The new export rule trips at one fifth of Singapore's own licensing gate.

From 19 November 2026 an exporter needs a licence for a data centre it knows will serve a Chinese or Macau company with more than 12,000,000 aggregate total processing performance. That applies wherever the hall is. On Nvidia's B200 accelerator it is 0.601 megawatts of electrical load, a fifth of Singapore's own 3 megawatt gate.

Leila Haddad, AI Analyst, Information Technology, Mosaic Singularity Strategic Intelligence Division | 2026-11-18



Key points

• No licence exception reaches a covered data-centre tenancy, so the only route is an application filed inside 90 days of publication.

• The trigger converts to 334 Nvidia B200 accelerators and 0.601 megawatts of load, against Singapore's own 3 megawatt licensing gate for a hall.

• The diligence that cures the red flag is ownership or control and location, over any customer the hall serves, not a shareholding test.

• On the Bureau of Industry and Security's 62-day average the last reliable filing date is 17 December 2026, and a tenancy in a new hall has no window.

• Attest each tenancy and the hall aggregate before supply, and record a written file-or-exit decision per tenancy by 17 December 2026.

0.601 MW	Load at which a B200 hall crosses the American threshold
334	B200 accelerators that cross 12,000,000 aggregate TPP

$17.6m	Revenue from one covered tenancy-year at 70 per cent use
$1.76bn	Industry total at 70 per cent use, on BIS's own count of applications

The United States now requires a licence to provide remote access to advanced computing power to a company headquartered in China, Macau or a Country Group D:5 destination, anywhere in the world. The licence applies once a data centre serves more than 12,000,000 aggregate total processing performance, the measure the rule uses for computing power [s1]. The rule is published on 19 November 2026, and no licence exception reaches the control [s1].

On Nvidia's B200 accelerator that threshold is 334 parts, or 0.601 megawatts of system load [s4] [s5]. Singapore's own data-centre licence starts at 3 megawatts of critical load, five times larger [s8].

One covered tenancy is worth $25.2m to $26.0m a year at captured list prices [s9] [s10]. At 70 per cent use it is $17.6m to $18.2m. The industry total is unmeasured: on the Bureau of Industry and Security's count of about 100 additional applications a year across the whole rule it is $1.76bn to $1.82bn [s6].

Halls that can attest take tenancies their competitors must refuse. Operators that rent graphics processing units and hold Chinese-headquartered tenants lose the tenancy or the revenue, and the tenants lose the capacity.

What to do first. The head of global trade compliance should make an attestation a condition of supply and decide each tenancy in writing by 17 December 2026. That is the last filing date that clears on the Bureau of Industry and Security's published 62-day average [s6].

The bottom line. The bar is not an exception but a filing date. Attest the tenancy and the hall before supply, and assume that no licence arrives before the tenancy starts.

How the control works

The control sits in section 744.24(a)(1) of the Export Administration Regulations [s1]. It prohibits export, reexport and in-country transfer without a licence where the exporter knows the items will be used in a data centre serving more than 12,000,000 aggregate TPP [s1]. That applies at one centre or through linked data centres, to or for a company headquartered in Macau or a Country Group D:5 destination [s1]. Aggregate TPP is the sum of the processing performance of all the covered chips in the group [s1]. Linked data centres include digitally linked infrastructure, so capacity at separate sites counts towards the threshold [s1]. The control is additional to the licence requirements elsewhere in the rules, not a replacement for them [s1].

Three further limbs arrive with it. One licenses exports to five destinations, Indonesia, Malaysia, Singapore, Brazil and Thailand [s1]. One is the worldwide entity control for entities headquartered in Macau or a D:5 destination [s1]. The third reaches front-end fabricators and packaging houses shipping parts designed by companies that are neither approved nor authorised designers [s1].

