




STRUCTURE
RESEARCH

MARKET CLASSIFICATION FRAMEWORK

GLOBAL DATA CENTRE MARKET TIERING 2026

A four-axis classification of 101 data centre markets — commercial colocation, hyperscale cloud and self-build, AI market maturity, and connectivity and network gravity.

101 MARKETS, RANKED BY COMPOSITE
SCORE AND COLOURED BY TIER



■ TIER 1 · 21 ■ TIER 2 · 29 ■ TIER 3 · 30 ■ TIER 4 · 21

EDITION

August 2026

STRUCTURERESEARCH.NET

CONTENTS

A Four-Axis Classification of 101 Data Centre Markets

Structure Research · August 2026

01	Executive summary	04
	Markets in the Tier 1 composite band	05
	Scoring basis & score distribution	06
02	Why a four-axis classification	07
	Colocation and self-build have decoupled	08
	AI demand has created markets from nothing	10
	Power has replaced land as the binding constraint	11
	Cloud infrastructure and colocation scale have come apart	12
03	The framework	14
A	Commercial Colocation — 40%	16
B	Hyperscale Cloud & Self-Build — 30%	18
C	AI Market Maturity — 15%	22
D	Connectivity & Network Gravity — 15%	24
	Tiers, scale gates & the announced-versus-committed rule	26
04	Results — 101 markets scored	28
05	Regional analysis	34
06	Coverage and sources	38

THE CLASSIFICATION AT A GLANCE

101

MARKETS
CLASSIFIED

16

METRICS, FOUR PER
AXIS


21

TIER 1 MARKETS


240MW

TIER 1 OPERATIONAL
COLOCATION FLOOR

TIER DISTRIBUTION BY REGION



AXIS WEIGHTS IN THE COMPOSITE



Public report · August 2026. Evidence reviewed through August 2026; operational commercial-colocation capacity for A1 and the MW gates is measured at 30 June 2026. Citation and media-reuse terms appear on page 39. Download, data and updates: datacentreresearch.com.

infra / STRUCTURE

Where AI and hyperscale meet land, energy and capital.

An AI infrastructure ecosystem gathering built for operating executives — private 1:1 meetings, roundtables and solution discovery in an intimate deal-making environment.

DATES

October 6–8, 2026

Wynn Las Vegas · Las Vegas, Nevada

700+

DIGITAL INFRASTRUCTURE
LEADERS ATTENDING

3

DAYS OF MEETINGS,
ROUNDTABLES AND DATA

2026 THEME

INTERLUDE

*an intervening or interruptive
period, space or event*

The sector is in a clear transitory period. Hyperscale cloud brought it here and that wave is far from over, but AI is already changing the vocabulary — AI factories for training and inference rather than availability zones, GPUs alongside CPUs, liquid cooling alongside air, and infrastructure moving to where the power is rather than bringing power in. Join us to make sense of the interlude between hyperscale cloud and AI.

WHAT THE THREE DAYS DELIVER

Private 1:1 buyer meetings

Scheduled sessions with buyers of capacity, cloud, interconnection, hardware and services.

Hyperscale & neocloud program

2026 participants include AWS, Google Cloud, Oracle, Meta, ByteDance, Groq, Lambda and Together AI.

Curated analyst data

Independent Structure Research insight — including this classification — as the backdrop for frank discussion.

C-level networking

Hyperscale, neocloud, real estate, energy, private equity, infrastructure and equipment vendors.

Register now for infra / STRUCTURE 2026

Secure access to the buyers, partners and insight defining the next wave of digital infrastructure.

REGISTER

AGENDA

infrastructuresummit.io/agenda-2026

ATTENDING COMPANIES

Request the 2026 list

SPONSORSHIP

Request a prospectus

SECTION 01

Executive summary

Structure Research classifies 101 data centre markets worldwide into four tiers. Each market is scored on four independent axes — commercial colocation, hyperscale cloud and self-build, AI market maturity, connectivity and network gravity — and a weighted Composite Score maps to a provisional tier through fixed thresholds; the applicable Tier 1 or Tier 2 gate then determines final classification.


The classification is built for an industry in which a single ordinal ranking no longer describes how markets work. A market can be a region’s largest AI destination and one of its thinnest enterprise ecosystems at the same time, or its densest interconnection hub and one of its weakest AI stories. The four-axis structure holds those facts separately, so a reader can rank on the dimension relevant to a siting, investment or competitive decision rather than anchoring on a single blended number.

Headline results. Of the 101 markets, 21 classify as Tier 1: six in North America, four in EMEA, ten in Asia Pacific and one in Latin America. Tier 2 holds 29 markets, Tier 3 holds 30 and Tier 4 holds 21. The Tier 1 set is listed overleaf, together with the one market that carries a Tier 1 composite but is held below it by a scale gate.

Asia Pacific holds the largest concentration of Tier 1 markets of any region, and also the largest Tier 4 cohort — a direct consequence of AI-driven capacity arriving in locations that were not data centre markets under any conventional definition three years ago. EMEA carries the widest spread, with 21 of its 36 markets in Tier 3 or Tier 4.

101 MARKETS CLASSIFIED 35 Asia Pacific · 36 EMEA · 25 North America · 5 Latin America	21 TIER 1 MARKETS Composite of 3.50 and above, both Tier 1 gate conditions met	16 SCORED METRICS Four per axis, against fixed global definitions	0.77 AXIS A-B CORRELATION Related enough to confirm, too low to collapse into one
--	---	--	--

EXHIBIT 1 · TIER DISTRIBUTION BY REGION, ALL 101 MARKETS



Bars are drawn on a common absolute scale, so region width reflects the number of markets covered rather than share. Source: Structure Research Market Tiering model, August 2026.

HEADLINE RESULTS

Markets in the Tier 1 composite band

Twenty-one markets classify Tier 1. Bangkok is shown in position 19 because it posts a Tier 1 composite; it is held at Tier 2 by the operational-colocation scale gate. Ranks follow unrounded composites, so rows displaying the same two-decimal score may still rank separately; equal ranks denote composites equal at full precision.

RANK	MARKET	REGION	COMPOSITE SCORE		CLASSIFICATION
1	Northern Virginia (incl. ext.)	North America	4.63		TIER 1
2	Tokyo	Asia Pacific	4.57		TIER 1
3	Singapore	Asia Pacific	4.43		TIER 1
4	London	EMEA	4.40		TIER 1
5	Sydney	Asia Pacific	4.26		TIER 1
6	Dallas-Fort Worth	North America	4.18		TIER 1
7	Sao Paulo	Latin America	4.18		TIER 1
8	Jakarta	Asia Pacific	4.15		TIER 1
9	Phoenix	North America	4.10		TIER 1
10	Mumbai	Asia Pacific	4.09		TIER 1
11	Atlanta	North America	4.06		TIER 1
12	Chicago	North America	4.02		TIER 1
12	Frankfurt	EMEA	4.02		TIER 1
14	Paris	EMEA	3.89		TIER 1
15	Melbourne	Asia Pacific	3.83		TIER 1
16	Osaka	Asia Pacific	3.81		TIER 1
17	Seoul	Asia Pacific	3.75		TIER 1
18	Hong Kong	Asia Pacific	3.66		TIER 1
19	Bangkok	Asia Pacific	3.64		TIER 2 · GATE
20	Johor	Asia Pacific	3.61		TIER 1
21	Silicon Valley	North America	3.61		TIER 1
22	Milan	EMEA	3.60		TIER 1

THE GATE, AND WHY IT MATTERS

Bangkok posts a Tier 1 composite of 3.64 on committed capital and operational hyperscaler capacity, but holds 86MW of operational colocation against a 240MW floor. The gate describes what a market *is* today; the composite also contains what it is becoming. That distance is informative in its own right — it identifies the clearest promotion candidate in the set.

Bars are drawn to a 0–5 scale. Rows classified Tier 1 clear both Tier 1 gate conditions: at least 240MW of operational colocation and a B1 hyperscaler-presence score of at least 3. Source: Structure Research Market Tiering model, August 2026.

SCORING BASIS

What the scores are measured on

Capacity is taken from the Structure Research Data Centre Model on a 2020–2031 series. Hyperscale presence, leasing and self-build are measured from Structure Research’s facility-level Hyperscale Self-Build Data Centre Model and from 16 deep-dive DCAI market models covering 17 of the markets classified here.


Connectivity is anchored on third-party sources: subsea systems from TeleGeography and submarinenetworks.com, interconnection and facility

network density from PeeringDB, cloud on-ramps from provider technical documentation.

Of the 16 metrics, 11 are measured directly on Structure Research data; power runway is scored against regulator and utility evidence, and the four connectivity metrics are anchored on auditable third-party sources so that axis can be reproduced independently.

EXHIBIT 2 · COMPOSITE SCORE DISTRIBUTION

101 markets in half-point bins, with the four fixed tier thresholds marked



Bins are half a point wide and coloured by the tier the band maps to. The distribution is bimodal around 1.0–1.5 and 2.0–3.0: the classification separates a large middle band of functioning secondary markets from a long tail of emerging entrants. One market, Bangkok, carries a Tier 1 composite while classifying below it on a scale gate. Source: Structure Research Market Tiering model, August 2026.

TIER 1 SCORE BAND

≥ 3.50

22 markets

TIER 2 SCORE BAND

2.50 – 3.49

28 markets

TIER 3 SCORE BAND

1.50 – 2.49

30 markets

TIER 4 SCORE BAND

< 1.50

21 markets

Final classifications after gates: Tier 1 · 21; Tier 2 · 29; Tier 3 · 30; Tier 4 · 21. Bangkok is the only reassignment.

02

SECTION TWO · METHOD RATIONALE

WHY A FOUR-AXIS CLASSIFICATION

Four structural changes since 2023 make a single-axis ranking unreliable. Each one is a case where two facts a single number would blend together have separated far enough to be worth reporting apart.

