golf Union.** Sustainability framework and water consumption data from member courses.
- **Confederation of Indian Industry.** Sectoral water benchmarking reports across manufacturing categories.
- **Indian Refineries Association.** Member sustainability data on water consumption per refinery output.
- **Steel Authority of India and Indian Steel Association.** Member water consumption data.
- **Indian Paper Manufacturers Association.** Industry water intensity benchmarks.
- **Confederation of Indian Textile Industry.** Member water audit and benchmarking data.
- **Indian Hotels Association.** Member sustainability filings.

Standards and frameworks

- **ISO/IEC 30134-9.** Information technology. Data centres key performance indicators. Part 9: Water Usage Effectiveness (WUE). The standard against which HyperNext reports site WUE quarterly.
- **ISO 14046:2014.** Environmental management. Water footprint. Principles, requirements and guidelines.
- **WRI Aqueduct Water Risk Atlas 4.0.** Catchment-level water stress mapping.

- **CDP Water Security disclosure framework.** The corporate water reporting framework on which the HyperNext WUE-3 categories map.

Adjacent HyperNext publications

- **HyperNext Research HN-RP-001.** The Nagmati Programme. The water architecture and watershed accounting framework that underpins this paper.
- **HyperNext Research HN-RP-002.** 800VDC Power Architecture. The power architecture that reduces embedded water consumption through PPA structure.
- **HyperNext Research HN-RP-006.** Liquid Cooling at 600 kW per Rack. The cooling architecture engineering that reduces direct site water consumption.
- **HyperNext Research HN-RP-007.** The India AI Compute Gap. The capacity context within which the water comparison sits.

HyperNext
Data Centers

HyperNext Research

We publish engineering and policy papers because the Indian conversation about AI infrastructure needs more substance than marketing material provides. The papers state methodology openly so other operators can run the same analysis on their own facilities. They report findings that may not flatter the HyperNext commercial position. They get review from the engineering team and the communications partners.

Correspondence on methods, figures, and conclusions: **hello@hypernxt.com**. We read every email.

HN-RP-008 · The Water Math
27 May 2026 · v1.0

www.hypernxt.com/research
hello@hypernxt.com · +91 99784 23333