There is no exception to reach for. The rule states that no licence exception may overcome the prohibition, except as noted when a destination is designated under paragraph (3) [s1]. Paragraph (3), read in full, is the transition for a centre already serving covered companies before publication: it must obtain a licence within 90 days of the effective date or face future export restrictions [s1]. It contains no destination designation, so on the face of the text that carve-out has no operative content. The transition is keyed to prior service, so a tenancy that starts after publication has no 90-day window at all [s1].

The diligence that decides the outcome is set by Note 3 to the same paragraph. A centre that declines or cannot supply sufficient information presumptively gives its supplier knowledge [s1]. To avoid that, the

supplier must obtain the centre's know-your-customer checks [s1]. Those checks must cover whether any data-centre customer is subject to direct or indirect ownership or control by persons located in, headquartered in, or whose ultimate parent is headquartered in Macau or a D:5 destination [s1]. Two features of that sentence decide what the register must hold. Control is not a shareholding test, and location is wider than headquarters. The preamble adds the population on the other side of the transaction: a non-D:5 customer that is in turn training models for, or providing models to, a D:5 company is inside the prohibition [s1].

What it means for the industry

The conversion is the most useful number here, and it runs the opposite way to the way a reader expects. On Nvidia's B200, 12,000,000 TPP is 334 accelerators and 0.601 MW. On AMD's MI355X at the top of its published range it is 199 accelerators and 0.358 MW; at the bottom of that range it is 300 accelerators and 0.543 MW. On H100 or H200 it is 758 accelerators and 1.359 MW [s4] [s5]. A United States control therefore bites at a fifth of the load at which Singapore licenses a hall [s8], and the part with the most processing performance per watt trips it at the lowest load of all.

Exhibit 1: Load at which a hall crosses the United States threshold, by accelerator type

	MW
B200, 36,000 TPP	0.601
MI355X, 60,600 TPP	0.358
MI355X, 40,000 TPP	0.543
H100 or H200, 15,832 TPP	1.359

One value per accelerator type, at 14.3 kW per eight-accelerator system. Singapore's own data-centre licence applies at 3 MW.

Source: Third-party TPP derivation [s4]; Nvidia DGX B200 system datasheet [s5]; Pinsent Masons on Singapore and Malaysia capacity [s8]; Mosaic Singularity Strategic Intelligence Division analysis

The gain goes to a hall that can attest. It can take tenancies its competitors must refuse, and charge for the certainty that carries. US-headquartered operators and tenants keep their route open by filing a monthly report [s1]. Colocation landlords gain wherever they do not themselves export or transfer the items. The loss falls on GPU-as-a-service operators with a concentrated roster of Chinese-headquartered tenants, which carry the licence obligation, the exit decision and the revenue. It falls on operators in the five destinations whose tenants are not US-headquartered, and on the tenants, who lose choice.

Exhibit 2: Four parties on the measures their own disclosures give

Company	Parts to cross the trigger	Disclosed measure, as the filing or rate card gives it	What it shows
Nvidia	B200: 334; H100 or H200: 758 [s4]	Revenue outside the United States 37.6 per cent of $96,221m in the quarter, designated by the headquarters of direct customers [s11]	The largest disclosed non-US base; the China line is not where the exposure sits
AMD	MI355X: 199 at 60,600 TPP, 300 at 40,000 [s4]	International sales 57 per cent of net revenue, three months to 26 September 2026 [s12]	The fewest parts to cross the trigger, so AMD-served halls are the smallest to become covered
CoreWeave	Not a vendor	Price: $8.60 per accelerator-hour on demand for an eight-accelerator B200 instance, $4.26 spot, $6.16 for H100 [s9]	Sets the revenue per tenancy: the committed rate is the base, the spot rate the downside

Company	Parts to cross the trigger	Disclosed measure, as the filing or rate card gives it	What it shows
Lambda	Not a vendor	Price: $8.87 per accelerator-hour at 256 or more accelerators on a two-week to one-year term [s10]	The top of the committed range, and the rate a large tenancy signs

Source: Company filings and published rate cards [s9] [s10] [s11] [s12]; third-party TPP derivation [s4]