2.1

Colocation and self-build have decoupled

Two channels, moving independently · $r = 0.77$

2.2

AI demand has created markets from nothing

Capacity arriving on four continents with no colocation base

2.3

Power has replaced land as the constraint

And it falls hardest on the markets at the top

2.4

Cloud infrastructure and scale have come apart

Presence is not production infrastructure

SECTION 02

Why a four-axis classification

Four structural changes since 2023 make a single-axis ranking unreliable.

2.1 Colocation and self-build have decoupled

For most of this industry's history, the size of a metro's commercial colocation market was a reasonable proxy for everything else about it: hyperscalers leased where colocation was deep, enterprises followed the hyperscalers and networks converged on the same buildings. That relationship no longer holds. Hyperscalers now express commitment through two channels that move independently — leased colocation capacity and self-build campuses — and the second is invisible to any ranking built on third-party supply.

Some markets are considerably more hyperscale than their colocation footprint suggests, because the self-build installed base was built early and third-party

growth has since been constrained by grid policy or land. Others are the inverse: deep, competitive colocation markets with several cloud regions and no hyperscale self-build capacity in operation, where the self-build commitment exists only as construction and land bank. Neither shape is unusual anymore, and both are common enough that an average of the two conceals more than it conveys.

A single blended score would describe neither channel, and would treat leased and self-build capacity as the same input when they carry different tenure, different capital at risk and different implications for anyone taking a position in the market.

0.77

AXIS A ↔ AXIS B
CORRELATION, N = 101

Axis A and Axis B are therefore held apart, so a reader can see which channel a market's capacity arrives through. The two axes correlate at 0.77 across the set — high enough to confirm they are related, and too low to justify collapsing them into one.

SHAPE ONE · SELF-BUILD HEAVY

More hyperscale than the colocation footprint suggests. The self-build installed base was built early; third-party growth has since been constrained by grid policy or land. Dublin is the European example — the deepest measured self-build base in the region at 733.6MW across four hyperscalers, against an AI axis of 0.40.

SHAPE TWO · LEASED ONLY


Deep, competitive colocation with several cloud regions and no hyperscale self-build capacity in operation — the commitment exists only as construction and land bank. Johor is the clearest case: 1,146MW operational, all hyperscaler capacity leased, one building under construction and a large Microsoft land bank.

EXHIBIT 3

Commercial colocation against hyperscale commitment

Each point is one of the 101 classified markets, plotted on its Axis A score (commercial colocation, horizontal) against its Axis B score (hyperscale cloud and self-build, vertical). The dashed line is the least-squares fit. Markets above the line carry more hyperscale commitment than their commercial colocation base would predict; markets below it carry less.

● NORTH AMERICA ● EMEA ● ASIA PACIFIC ● LATIN AMERICA



Northern Virginia and Tokyo both score 5.00 on each axis and therefore plot at the same location; the two dots are drawn side by side inside a double ring. Labelled markets are those with the widest divergence between the two axes, plus those carrying the strongest recent commitment activity. Axis scores are weighted averages of four metrics each, on a 0–5 scale against fixed global rubrics. Pearson correlation across the full set is 0.77. Source: Structure Research Market Tiering model, August 2026.

TOP RIGHT · BOTH CHANNELS DEEP

Northern Virginia, Tokyo and Singapore score at or near maximum on both. Jakarta, Bangkok and Johor are the set's clearest recent risers, carrying committed pipeline on both axes at once.

BOTTOM RIGHT · COLOCATION WITHOUT HYPERSCALE

New York Metro (A 4.50, B 1.65) and Chennai (A 3.60, B 1.00) hold real commercial depth with little or no hyperscale commitment attributable to the metro.

TOP LEFT · HYPERSCALE WITHOUT A MARKET

Dublin, Zaragoza, Queretaro and Riyadh carry hyperscale commitment ahead of their commercial base, in most cases because policy or grid position constrains third-party supply.

2.2 AI demand has created markets from nothing

Demand has begun arriving in places with no colocation market to absorb it, and it has done so on four continents. The six markets in Exhibit 4 were not data centre markets under any conventional definition three years ago.

They divide into two shapes. The first is capacity arriving where there is almost no commercial colocation at all — a single hyperscaler campus, a national AI programme or a financed neocloud build, sited for power and land rather than for customers, in a location with no operator bench, no exchange and no enterprise base.

The second is more awkward for a size-based ranking: locations that now hold large single-tenant AI capacity behind one or two developers, where no third party can purchase capacity and no interconnection ecosystem exists. Ranked on total operational data-centre capacity, including single-tenant AI capacity, two of these locations would sit among the 20 largest in the world; ranked on sellable commercial colocation capacity, neither would.

The classification places the first group in Tier 4, because Axis A registers that there is almost nothing commercial in them and the AI axis carries only 15% of the composite, and the second in Tier 2 and Tier 3, because scale without operator depth, enterprise demand or interconnection describes a site rather than a market.

What the four-axis structure adds is that a reader can see these locations at all, can distinguish single-tenant capacity from a functioning market and can see which are backed by construction rather than by announcement — gigawatt ambitions with no tenant, no capital commitment and no power contract are carried at zero. A framework that admits markets only once they have a colocation base admits them several years after the relevant siting decisions have been taken.

EXHIBIT 4 · THE AI-CREATED MARKETS, AND WHERE THE FRAMEWORK PUTS THEM

MARKET	AXIS C	AXIS A	COMPOSITE	TIER	WHAT THE AXIS PROFILE RECORDS
West Texas (incl. Abilene)	4.60	2.50	2.46	TIER 3	Largest operational AI campus in the US; no colocation market, no interconnection, no cloud
Ellendale (ND)	3.60	2.00	1.57	TIER 3	The same shape at smaller scale; Axis D scores zero
Hokkaido	3.20	1.20	1.40	TIER 4	Real AI capacity against almost no hyperscale presence
Zaragoza	2.90	0.80	1.48	TIER 4	Hyperscale commitment ahead of any commercial base
Tasmania	2.20	0.70	0.84	TIER 4	Hydro position and consented site; short on tenancy and construction
Jeollanam-do (Solaseado)	1.60	0.20	0.56	TIER 4	National programme; nothing yet operational

Axis C is AI market maturity, Axis A commercial colocation; cells are tinted by score. Readers seeking an AI siting ranking should sort on Axis C rather than the composite. Source: Structure Research Market Tiering model.

2.3 Power has replaced land as the binding constraint, asymmetrically

Until roughly 2022, the scarce input in a data centre market was suitable land near fibre and customers. It is now grid interconnection capacity — and this matters for classification because the constraint does not fall evenly. It falls hardest on the markets a size-based ranking puts at the top.

Several of the largest and most established metros in Asia, Europe and North America now operate under connection moratoria, refusal regimes, multi-year queues or conditional frameworks that make new capacity slow, costly or contingent on generation the developer has to bring itself.

A separate and much smaller set of markets — some with behind-the-meter generation, some with hydro or curtailed renewables, some with restarted nuclear —

can energise load in a fraction of the time.

Power runway is therefore scored as its own metric rather than folded into a growth forecast, because a market's installed base describes what has already been built while its grid interconnection position determines whether the next tranche can be. Those two facts have now separated far enough that a single ordinal number reports the past while implying a future for which it holds no evidence.

EXHIBIT 5 · THE CONSTRAINT IS INVERTED

Across the 101 markets, colocation scale correlates negatively with power runway. The pattern is measurable rather than anecdotal.

16 MARKETS · WEAKEST POWER POSITION

11

are Tier 1 or Tier 2 — moratoria, refusal regimes and multi-year queues concentrate in the largest markets




Power runway is metric C4, scored on evidenced grid policy and physics — published connection policies, queue disclosures, tariff rulings and generation positions — rather than on developer statements or utility marketing. Announced generation projects do not count until they are contracted. Source: Structure Research, regulator, utility and system-operator evidence, August 2026.

11 MARKETS · STRONGEST POWER POSITION

0

are Tier 1 — structural surplus and dedicated behind-the-meter generation sit almost entirely outside the established metros



Installed base is Axis A's job. On current evidence, installed base and power runway point in opposite directions **more often than not**.

2.4 Cloud infrastructure and colocation scale have come apart

A hyperscaler being present in a market and a hyperscaler having built production cloud infrastructure there are no longer the same fact.

The unit that separates them is the availability-zone region: three or more physically separated zones inside one metro, which is what allows a customer to run a production workload with failover that stays in the market. A cache node, a local zone or a single-zone region cannot do that. If the single zone fails there is nowhere inside the market for the workload to fail over to, so the customer has to keep its production copy in another market.

Across the 101 markets, 55 carry at least one full availability-zone region and 22 carry three or more from different platforms. A further 28 carry no cloud region at all, of which 13 hold edge presence only. The remaining 18 carry live regions that are not zone-redundant. The distribution does not track colocation scale, and in several large markets it runs directly against it.

EXHIBIT 6 · CLOUD REGION DEPTH ACROSS THE CLASSIFIED SET

All 101 markets, graded on metric B2. The unit is the availability-zone region, not the presence of a platform.

55 MARKETS CARRY AT LEAST ONE FULL AVAILABILITY-ZONE REGION

28 CARRY NO CLOUD REGION AT ALL



Region depth is classified from what each platform publishes about its own infrastructure — region lists, zone counts and availability-zone support flags — so the metric can be reproduced independently. Where a platform does not publish a zone count, its region is treated as not zone-redundant. Source: provider infrastructure documentation compiled by Structure Research, August 2026.

WHY THIS BELONGS IN A CLASSIFICATION RATHER THAN A FOOTNOTE

The classification scores cloud region depth as its own metric, counting how many platforms have committed a full availability-zone region rather than how many have a presence. It is nationality-blind: an Alibaba or Huawei region counts on the same test as an AWS or Azure one.

A full region anchors production enterprise workloads and generates the cloud on-ramp gravity that pulls third-party colocation demand behind it over the following years. Installed megawatts describe what a market has already absorbed; region depth describes what kind of demand it is equipped to absorb next.

EXHIBIT 7

Scale without regions, and regions without scale


Atlanta is the clearest case. It classifies Tier 1 on 891MW of installed colocation and one of the strongest AI profiles in the United States, and it holds no cloud region a customer can buy — its hyperscaler footprint is self-build and leased capacity, its generally available cloud presence is an AWS Local Zone and a Google cache node, and the CoreWeave regions in the market are reserved for selected customers rather than sold openly. Hillsboro at 647MW has the same shape. Johor holds 1,146MW and no zone-redundant region at all: its cloud anchors are a single-domain Oracle region and a two-zone Alibaba region. Silicon Valley carries six platforms with regions and only one of them, AWS, is zone-redundant.

The inverse is just as common. Zurich, Warsaw, Brussels and Auckland each carry two or three full availability-zone regions on colocation bases between 74MW and 191MW. These are not large markets, and on installed megawatts alone none of them looks consequential.

SCALE WITHOUT REGIONS · FOUR LARGE MARKETS, NO ZONE-REDUNDANT CLOUD


Atlanta

891MW installed · no cloud region a customer can buy




Hillsboro

647MW installed · no cloud region in the metro




Johor

1,146MW installed · no zone-redundant region



Silicon Valley


Six platforms with regions · one zone-redundant



REGIONS WITHOUT SCALE · SMALL COLOCATION BASES, GENUINE MULTI-PLATFORM DEPTH


Zurich

191MW · AWS, Azure and Google full regions




Warsaw

Two to three full availability-zone regions




Brussels

Two to three full availability-zone regions



Auckland

Two to three full availability-zone regions



Bars are the four axis scores on a 0–5 scale: **A** commercial colocation, **B** hyperscale cloud & self-build, **C** AI market maturity, **D** connectivity & network gravity. Colocation figures are installed commercial megawatts on the current Structure Research Data Centre Model cycle. Source: Structure Research Market Tiering model, August 2026.

03

SECTION THREE · DEFINITIONS

THE FRAMEWORK

Every market is scored 0–5 on 16 metrics, four per axis, against fixed definitions. Axis scores are weighted averages of their metrics; the four axes carry fixed weights in the composite.

A 40%

Commercial Colocation

Capacity, growth, operator density,
enterprise depth

B 30%

Hyperscale Cloud & Self-Build

Physical capacity, region depth,
self-build, investment

C 15%

AI Market Maturity

Live capacity, committed pipeline,
demand diversity, power runway

D 15%

Connectivity & Network Gravity

Subsea, peering, facility density,
route diversity

SECTION 03

The framework

Every market is scored 0–5 on 16 metrics, four per axis, against fixed definitions. Axis scores are weighted averages of their metrics, and the four axes carry the weights shown below.

EXHIBIT 8 · AXIS WEIGHTS IN THE COMPOSITE SCORE



The four weights sum to 100%. Axis A carries the largest share because commercial colocation remains the best single description of a functioning market; Axis C carries 15%, which is why AI capacity alone cannot lift a market with no commercial base out of Tier 4.

EXHIBIT 9 · THE 16 METRICS AND THEIR INTRA-AXIS WEIGHTS

Axis A40%			Axis B30%			Axis C15%			Axis D15%		
A1	Operational colocation capacity	30%	B1	Hyperscalers with physical capacity	45%	C1	Operational AI capacity	30%	D1	Subsea cable systems	25%
A2	Committed colocation growth	20%	B2	Cloud region depth	15%	C2	Committed AI pipeline	30%	D2	Peering depth	30%
A3	Competitive operator density	25%	B3	Self-build commitment	30%	C3	AI demand diversity	20%	D3	Facility network density	25%
A4	Enterprise workload depth	25%	B4	Committed hyperscaler investment	10%	C4	Power runway for AI-scale loads	20%	D4	Route diversity and openness	20%

11

MEASURED DIRECTLY ON
STRUCTURE RESEARCH DATA

4

ANCHORED ON AUDITABLE
THIRD-PARTY SOURCES

1

SCORED ON REGULATOR
AND UTILITY EVIDENCE

Of the 16 metrics, 11 are measured directly on Structure Research data; power runway is scored against regulator and utility evidence, and the four connectivity metrics are anchored on auditable third-party sources so that axis can be reproduced independently by any reader.

A

Commercial Colocation

Measured on the Structure Research Data Centre Model, on the current 2020–2031 series. Four metrics: installed capacity, committed growth, operator density and enterprise workload depth.

40%
OF COMPOSITE

A1

30% OF
THE AXIS

Operational colocation capacity

Installed commercial colocation megawatts in the market on the current Structure Research Data Centre Model cycle, counting third-party capacity that is built and sellable today. The bands are logarithmic rather than linear — under 10MW at the bottom, 700MW and above at the top — because market maturity scales sub-linearly with megawatts, and the difference between 20MW and 60MW says far more about a market than the difference between 620MW and 660MW.

Hyperscale self-build campuses are not counted here at any size, because they are not available to a third party; they belong to Axis B. Capacity under construction, permitted or announced is also excluded, and is picked up by A2.

A1 and both scale-gate MW tests use operational commercial-colocation capacity as at 30 June 2026 — H1 2026 actuals on the current model cycle.

A2

20% OF
THE AXIS


Committed colocation growth

Absolute megawatts of commercial colocation added between 2026 and 2031 on the same Structure Research series, taken as the difference between the two endpoint years. The measure is deliberately absolute rather than percentage, because percentage growth overstates small bases — a market going from 2MW to 20MW posts 900% growth and adds almost nothing usable, while one adding 400MW off a large base changes the regional supply picture.

Bands run from under 10MW at the bottom to 1,000MW and above at the top. Seven emerging markets that the Data Centre Model does not break out separately carry Structure Research estimates rather than modelled figures, and are identified as such.

EXHIBIT 10 · WHY THE A1 BANDS ARE LOGARITHMIC

The six A1 score bands, plotted on a logarithmic megawatt axis. Equal score steps span very unequal absolute ranges by design.



Logarithmic megawatt axis, 1MW to 2,000MW. Source: Structure Research Market Tiering model rubric, August 2026.

A3

25% OF
THE AXIS

Competitive operator density

The count of distinct commercial operators with operational capacity in the market, from Structure Research facility-level coverage, running from none at the bottom to 12 or more at the top. The count is of operators at scale rather than of every registered entity, so a market with a long tail of sub-megawatt providers is not credited for them; the raw bench, the scored band and the 15 markets where the two differ are all recorded.

State and incumbent telcos are excluded by design: they hold capacity in almost every market on earth, so counting them adds a constant to every score and discriminates between nothing. Operators with only announced or under-construction capacity do not count until something is operating, and multiple facilities held by one operator count once.

The metric measures whether a buyer has genuine alternatives in the market, which is a different question from how many megawatts the market holds — several markets in this set carry large capacity behind a single developer and score zero here.

A4

25% OF
THE AXIS

Enterprise workload depth

An assessment of the enterprise demand base underneath the market: regulated financial services, government and sovereign workloads, regional headquarters and latency-sensitive trading ecosystems, assessed from Structure Research market coverage. The lowest score means no enterprise base at all and the highest means a market that serves as a regional enterprise capital. The steps in between measure how many industries a market serves and how deeply.

Network and content workloads are deliberately not counted here — they were moved to Axis D, because a heavy CDN or carrier presence describes a market's connectivity, not its enterprise depth, and counting it in both places would double-weight the same fact. Hyperscaler demand is likewise excluded, since it is measured in full on Axis B.

EXHIBIT 11 · SCORING BANDS FOR AXIS A

METRIC	0	1	2	3	4	5
A1 Colocation	<10MW	10 to <50MW	50 to <100MW	100 to <300MW	300 to <700MW	≥700MW
A2 Growth	<10MW	10 to <50MW	50 to <150MW	150 to <400MW	400 to <1,000MW	≥1,000MW
A3 Operators	0	1–2	3–4	5–7	8–11	≥12
A4 Enterprise	None	Single niche vertical	Emerging base	Meaningful enterprise demand	Deep multi-vertical enterprise	Regional enterprise capital

Rubrics are global across regions. Eleven documented A1 analyst overlays apply in this model cycle. Detailed source values, mechanical scores, final scores and analyst rationales are maintained in the underlying research model. Source: Structure Research Market Tiering model, August 2026.

EXCLUDED FROM AXIS A BY DESIGN

Hyperscale self-build

Not available to a third party at any size. Measured in full on Axis B.

Pipeline capacity

Under construction, permitted or announced. Picked up by A2, never by A1.

State and incumbent telcos

Present almost everywhere, so counting them discriminates between nothing.

Network and content demand

Describes connectivity, not enterprise depth. Scored on Axis D instead.

Hyperscale Cloud & Self-Build

B

Drawn on two Structure Research proprietary datasets — the deep-dive DCAI market models, which carry leased hyperscaler capacity as well as self-build, and the Hyperscale Self-Build Data Centre Model — and, for cloud region depth, on the infrastructure documentation the cloud platforms publish themselves. Hyperscaler presence is quality-weighted rather than counted flat.

30%
OF COMPOSITE

B1

45% OF
THE AXIS

LARGEST
SINGLE METRIC
IN THE
FRAMEWORK

Hyperscalers with physical capacity

A quality-weighted score for the hyperscalers holding capacity physically in the market, leased or self-build, taken from the Structure Research deep-dive DCAI market models. Operators are not counted flat but weighted by the depth of ecosystem each pulls with it. Points are then totalled and banded, running from none at the bottom to 20 or more at the top — a total that a market anchored by Western clouds alone cannot reach. It carries the largest weight of any single metric in the framework, because it is the hardest evidence available: an operating presence or a signed contract, not a forecast.

Sales offices, edge caches and CDN points of presence do not count, and an unnamed residual entry in a market model is not counted as an operator. The count is therefore a minimum rather than a guess: it can understate a market but will not overstate one. Where no deep-dive DCAI market model exists the count is an analyst judgement rather than a measurement, and those markets are marked as such.

Contracted-but-unbuilt positions score at half credit, so a signed deal never counts for as much as capacity already running. They qualify on the same test D1 applies to financed subsea systems under construction: a named operator, a stated capacity or capital figure and a signed lease, financing close or started construction.