Where the answer differs. By position, an operator holds the tenancy record and cannot pass the duty to its tenant, a chipmaker needs a per-product processing-performance field in its monthly report, and a system maker is the exporter of record that files [s1]. By scale, a neocloud with one or two non-US tenants decides by hand, while a hyperscaler turns dozens of tenancy records into a workflow. By jurisdiction, a tenant headquartered in Singapore with a Chinese ultimate parent is inside the control wherever it sits [s1]. Singapore's 3 MW regime is a separate and larger gate that a hall there must clear anyway [s8]. Outside Country Groups D:1, D:4 and D:5, and outside A:5 and A:6, the entity control adds a licence for D:5-headquartered entities [s1]. Allied destinations therefore do not carry that limb, and the country list that would enumerate the rest did not render in our capture [s16].

One qualification bounds the recommendation. The rule binds a person in possession of the items who was not party to the original shipment [s1], and no captured text decides whether an operator that never holds title is such a person. The operator's duty is therefore argued rather than settled, and the register is a commercial instrument for a party whose supplier must obtain the diligence.

What it is worth

A tenancy at the trigger is 334 B200 accelerators [s4]. At captured list rates of $8.60 to $8.87 per accelerator-hour over 8,760 hours, that is $25.2m to $26.0m a year [s9] [s10]. At 70 per cent use it is $17.6m to $18.2m, and at 50 per cent $12.6m to $13.0m. It is revenue on the income statement, recognised over the contract term; on an exit it reaches remaining performance obligations and, where prepayments were taken, cash already received.

A typical company. Take an operator with 30 MW of covered capacity, which is 50 units at the 0.601 MW trigger. At a labelled 10 per cent covered share that is five tenancies and $88m to $91m a year at 70 per cent use. At 30 per cent it is 15 tenancies and $264m to $273m. No captured source measures the covered share of any hall, so the reader supplies it.

The industry total is a construction, not a measurement. BIS expects the rule to add about 100 licence applications a year [s6]. That figure covers the whole rule, which creates at least four licence requirements, and no captured source allocates it between them. Putting all of it into covered tenancies, one per application, gives $1.76bn to $1.82bn a year at 70 per cent use: an upper construction, since one application can cover several halls. From installed capacity, Singapore's 1.4 GW is 2,331 trigger units [s8]; ten points of it let to covered tenants is 233 units and $4.11bn to $4.23bn. The two constructions differ by 2.3 times because the covered share is unmeasured, and the recommendation does not rest on either.

Your own figure. Revenue at risk is the accelerators a tenant draws, times your price per accelerator-hour, times 8,760 hours, times your utilisation. For all covered tenancies, multiply that by your covered capacity in MW divided by 0.601, and by the covered share.

The options and our recommendation

The options open to an operator

Option	Cost	Risk	What it forecloses
Attestation register, condition of supply, and a file-or-exit decision per tenancy (recommended)	The register, the contract changes and the monthly report, which BIS prices at $17.24 per submission [s6]. A $2m annual programme, illustrative, is 0.11 of one tenancy-year at 70 per cent use.	Exiting a paying tenant forgoes $17.6m to $18.2m a year, and is unnecessary if BIS narrows the control.	Nothing permanent: the register is an asset under any successor rule.
File for the tenancies worth keeping, assuming no licence is decided before the tenancy starts	The application work, plus the terms a grant carries on the only captured precedent: inspection before shipment, a 25 per cent tariff and a 15 per cent revenue expectation [s12].	The presumption of denial governs where a destination licence is also required, and no data-centre application has been decided [s1].	Re-letting at market while the application is pending. Live only where the tenant accepts a contingent start.
Defer, and underwrite another stay of the Affiliates Rule	Nothing now.	A stay governs the ownership screen only, and leaves the data-centre control, the destination licence, the monthly report and the 90-day deadline running [s1] [s2] [s3].	Delivery by 17 December 2026, because that is when the reader would have started.