EXHIBIT 12 · THE EVIDENCE LADDER

One discipline runs through every axis. Construction status and capital at risk are scored; announcements are not.

1.0x FULL CREDIT

Capacity operating in the market today. An operating presence, energised and serving workloads.

0.5x HALF CREDIT

Contracted but unbuilt: a named operator, a stated capacity or capital figure, and a signed lease, financing close or started construction.

0x CARRIED AT ZERO

Announcements, memoranda of understanding, land options, pre-construction without a contract, suspended programmes. All can be withdrawn at no cost, and many are.

Controlled land means land owned or held under a binding long-term agreement. A revocable option or non-binding site announcement scores zero. For B4, a government memorandum counts only when it is binding, location-specific and states a quantified capital commitment; a non-binding MoU scores zero.

The same rule governs subsea systems, self-build ladders and investment commitments, and it is the main reason this classification diverges from headline-driven rankings. Source: Structure Research Market Tiering model, August 2026.

B2

15% OF
THE AXIS

Cloud region depth

Not how many cloud platforms are present in the market, which B1 already measures, but how deep the infrastructure they built there actually is. The unit is the availability-zone region: three or more physically separated zones inside the metro, which is what lets a customer run production workloads with failover that stays in the market. None of a cache node, a local zone or a single-zone region can support that, so all three score below any full region. The metric counts distinct platforms rather than regions, so a platform running two regions in one metro counts once and cannot inflate the score.

The scale runs from no cloud infrastructure at all, through edge presence only, to a single region that is not zone-redundant, to a market carrying either one full availability-zone region or two or more live regions that are not zone-redundant, and finally to two and then three or more platforms with full regions at the top. The top two bands are reachable only on full availability-zone regions, so a market with several thin regions cannot out-score a market with genuine multi-platform depth.

The metric is nationality-blind by design: an Alibaba or Huawei region counts on exactly the same test as an AWS or Azure one, because the question is what was built and not who built it. Every region is classified from what the platform publishes about its own infrastructure — region lists, zone counts and availability-zone support flags, including the regions that are present but do not clear the zone-redundancy line. Where a platform does not publish a zone count, its region is treated as not zone-redundant. That is a deliberately conservative call, consistent with the rest of the framework: the metric will understate a platform before it overstates one.

EXHIBIT 13 · THE B2 LADDER

0	No cloud infrastructure at all
1	Edge presence only — local zone, extended zone or cache PoP; no region
2	One region, not zone-redundant
3	One platform with a full availability-zone region, or two or more platforms with regions that are not zone-redundant
4	Two platforms with full availability-zone regions
5	Three or more platforms with full availability-zone regions

The top two bands are reachable only on full availability-zone regions. Frankfurt, Singapore, Jakarta and Hong Kong each carry six platforms running full regions — the joint-highest in the classification. Source: provider infrastructure documentation compiled by Structure Research, August 2026.

THE INDEPENDENT AI CLOUDS ARE MEASURED ON THE SAME TEST

CoreWeave, Nebius, Crusoe and Lambda all run multi-region public clouds, and none of them operates a full availability-zone region anywhere: CoreWeave documents availability zones but no region carries more than two, and the other three use no zone construct at all. Their regions therefore score as regions without zone redundancy and cannot lift a market into the top two bands.

Providers that publish no region identifiers are not counted, because there is nothing to verify. Only capacity a customer can actually buy counts, so CoreWeave regions reserved for selected customers and Nebius private regions score zero — on the same test that governs operational colocation capacity and operator density.

B330% OF
THE AXISRAISED FROM
20% THIS
CYCLE**Self-build commitment**

Operational hyperscale self-build megawatts and the number of operators holding them, measured on the Structure Research Hyperscale Self-Build Data Centre Model, with the top band requiring three or more operators and 400MW or more in operation. For scores 3–5, both the operational-MW threshold and the operator-count threshold apply, so a single operator with a very large campus cannot reach the top — the metric is measuring committed breadth, not one company's decision. Scores 1–2 represent land-bank or construction-only commitment and do not require operational capacity.

Controlled land banks score in the lowest positive band however large, and capacity that is only under construction scores no higher than the band below 40MW operational, because land and construction are commitments rather than capacity. Leased capacity is excluded here in full and is captured by B1 instead, as is developer-built build-to-suit, which is a lease with extra steps rather than a self-build position.

Its weight rose from 20% to 30% because it is measured independently of the operator lists that drive B1, and correlates least with the rest of the axis.

B410% OF
THE AXISSOFTEST
EVIDENCE IN
THE
FRAMEWORK**Committed hyperscaler investment**

Verified public investment commitments attributable to the market, in US dollars, tracked by Structure Research against primary disclosure and banded from under \$0.5bn to \$15bn and above. A commitment counts when it is stated by the operator or a government counterparty and tied to a named location — a signed memorandum with a state government, a disclosed capital programme, a permitted campus with a published figure.

Aggregate country-level or multi-year global capex numbers are not apportioned across markets, and press estimates of unannounced spending are not used. This metric carries the smallest weight on the axis at 10%, because announced money is the softest evidence in the framework: it is revised, restated and occasionally withdrawn, and the same dollars are frequently announced more than once.

EXHIBIT 14 · B1 OPERATOR WEIGHTING

WEIGHTING TIER	OPERATORS	POINTS
Anchor clouds	AWS, Microsoft	4
Frontier AI labs	OpenAI, Anthropic, xAI	4
Major cloud	Google	3
Asian hyperscalers	Alibaba, Tencent, Huawei, ByteDance, NAVER, Kakao, Sea Group and others	2
Other Western platforms	Oracle, Meta, IBM, Apple	1
Neoclouds	CoreWeave, Nebius, Crusoe, Nscale, Firmus, Lambda, Yotta, Sharon AI and others	1

Western platforms are split across three tiers by the depth of ecosystem each carries; Asian hyperscalers are not ranked among themselves and each scores 2. A contracted position scores half the points shown; an operator both operating and contracting in one market counts once, at full credit.

Source: Structure Research Market Tiering model, August 2026.

EXHIBIT 15 · HOW B2 CLASSIFIES EACH PLATFORM'S REGIONS

Depth is read from what the platform publishes about its own infrastructure, so the metric can be reproduced independently. Every region is recorded by name against its source.

PLATFORM	HOW REGION DEPTH IS ESTABLISHED
Amazon Web Services	Every commercial region carries three or more availability zones and counts as a full region. Local Zones and Wavelength Zones are edge presence.
Microsoft Azure	Microsoft publishes an availability-zone support flag per region rather than a zone count. Flagged regions count as full; the rest are regions without zone redundancy.
Google Cloud	Every launched region carries three or more zones and counts as a full region. CDN, Media CDN and interconnect locations are edge presence.
Oracle OCI	Graded on availability domains. Only five OCI regions worldwide carry three; Oracle states that it launches new geographies with a single domain.
Alibaba Cloud, Tencent Cloud	Graded on the zone count each publishes per region. Three or more counts as a full region; one or two does not.
Huawei Cloud	Graded on per-region launch releases and Huawei's published global-infrastructure dataset. Four regions whose count is unpublished or internally contradicted are treated as not zone-redundant.
IBM Cloud	Multizone regions count as full. Single-campus multizone regions and classic data centres do not.
CoreWeave	The only AI cloud that documents availability zones as independent failure domains, but no region carries more than two and most carry one, so every region counts as a region without zone redundancy. Regions reserved for selected customers score zero.
Nebius, Crusoe, Lambda	None publishes a zone construct — the region is the smallest unit each of them exposes — so every region counts as a region without zone redundancy. Nebius private regions score zero.
Not counted	AI clouds that publish no region identifiers: Nscale, Firmus, Yotta, Sharon AI, Together AI and Fluidstack. Where there is no published region list, there is nothing a reader can verify, so the metric scores zero rather than relying on an analyst estimate.

Source: provider infrastructure documentation — region lists, availability-zone counts and availability-zone support flags — compiled by Structure Research in August 2026 and cross-checked against 19 Structure Research DCAI Market reports.

EXHIBIT 16 · SCORING BANDS FOR AXIS B

METRIC	0	1	2	3	4	5
B1 Hyperscalers	0 pts	1–3 pts	4–7 pts	8–12 pts	13–19 pts	≥20 pts
B2 Cloud region depth	No cloud infrastructure	Edge presence only	1 region, not zone-redundant	1 platform with a full AZ region, or 2+ non-redundant	2 platforms with full AZ regions	3+ platforms with full AZ regions
B3 Self-build	No footprint	Land bank only	UC only, or <40MW operational	1+ operator, ≥40MW operational	2+ operators, ≥150MW operational	3+ operators, ≥400MW operational
B4 Investment US\$bn	<0.5	0.5–1	1–3	3–8	8–15	≥15

B1 bands are point totals from the operator weighting. B2 counts distinct platforms, and its top two bands are reachable only on full availability-zone regions. Source: Structure Research Market Tiering model, August 2026.

C

AI Market Maturity

Measured on Structure Research AI deployment tracking, with power runway scored against regulator, utility and system-operator evidence. Live capacity and forward pipeline are held apart and never counted in both.

15%

OF COMPOSITE

C1

30% OF
THE AXIS

Operational AI capacity

Accelerated-compute capacity that is installed, energised and running in the market today, measured in megawatts from Structure Research deployment tracking. The test is whether the hardware is drawing power and serving workloads: a pilot below 5MW sits in the bottom band, and 150MW or more scores at the top.

Capacity that is ordered, financed, under construction or merely announced is excluded from C1 without exception — it is scored in C2 instead, and never in both. Holding the two apart matters because a market's live AI base and its forward pipeline routinely diverge by an order of magnitude, and a blended figure would let a construction site read as a deployment.

C2

30% OF
THE AXIS

Committed AI pipeline

Accelerated-compute capacity that is under construction or financed but not yet operational, on the same Structure Research basis. Capacity qualifies once construction has physically started or capital is contractually committed — a signed tenant, a financing close or a groundbreak with money at risk.

Announced programmes, memoranda of understanding, land options and gigawatt ambitions without a tenant, a capital commitment or a power contract score zero here, however widely reported. Bands run from under 20MW at the lowest positive band to 400MW and above at the top, so a market carrying a very large but purely announced pipeline scores nothing on this metric while a smaller financed one scores well.

C3

20% OF
THE AXIS

AI demand diversity

The number of distinct classes of AI demand physically present in the market, rather than the volume of any one of them. Five classes are recognised. Scoring is one point per class present, so a market holding a single class scores 1 and one holding all five scores 5. Breadth is rewarded over concentration, and a very large single-buyer market can score below a smaller diversified one.

The reasoning is the same one that runs through the framework: a market whose entire AI base is one buyer class is exposed to that class's capital cycle, export-control position and procurement decisions in a way a diversified market is not.

EXHIBIT 17 · THE FIVE CLASSES COUNTED IN C3

A class counts once however many of its members are present, and only where capacity is physically deployed in the market.

01	02	03	04	05
Western hyperscalers	Asian hyperscalers	Nvidia-partner neoclouds	Sovereign or state-backed	Domestic platform operators
AWS, Microsoft, Google, Oracle, Meta	Alibaba, Tencent, Huawei, ByteDance	CoreWeave, Nebius, Firmus, Nscale, Sharon AI	National AI computing centres and state-funded campuses	NAVER, Kakao, Yotta, GMO

Scoring is one point per class present. Source: Structure Research AI deployment tracking, August 2026.

C4

20% OF
THE AXIS

Power runway for AI-scale loads

An assessment of whether the market can energise the next tranche of large load, scored on evidenced grid policy and physics rather than on developer statements or utility marketing. The scale runs from a formal moratorium or hard connection cap at the bottom, through documented refusal regimes and multi-year grid interconnection queues, to adequate headroom and then to structural surplus or dedicated behind-the-meter generation at the top.

Evidence is taken from regulators, utilities and system operators — published connection policies, queue disclosures, tariff rulings and generation positions — and announced generation projects do not count until they are contracted.

What this metric deliberately does not measure is how much capacity a market has already built: installed base is Axis A's job, and on current evidence the two point in opposite directions more often than not.

EXHIBIT 18 · SCORING BANDS FOR AXIS C

METRIC	0	1	2	3	4	5
C1 AI live	None	Pilot <5MW	5–20MW	20–60MW	60–150MW	≥150MW
C2 AI pipeline	None	<20MW	20–60MW	60–150MW	150–400MW	≥400MW
C3 AI classes	0 classes	1 class	2 classes	3 classes	4 classes	All five classes
C4 Power runway	Moratorium	Refusals	Long queues	Adequate	Strong	Surplus / dedicated

C3 counts distinct classes of AI demand present, not megawatts — one point per class, from none to all five. Source: Structure Research Market Tiering model, August 2026.

EXHIBIT 19 · STRONGEST AI AXIS SCORES IN THE CLASSIFICATION

Axis C carries 15% of the composite, which is why a strong AI profile alone cannot lift a market with no commercial base. Sort on Axis C for an AI siting view.

West Texas	Tier 3		4.60
Abu Dhabi	Tier 2		4.40
Johor	Tier 1		4.40
Atlanta	Tier 1		4.30
Bangkok	Tier 2		4.20
Riyadh	Tier 2		4.10
Paris	Tier 1		4.00
Phoenix	Tier 1		3.90

Bars are drawn on the 0–5 axis scale. Four of the eight strongest AI markets classify Tier 2 or Tier 3 on the composite: West Texas, Abu Dhabi, Bangkok and Riyadh. Source: Structure Research Market Tiering model, August 2026.

D

Connectivity & Network Gravity

Deliberately anchored on auditable third-party sources rather than on Structure Research data, so that the connectivity axis can be reproduced independently by any reader. Subsea credit is graded by distance from the metro.

15%

OF COMPOSITE

D1

25% OF
THE AXISSOURCE
TeleGeography

Subsea cable systems

In-service international submarine systems credited by how far their landing station sits from the metro. A landing within 100km scores a full point; one between 100 and 250km scores half, and only where a dedicated backhaul route into the metro is documented; beyond 250km the system scores nothing here and the market takes that credit through route diversity in D4 instead.

The two thresholds are the same distances the framework already uses to judge hub adjacency, reused rather than invented, and they replace an earlier rule that let some metros count landings hundreds of kilometres away while denying the same treatment to others.

Financed systems under construction still score half and the two discounts compound, so a financed system landing 150km out contributes a quarter point; announced systems score zero. Domestic and festoon systems that carry no international traffic are never counted, and a system with two landings serving the same metro counts once, at the better rate. Bands are point totals rather than system counts, running from none to 18 or more.

D2

30% OF
THE AXISSOURCE
PeeringDB

Peering depth

The number of connected networks at the metro's largest internet exchange, taken from PeeringDB and banded from no exchange at all to more than 1,000 networks. Membership is used rather than traffic because membership is community-maintained, consistently published and hard to game, whereas exchange traffic figures are inconsistently reported and self-serving where they are.

Only the single largest exchange in the metro is counted, not the sum of all exchanges, so a market with several small exchanges does not out-score one with a single deep one. PeeringDB relies on self-reporting and undercounts closed or state-controlled markets, so scores in those markets are treated as floors and flagged as such rather than presented as complete.

EXHIBIT 20 · THE D1 DISTANCE RULE

Subsea credit is graded by the distance from the metro to the landing station, and the construction discount compounds with it.

LANDING WITHIN 100 KM	100 – 250 KM	BEYOND 250 KM
1.0 point per in-service system 	0.5 and only with documented backhaul 	0 credit taken through D4 instead 
Financed and under construction: 0.5 points.	Financed and under construction: 0.25 points — the discounts compound.	Announced systems score zero at any distance.

Domestic and festoon systems carrying no international traffic are never counted. A system with two landings serving the same metro counts once, at the better rate. Sources: TeleGeography, submarinenetworks.com.

D325% OF
THE AXISSOURCE
PeeringDB**Facility network density**

The number of connected networks at the metro's most network-dense carrier-neutral facility, also from PeeringDB, running from none to more than 250 networks. This measures where interconnection actually happens — the carrier hotel or campus that a new entrant has to be in to reach everyone else — which is a different property from exchange membership and frequently sits in a different building.

Operator-owned and single-tenant facilities are excluded because they are not open to third parties, however large they are. The metric will read low in markets where interconnection is deliberately distributed rather than concentrated, and that is the intended behaviour: a market with no dominant meet-me point is harder to enter.

D420% OF
THE AXISTHE ONLY
METRIC THAT
MARKS A
MARKET DOWN
FOR ITS
STRUCTURE**Route diversity and openness**

A rubric score built on three things: how many independent routes carry traffic in and out of the metro, how many cloud on-ramps the market has, and whether the market depends on a single point of failure or a single provider. A route counts as independent only where it is physically separate — a different landing area, a different border crossing, a different overland path — so several cables sharing one trench count once. The top score requires five or more independent routes, five or more cloud on-ramps and no single point through which all traffic must pass.

A market scores 1 in any of three cases: it has only one route in and out; traffic passes through it on the way somewhere else rather than terminating there, so the market is a waypoint and not a destination; or one carrier controls transit and can therefore set the price of access or withhold it. The bottom score describes a market with no meaningful route out at all.

This is the metric that credits inland markets with genuine overland route diversity, which D1 cannot see. It is also the only place the framework marks a market down for its structure rather than counting what it has — a market with plentiful capacity that all runs through one strait, one border or one incumbent is penalised for it.

EXHIBIT 21 · SCORING BANDS FOR AXIS D

METRIC	0	1	2	3	4	5
D1 Subsea cable	0 pts	1–2 pts	3–5 pts	6–9 pts	10–17 pts	≥18 pts
D2 Peering depth	No IX	<50	50–150	150–400	400–1,000	>1,000
D3 Network density	None	<20	20–50	50–100	100–250	>250
D4 Route diversity	Isolated	Single corridor, pass-through or monopoly	2 corridors or chokepoint-heavy	3 corridors, moderate on-ramps	4+ corridors, strong on-ramps	5+ corridors, 5+ on-ramps, no chokepoint

D1 bands are point totals under the distance rule, not raw system counts; a market whose landings all sit beyond 250km scores 0 and takes that credit in D4. Sources: TeleGeography, submarinenetworks.com, PeeringDB, provider technical documentation.

METHODOLOGY · HUB-ADJACENCY CREDIT

Credit applies only where the qualifying hub's own Axis D score is at least 3.50 and a documented route connects the markets. A satellite within 100km receives 50% of the positive Axis D gap to the hub; one between 100km and 250km receives 25%; beyond 250km it receives no adjacency credit.

Credit = coefficient × max(0, hub Axis D – satellite Axis D)

SATELLITE	HUB	DISTANCE	COEF.	CREDIT
Johor	Singapore	<100km	50%	+1.75
Batam	Singapore	<100km	50%	+1.25
Cardiff / South Wales	London	100–250km	25%	+0.6875
Brussels	Amsterdam	100–250km	25%	+0.3875
Pune	Mumbai	100–250km	25%	+0.8125

Five applications in this model cycle. Published Axis D scores are inclusive of the credit shown. Marseille and Lisbon take no adjacency credit: both are significant landing points, and landing points that have not become interconnection ecosystems score on exchange and facility density rather than on cable count. Source: Structure Research Market Tiering model, August 2026.

CLASSIFICATION RULES

Tiers and scale gates

The Composite Score maps to a provisional tier through four fixed thresholds. Two tier-specific gates then apply: one for Tier 1 and one for Tier 2. Both conditions in the applicable gate must be met. A market that fails either condition cascades to the highest lower tier whose conditions it meets. Gate applications and emerging-entrant status are recorded market by market.

COMPOSITE THRESHOLDS

T1	Composite Score of 3.50 and above
T2	2.50 to 3.49
T3	1.50 to 2.49
T4	Below 1.50

SCALE GATES — BOTH CONDITIONS MUST BE MET

TIER 1 GATE 240MW operational colocation capacity, at minimum	≥ 3 B1 hyperscaler-presence score, at minimum
TIER 2 GATE 50MW operational colocation capacity, at minimum	≥ 1 B1 hyperscaler-presence score, at minimum


EXHIBIT 22 · THE GATE IN PRACTICE — BANGKOK

A tier label describes what a market *is*; the Composite Score also contains what it is becoming.

COMPOSITE SCORE · CLEARS THE TIER 1 THRESHOLD



OPERATIONAL COLOCATION · FAILS THE TIER 1 GATE



Bangkok posts a Tier 1 Composite Score of 3.64 on committed capital and operational hyperscaler capacity, but holds 86MW of operational colocation, and is classified Tier 2 on the gate. That distance between Composite Score and tier is informative in its own right: it identifies the clearest promotion candidate in the set.

<25MW

EMERGING ENTRANT THRESHOLD

Separately, markets with under 25MW of operational colocation are classified as emerging entrants. That status is orthogonal to tier: a market can be an emerging entrant and score Tier 3, or be established and score Tier 4.

CLASSIFICATION RULES

Announced versus committed

One discipline runs through every axis. Construction status and capital at risk are scored; announcements are not.

A financed campus under construction counts, a signed power allocation counts and a groundbroken programme counts. A multi-gigawatt ambition without a tenant, capital or power contract is carried at zero, however widely reported.

The same rule governs subsea systems, self-build ladders and investment commitments, and it is the main reason this classification diverges from headline-driven rankings.

EXHIBIT 23 · HOW THE RULE APPLIES, AXIS BY AXIS

METRIC	SCORED AT FULL CREDIT	SCORED AT HALF CREDIT	CARRIED AT ZERO
B1 Hyperscalers	Capacity operating in the market, leased or self-build	Contracted but unbuilt — named operator, stated figure, signed lease or financing close	Announcements, MoUs, land options, pre-construction without a contract, suspended programmes
B3 Self-build	Operational self-build megawatts, banded with an operator-count condition	Under construction scores no higher than the band below 40MW operational	No qualifying self-build capacity. Controlled land is the lowest-positive-band exception described below.
B4 Investment	Verified commitment stated by the operator or a government counterparty and tied to a named location	—	Aggregate country-level or global capex; press estimates of unannounced spending
C1 AI live	Hardware drawing power and serving workloads today	—	Ordered, financed, under construction or announced — scored in C2 instead, never in both
C2 AI pipeline	Construction physically started, or capital contractually committed	—	Announced programmes, MoUs, land options, gigawatt ambitions with no tenant, capital or power contract
D1 Subsea	In-service system landing within 100km of the metro	Financed and under construction, and again for landings 100–250km out — the discounts compound	Announced systems; domestic and festoon systems carrying no international traffic

B3 exception: controlled land banks receive the lowest positive band; revocable options and non-binding site announcements score zero. For B4, a binding, location-specific memorandum that states a quantified capital commitment counts as a verified commitment; non-binding or unquantified MoUs score zero. Source: Structure Research Market Tiering model, August 2026.

EXHIBIT 24 · THE RULE AT THE BOTTOM OF THE TABLE

Five markets sit at the bottom of Tier 4: each has something real, and each falls short on the specific test the framework applies.

Tasmania 0.84 A hydro position and a consented site; short on tenancy and construction.	Sarawak 0.75 Hydro capacity and a programme; nothing operational, no committed tenant.	Jeollanam-do 0.56 A national programme at Solaseado; Axis A 0.20 and Axis D zero.	Toyama 0.35 A financed construction programme against no measured hyperscale position.	Southland 0.25 A consented site and a power position; Axis B and Axis D both zero.
---	--	---	--	--

Composite scores shown. Visakhapatnam, by contrast, scores 1.37 on groundbroken hyperscaler capital against a land bank with nothing yet built — the groundbreak counts, the land bank does not. Source: Structure Research Market Tiering model, August 2026.

04 SECTION FOUR · SCORES

RESULTS — 101 MARKETS SCORED

Scores are produced by the Structure Research Market Tiering model. Axis D is shown inclusive of the hub-adjacency credit: where a documented route connects a satellite to a hub scoring at least 3.50 on Axis D, the satellite takes 50% of the positive Axis D gap within 100km, or 25% within 100–250km, and nothing beyond. Five markets qualify — Johor, Batam, Cardiff, Brussels and Pune. The full rule is set out on page 25.

25

North America

6 · 10 · 8 · 1

36

EMEA

4 · 11 · 12 · 9

35

Asia Pacific

10 · 6 · 9 · 10

5

Latin America

1 · 2 · 1 · 1

101

Total

21 · 29 · 30 · 21

COUNTS SHOWN AS TIER 1 · TIER 2 · TIER 3 · TIER 4

SECTION 04

Global summary

EXHIBIT 25 · TIER COUNTS BY REGION

REGION	TIER 1	TIER 2	TIER 3	TIER 4	TOTAL	TIER 1 SHARE
North America	6	10	8	1	25	24%
EMEA	4	11	12	9	36	11%
Asia Pacific	10	6	9	10	35	29%
Latin America	1	2	1	1	5	20%
Total	21	29	30	21	101	21%

Asia Pacific holds the largest concentration of Tier 1 markets of any region and also the largest Tier 4 cohort. EMEA carries the widest spread, with 21 of its 36 markets in Tier 3 or Tier 4. Source: Structure Research Market Tiering model, August 2026.

EXHIBIT 26 · TIER COMPOSITION BY REGION

Each region normalised to 100%, so the shape of a region's distribution can be compared independently of how many markets it contains.





EXHIBIT 27 · COMPOSITE SCORE RANGE BY REGION

Lowest to highest composite in each region, with the regional median marked.



Published composite scores, ranges and medians are rounded to two decimal places using exact decimal, half-up rounding. Rankings, score-band placement and gate outcomes use unrounded composite scores. Exact ties receive the same competition rank.

RESULTS · REGION 01 OF 04

North America

25 MARKETS 6 TIER 1 10 TIER 2 8 TIER 3 1 TIER 4

MARKET	A	B	C	D	COMPOSITE SCORE	TIER
Northern Virginia (incl. ext.)	5.00	5.00	4.40	3.15	4.63	TIER 1
Dallas-Fort Worth	4.75	4.30	3.70	2.90	4.18	TIER 1
Phoenix	4.50	4.75	3.90	1.95	4.10	TIER 1
Atlanta	4.75	3.85	4.30	2.40	4.06	TIER 1
Chicago	5.00	3.60	3.10	3.15	4.02	TIER 1
Silicon Valley	4.55	3.25	2.20	3.20	3.61	TIER 1
Columbus (OH)	3.70	3.95	3.20	1.70	3.40	TIER 2
Austin-San Antonio	3.70	3.70	3.50	1.20	3.30	TIER 2
New York Metro	4.50	1.65	2.50	3.90	3.26	TIER 2
Toronto	3.80	3.30	1.90	2.95	3.24	TIER 2
Hillsboro	3.25	3.20	2.50	2.70	3.04	TIER 2
Montreal + Quebec City	3.35	3.35	2.00	1.95	2.94	TIER 2
Salt Lake City	3.45	2.75	3.00	1.20	2.84	TIER 2
Charlotte-Raleigh	3.15	3.55	2.30	1.00	2.82	TIER 2
Las Vegas	3.25	2.35	2.80	1.20	2.61	TIER 2
Reno / N. Nevada	2.65	2.75	3.50	0.75	2.52	TIER 2
West Texas (incl. Abilene)	2.50	2.45	4.60	0.20	2.46	TIER 3
Houston	3.45	1.40	2.60	1.50	2.42	TIER 3
Minneapolis	3.20	1.60	1.80	1.95	2.32	TIER 3
Los Angeles	3.05	0.90	0.70	3.90	2.18	TIER 3
Buffalo / Western NY	2.75	1.10	3.30	1.20	2.11	TIER 3
Denver	3.25	0.15	1.80	2.20	1.95	TIER 3
Calgary	2.25	0.90	2.40	1.50	1.76	TIER 3
Ellendale (ND)	2.00	0.75	3.60	0.00	1.57	TIER 3
Vancouver	1.75	0.15	0.40	2.20	1.14	TIER 4

Axis scores 0–5: **A** commercial colocation, **B** hyperscale cloud & self-build, **C** AI market maturity, **D** connectivity & network gravity. Axis cells are tinted by score; composite bars are drawn to a 0–5 scale. Markets are ordered by composite. Source: Structure Research Market Tiering model, August 2026.

RESULTS · REGION 02 OF 04

EMEA

36 4 11 12 9
MARKETS TIER 1 TIER 2 TIER 3 TIER 4

MARKET	A	B	C	D	COMPOSITE SCORE	TIER
London	5.00	4.40	3.00	4.20	4.40	TIER 1
Frankfurt	5.00	3.90	1.90	3.75	4.02	TIER 1
Paris	4.70	2.85	4.00	3.70	3.89	TIER 1
Milan	4.15	3.45	2.60	3.40	3.60	TIER 1
Stockholm	3.45	3.05	2.90	3.25	3.22	TIER 2
Amsterdam	3.75	3.05	1.00	4.20	3.20	TIER 2
Madrid	3.70	3.20	2.20	2.70	3.18	TIER 2
Dublin	3.25	4.15	0.40	2.50	2.98	TIER 2
Abu Dhabi	2.70	3.35	4.40	1.40	2.96	TIER 2
Tel Aviv / Israel	3.65	3.00	2.20	1.70	2.95	TIER 2
Johannesburg	3.00	2.60	2.10	3.05	2.75	TIER 2
Helsinki	2.65	2.45	3.60	2.75	2.75	TIER 2
Riyadh	2.40	2.80	4.10	0.95	2.56	TIER 2
Dubai	2.75	2.45	2.30	2.50	2.56	TIER 2
Zurich	3.25	2.10	1.20	2.70	2.52	TIER 2
Berlin	2.95	2.10	2.10	2.00	2.43	TIER 3
Moscow	3.05	1.65	2.10	2.35	2.38	TIER 3
Cape Town	1.80	2.80	2.10	3.25	2.36	TIER 3
Warsaw	2.50	2.30	1.60	2.55	2.31	TIER 3
Oslo	2.55	1.35	2.40	2.20	2.12	TIER 3
Lisbon + Sines	2.10	1.00	3.70	2.75	2.11	TIER 3
Copenhagen	2.20	1.50	1.80	2.95	2.04	TIER 3
Cardiff / South Wales	1.70	2.25	2.40	2.14	2.04	TIER 3
Brussels	2.50	1.45	0.60	3.04	1.98	TIER 3
Barcelona	2.70	0.30	2.70	2.20	1.91	TIER 3
Vienna	2.00	1.60	0.60	2.70	1.78	TIER 3
Dammam	1.20	2.15	3.30	0.95	1.76	TIER 3
Zaragoza	0.80	2.30	2.90	0.20	1.48	TIER 4
Athens	1.20	1.40	1.80	1.75	1.43	TIER 4
Lagos	2.20	0.15	0.90	2.30	1.41	TIER 4
Marseille	1.20	0.90	0.90	3.40	1.40	TIER 4
Reykjavik	1.80	0.30	2.70	1.00	1.37	TIER 4
Jeddah	1.20	0.75	1.30	1.75	1.16	TIER 4
Nairobi	1.70	0.00	0.60	2.00	1.07	TIER 4
Cairo	0.90	0.75	1.40	1.50	1.02	TIER 4
Bilbao	1.40	0.00	2.20	0.70	1.00	TIER 4

Axis cells are tinted by score; composite bars are drawn to a 0–5 scale. Markets are ordered by composite. Source: Structure Research Market Tiering model, August 2026.

RESULTS · REGION 03 OF 04

Asia Pacific

35 10 6 9 10
MARKETS TIER 1 TIER 2 TIER 3 TIER 4

MARKET	A	B	C	D	COMPOSITE SCORE	TIER
Tokyo	5.00	5.00	3.00	4.15	4.57	TIER 1
Singapore	4.80	5.00	2.00	4.70	4.43	TIER 1
Sydney	4.75	4.25	3.30	3.90	4.26	TIER 1
Jakarta	4.45	4.30	3.20	4.00	4.15	TIER 1
Mumbai	5.00	3.30	3.30	4.00	4.09	TIER 1
Melbourne	4.35	4.00	3.50	2.45	3.83	TIER 1
Osaka	4.00	4.15	3.80	2.65	3.81	TIER 1
Seoul	4.30	4.30	2.50	2.40	3.75	TIER 1
Hong Kong	4.40	3.70	1.30	3.95	3.66	TIER 1
Bangkok	3.65	3.95	4.20	2.40	3.64	TIER 2 · GATE
Johor	3.75	3.35	4.40	2.95	3.61	TIER 1
Kuala Lumpur	3.65	3.85	2.40	2.45	3.34	TIER 2
Hyderabad	3.55	3.25	3.40	1.20	3.09	TIER 2
Taipei	3.00	3.10	2.30	3.30	2.97	TIER 2
Delhi NCR	3.60	1.35	3.00	2.30	2.64	TIER 2
Chennai	3.60	1.00	2.50	2.70	2.52	TIER 2
Auckland	2.65	2.40	1.70	2.45	2.40	TIER 3
Busan-Ulsan	2.25	2.40	3.10	1.40	2.30	TIER 3
Canberra	2.80	1.70	3.50	0.65	2.25	TIER 3
Manila	2.70	2.10	1.00	2.50	2.24	TIER 3
Batam	2.15	1.05	3.50	3.45	2.22	TIER 3
Pune	2.00	1.40	1.40	1.56	1.66	TIER 3
Ho Chi Minh City	2.40	0.40	2.50	1.25	1.64	TIER 3
Perth	2.00	0.60	1.50	2.40	1.57	TIER 3
Hanoi	2.25	0.60	2.00	0.95	1.52	TIER 3
Brisbane	2.25	0.15	1.60	2.00	1.49	TIER 4
Kyushu	1.70	0.85	2.60	1.00	1.48	TIER 4
Bengaluru	2.65	0.00	1.20	1.20	1.42	TIER 4
Hokkaido	1.20	1.10	3.20	0.75	1.40	TIER 4
Visakhapatnam	1.40	1.25	2.70	0.20	1.37	TIER 4
Tasmania	0.70	0.65	2.20	0.20	0.84	TIER 4
Sarawak	1.00	0.00	1.30	1.00	0.75	TIER 4
Jeollanam-do (Solaseado)	0.20	0.80	1.60	0.00	0.56	TIER 4
Toyama	0.50	0.00	0.80	0.20	0.35	TIER 4
Southland	0.20	0.00	1.10	0.00	0.25	TIER 4

Bangkok posts a Tier 1 composite of 3.64 and is held at Tier 2 by the 240MW operational-colocation gate. Axis cells are tinted by score. Source: Structure Research Market Tiering model, August 2026.

RESULTS · REGION 04 OF 04

Latin America


5 1 2 1 1
MARKETS TIER 1 TIER 2 TIER 3 TIER 4

MARKET	A	B	C	D	COMPOSITE SCORE	TIER
Sao Paulo	4.70	3.75	3.50	4.30	4.18	TIER 1
Queretaro	2.95	3.75	2.80	1.45	2.94	TIER 2
Santiago	3.00	3.25	2.10	2.55	2.87	TIER 2
Rio de Janeiro	2.95	0.75	2.50	3.55	2.31	TIER 3
Bogota	1.95	0.75	1.20	1.55	1.42	TIER 4

Sao Paulo carries the largest single-metro internet exchange in the world by traffic. Source: Structure Research Market Tiering model, August 2026.

EXHIBIT 28 · THE 20 HIGHEST COMPOSITE SCORES, ALL REGIONS

Stacked bars show each axis's weighted contribution to the composite, so the shape of a market's score is visible alongside its total.



Segment length is the axis score multiplied by its composite weight; the full bar width equals a composite of 5.00. Source: Structure Research Market Tiering model.

05

SECTION FIVE · COMMENTARY

REGIONAL ANALYSIS

What the scores describe, region by region: where the leaders sit, which markets are constrained by policy rather than demand, and which of the AI-created locations are backed by construction rather than by announcement.

North America

Northern Virginia leads at 4.63; both dedicated AI campuses classify Tier 3

EMEA

Four Tier 1 markets; the Gulf is the momentum cluster, its AI profile ahead of installed depth

Asia Pacific

Ten Tier 1 markets, the largest concentration of any region

Latin America

Sao Paulo is the region's single Tier 1 market at 4.18

SECTION 05 · REGION 01

North America

Northern Virginia leads North America and ranks first globally at 4.63. It scores 5.00 on Axis A and 5.00 on Axis B; power runway is its principal constraint, with Dominion facing connection requests several times its all-time system peak, substation timelines beyond three years and vacancy under 1%.

Atlanta at 4.06 has the strongest riser profile in the region: five hyperscalers, an operational frontier-training flagship, a two-gigawatt construction pipeline and the best large-market power position in the United States. It is also the sharpest illustration of the divergence described in section 2 — on 891MW of installed colocation it holds no cloud region of any kind, only an AWS Local Zone and a Google cache node.

Dallas-Fort Worth, Phoenix, Chicago and Silicon Valley complete the Tier 1 set. Chicago scores 4.02 on 1,288MW of measured operational self-build across three operators, against one of the weakest power runways in the American Tier 1 cohort. Silicon Valley’s 3.61 composite is supported by installed interconnection depth and three operators with operational self-build capacity, against effectively halted new grid

connections; six cloud platforms run regions there and only one, the oldest AWS region in the western United States, is zone-redundant.

Columbus and Austin-San Antonio classify Tier 2 despite shallow connectivity scores: both carry fibre-crossroads value but limited exchange, facility and on-ramp depth. Reno enters Tier 2 at 2.52 on a cloud position built entirely by the AI clouds — it holds no hyperscale cloud region, and CoreWeave and Crusoe both run regions in the metro. Hillsboro scores 3.04 and holds one of the deepest hyperscale self-build complexes in the world against a thin commercial colocation base and no cloud region in the metro.

The dedicated AI campuses both classify Tier 3. West Texas hosts the largest operational AI campus in the United States and a multi-gigawatt forecast, and scores 2.46 on no competitive colocation market, no interconnection, no hyperscale self-build capacity and no cloud infrastructure of any kind, against an operating OpenAI position and a contracted Microsoft one. Ellendale scores 1.57 on the same shape at smaller scale. Both are deployment sites of significant consequence, and the axis profile records that directly: high on AI, low on everything commercial.

Readers seeking an AI siting ranking should sort on **Axis C** rather than the composite.

WHERE THE REGION’S SCORES DIVERGE MOST

MARKET	COMPOSITE	THE DIVERGENCE THE AXIS PROFILE RECORDS
Atlanta	4.06 · Tier 1	891MW installed and one of the strongest US AI profiles, against no cloud region of any kind a customer can buy
New York Metro	3.26 · Tier 2	Axis A 4.50 and Axis D 3.90 against Axis B 1.65 — commercial and network depth without hyperscale commitment
Hillsboro	3.04 · Tier 2	One of the deepest hyperscale self-build complexes in the world on a thin commercial base and no cloud region
West Texas	2.46 · Tier 3	Axis C 4.60 against Axis D 0.20 — the largest operational US AI campus with no interconnection at all

SECTION 05 · REGION 02

EMEA

EMEA holds four Tier 1 markets. London leads at 4.40 on maximum installed colocation, Telehouse-anchored interconnection depth and Atlantic landings. Frankfurt scores 4.02 with the highest European hyperscaler count, six platforms running full availability-zone regions — the joint-highest in the classification, alongside Singapore, Jakarta and Hong Kong — the only Tencent region in Europe, and DE-CIX.

Paris scores 3.89 on Europe’s strongest AI profile — a large operational GB200 estate, a gigawatt-scale campus with a direct transmission partnership and grid-secured capacity — layered on a low-carbon grid at a fraction of German power prices. Milan enters Tier 1 at 3.60, carrying full AWS, Azure and Google regions on a 250MW colocation base.

Amsterdam at 3.20 and Dublin at 2.98 fall short for the same structural reason: deep installed ecosystems whose growth vectors are policy-constrained. Dublin holds the deepest measured self-build base in Europe at 733.6MW across four hyperscalers, against an AI axis of 0.40. Madrid scores 3.18 and Stockholm 3.22. Zurich enters Tier 2 at 2.52 on cloud infrastructure rather than scale: AWS, Azure and Google all run full availability-zone regions there on a 191MW colocation base.

The Gulf is the region’s momentum cluster. Strong AI profiles outpace installed commercial colocation depth, which limits the markets’ composite scores. None is gate-demoted; final tiers match composite-score bands. Abu Dhabi scores 2.96 with a Tier 1 AI axis of 4.40 on a large sovereign programme and associated power, against 126MW of installed colocation in a single-operator-dominated market and three cloud regions, none of them zone-redundant. Riyadh scores 2.56 with the largest committed AI pipeline outside the United States against 74MW installed. Tel Aviv scores 2.95, with a national grid-application freeze constraining its forward position.

Cape Town classifies Tier 3 at 2.36 on cable landings and interconnection density. Johannesburg scores 2.75, carrying three full availability-zone regions. Marseille and Lisbon illustrate the distinction the connectivity axis draws: both are significant landing points, and both score modestly because landing points that have not become interconnection ecosystems score on exchange and facility density rather than on cable count.

EXHIBIT 29 · THE GULF MOMENTUM CLUSTER

Four Gulf markets whose AI profiles outpace installed commercial colocation depth. None is gate-demoted; each final tier matches its composite-score band.

Abu Dhabi 4.40 2.70 AXIS C AXIS A 126MW installed in a single-operator-dominated market; three cloud regions, none zone-redundant.	Riyadh 4.10 2.40 AXIS C AXIS A The largest committed AI pipeline outside the United States, against 74MW installed.	Dammam 3.30 1.20 AXIS C AXIS A Tier 3 at 1.76; hyperscale commitment well ahead of the commercial base.	Dubai 2.30 2.75 AXIS C AXIS A Tier 2 at 2.56 on a more balanced profile than its neighbours.
--	---	---	--

Source: Structure Research Market Tiering model, August 2026.

SECTION 05 · REGION 03

Asia Pacific

Asia Pacific holds ten Tier 1 markets, the largest concentration of any region. Tokyo leads at 4.57 on maximum scores for installed colocation and hyperscale presence, with three hyperscalers holding operational self-build capacity alongside domestic platform self-build. Singapore follows at 4.43, combining the region's deepest connectivity profile with 434MW of operational hyperscaler self-build across four operators, against an AI axis of 2.00 reflecting administrative allocation of new capacity. Sydney scores 4.26 on seven hyperscalers with physical capacity and four platforms running full availability-zone regions.

Jakarta classifies Tier 1 at 4.15, the strongest composite in the region outside the top three, on roughly 13 competitive platforms across 113 facilities, five operational cloud regions and eight hyperscalers with physical capacity. Mumbai scores 4.09 and Seoul 3.75; Seoul rests on installed scale, eight cloud platforms with regions — four of them zone-redundant, including the domestic Kakao platform — and deep domestic platform self-build, with an AI axis of 2.50 that reflects capital-area grid constraint rather than demand. Osaka follows at 3.81 on the strength of its self-build and AI profile.

Johor classifies Tier 1 at 3.61 on 1,146MW operational, the region's largest committed growth pipeline, six hyperscalers leasing at scale and the deepest operational AI base in the region outside Japan. It holds no zone-redundant cloud region: its two cloud anchors are a single-domain Oracle region at Kulai and a two-zone Alibaba region. Its measured self-build is negligible — the market holds a large Microsoft land bank and one building under construction, while all operating hyperscaler capacity is leased — and its connectivity score rests substantially on engineered access to Singapore rather than on landings of its own.

Hong Kong classifies Tier 1 at 3.66 on 567MW installed, around 14 competitive operators, the region's second-deepest interconnection ecosystem and 13 hyperscalers with capacity, the deepest operator field in the classification. Its AI axis of 1.30 is the lowest of any Tier 1 market: Western hyperscaler expansion has been static since around 2020 and US advanced-GPU export controls apply to the territory.

Bangkok posts a Tier 1 composite of 3.64 and is held at Tier 2 by the scale gate on 86MW of operational colocation; it carries three full availability-zone regions, from AWS, Google and Huawei. Kuala Lumpur scores 3.34 on measured self-build of 66MW across two operators. Chennai enters Tier 2 at 2.52: it holds no zone-redundant region, but Azure South India and an IBM single-campus region together give it two live cloud anchors on a 215MW colocation base.

Manila classifies Tier 3 at 2.24 with three hyperscalers holding capacity — Alibaba, Huawei Cloud and an AWS Local Zone — all on leased colocation, and no committed hyperscaler capital attributable to the market; the Huawei region is a full three-zone deployment. Batam scores 2.22, with its connectivity carried by proximity to Singapore and its hyperscaler presence limited to one operator in 2025.

The AI-driven emerging markets stratify by evidence rather than ambition. Kyushu leads at 1.48 and Hokkaido scores 1.40, both carrying real AI capacity against almost no hyperscale presence. Visakhapatnam scores 1.37 on groundbroken hyperscaler capital against a land bank with nothing yet built. Tasmania, Jeollanam-do, Sarawak, Toyama and Southland sit at the bottom of Tier 4: each has a financed construction programme, a national programme, a hydro position or a consented site, and each falls short on tenancy, capital or construction status.

The AI-driven emerging markets stratify by **evidence rather than ambition**.

SECTION 05 · REGION 04

Latin America

Sao Paulo is the region’s single Tier 1 market at 4.18, on six hyperscalers with capacity including operational self-build AI halls, more than 12 competitive operators, five platforms running full availability-zone regions and the largest single-metro internet exchange in the world by traffic.

Queretaro at 2.94 and Santiago at 2.87 hold Tier 2 with real hyperscaler depth, each carrying full regions from three and two platforms respectively; Queretaro’s water and grid externalities are its principal structural constraint. Rio de Janeiro scores 2.31 and Bogota 1.42.

SECTION 06

Coverage and sources

Coverage

101 markets: 35 in Asia Pacific, 25 in North America, 36 in EMEA and 5 in Latin America. Market definitions follow the Structure Research Data Centre Model. Three consolidations apply where markets function as single hubs: Northern Virginia includes the NoVA Extension corridor, Busan and Ulsan are scored as a single southeastern Korean hub and Stockholm includes its hinterland.

Emerging AI markets are added under a two-part inclusion test — at least one project with committed capital, land and power, or operational AI-specific capacity today — which is why several markets with large announced programmes are tracked on a watchlist rather than scored.

Sources

Capacity from the Structure Research Data Centre Model, 2020–2031 series. Hyperscale leasing, self-build and land positions from Structure Research’s Hyperscale Self-Build Data Centre Model and from 16 deep-dive DCAI market models. Neocloud tenancies and contracted-but-unbuilt positions are carried separately, each recorded against a named operator with a stated capacity or capital figure and a source.

Cloud region depth from the infrastructure documentation each platform publishes — region lists, availability-zone counts and availability-zone support flags for AWS, Microsoft Azure, Google Cloud, Oracle,

Alibaba, Tencent, Huawei, IBM, NAVER, Kakao, CoreWeave, Nebius, Crusoe and Lambda — compiled in August 2026 and cross-checked against 19 Structure Research DCAI Market reports.

Subsea systems from TeleGeography and submarinenetworks.com. Interconnection and facility network density from PeeringDB. Cloud on-ramp counts from provider technical documentation. Investment commitments from verified public disclosures. Evidence reviewed through August 2026; operational commercial-colocation capacity for A1 and the MW gates is measured at 30 June 2026.

STRUCTURE RESEARCH
DATA

Data Centre Model (2020–2031) · Hyperscale Self-Build Data Centre Model · 16 deep-dive DCAI market models · AI deployment tracking

THIRD-PARTY,
AUDITABLE

TeleGeography · submarinenetworks.com · PeeringDB · provider technical documentation

REGULATOR AND
UTILITY

Published connection policies · queue disclosures · tariff rulings · generation positions

AS-OF DATE

Evidence reviewed through August 2026; operational commercial-colocation capacity for A1 and the MW gates is measured at 30 June 2026.



**STRUCTURE
RESEARCH**

Global perspective on the process of taking cloud to the ground.

Structure Research is an independent research and consulting firm covering the digital infrastructure value chain — data centres, hyperscale, cloud, edge and AI infrastructure. Clients receive a regular flow of global digital infrastructure intelligence, insight and opinion through syndicated publications, market share databases and analyst notes.

THIS REPORT

Global Data Centre Market Tiering 2026

A four-axis classification of 101 data centre markets

Edition · August 2026

RELATED COVERAGE

Data Centre Model, 2020–2031

Hyperscale Self-Build Data Centre Model

DCAI Market Reports · 19 markets

Market Share Reports

OFFICES

Toronto, Canada

Singapore

This is a public Structure Research report. Journalists and publications may quote its findings and reproduce individual charts or tables in editorial coverage with clear attribution to Structure Research and a link to datacentreresearch.com. The complete report may be downloaded from and linked to at that address. Commercial reuse, adaptation, resale or redistribution of the underlying dataset requires prior written permission from Structure Research.

Suggested citation: Structure Research. *Global Data Centre Market Tiering 2026: A Four-Axis Classification of 101 Data Centre Markets*. Version 1.0, August 2026. <https://datacentreresearch.com/>

Media, data and permissions enquiries: <https://www.structureresearch.net/contact-us/>