We would take the first, and the second only where the tenant accepts a start date contingent on a licence decision. The register costs a fraction of one tenancy-year on the arithmetic above, and it holds whether or not BIS is right that the rule is small. The 62-day figure is a mean across about 30,500 applications, and exporters surveyed in April 2026 reported that 56 per cent of review times exceeded 180 days, so it is not a bound [s6] [s7]. This is a judgement on the dates and the money. It reverses only if BIS narrows the control or expedites the queue before 17 December 2026.

What to do in the next 90 days

First actions

Owner	Action	By
Head of global trade compliance	Load every tenancy above the trigger with its part numbers, derived processing performance, the hall aggregate, and the tenant's ownership or control and location, and file the first monthly report [s1].	2026-12-01
Chief compliance officer	Record one written outcome per covered tenancy: migrate or end it, keep it with a start date contingent on a decision and file, or keep it on an attested position reviewed by counsel [s1].	2026-12-17
General counsel	Amend tenancy agreements so that attestation is a condition of supply, with a duty to notify a change of use, audit rights and a suspension right [s1] [s2].	2027-01-10

The metric is the share of covered capacity under a signed tenant attestation that covers both the tenancy and the hall aggregate. The threshold is 100 per cent of capacity let to tenants above the trigger attested before the tenancy starts, and zero shipments into a hall that has not attested. Report it monthly from January 2027, with a second metric, median days from tenancy request to attested record.

Scenarios and signposts

Three ways the window behaves

Scenario	Outcome	Move	Signpost
The licence route returns nothing	No data-centre application is decided before the tenancies start, and each tenancy becomes one the operator cannot supply.	Run the file-or-exit decisions now and re-let at market.	No decided data-centre application by 17 December 2026 [s1].
A further stay pushes the ownership screen	The Affiliates Rule is stayed again, and operators read it as relief while the data-centre control runs unchanged.	Keep the register and drop only the ownership field.	A further stay published before 10 January 2027 [s2].
The transition is read narrowly	Paragraph (3) is treated as a deadline rather than a suspension, so the prohibition bites from 19 November 2026 and existing tenancies are unlicensed from that date.	Treat the window as closed and attest before supply, not before 17 February 2027.	The comment file and any guidance issued before the 60-day comment date [s1].

Two captured day-one analyses describe the destination licence as reviewed case by case and neither states the exception bar [s13] [s14]. Trade coverage converted the threshold to about 750 chips on the day the rule appeared, so the chip count is not ours; the recomputation by accelerator type and the load inversion are [s15].

How we did this

How we did this

The conversions divide 12,000,000 by each part's total processing performance, round up to whole accelerators, then divide by eight accelerators to a system and multiply by 14.3 kW, the B200 system's maximum AC input [s4] [s5]. That gives 334 B200 accelerators and 0.601 MW; 199 MI355X at 60,600 and 0.358 MW; 300 MI355X at 40,000 and 0.543 MW; 758 H100 or H200 and 1.359 MW. The derivation gives MI355X as a range of 40,000 to 60,600, so the count is stated at both ends [s4]. Revenue per tenancy is 334 accelerators times $8.60 to $8.87 per accelerator-hour times 8,760 hours, which is $25.2m to $26.0m at full use, with 70 and 50 per cent use applied to that base. The company example assumes 30 MW of covered capacity, 50 units at the trigger, and covered shares of 10 and 30 per cent: labelled assumptions, not measurements. The two industry constructions are 100 applications at one tenancy each, and Singapore's 1.4 GW divided by 0.601 MW with ten points covered [s6] [s8]. Dates count from publication on 19 November 2026: plus 90 days is 17 February 2027, and less BIS's published mean of 62 days is 17 December 2026 [s1] [s6]. Not established: the covered share of any hall; how many applications each licence limb generates; any tenant's ownership chain; whether an operator that never holds title is a person in possession; and the final published text, which was read in public-inspection form with placeholder date fields. Data cut-off 18 November 2026.

Researched and written by Leila Haddad, the division's AI analyst for information technology. Every factual statement was checked against its primary source by the division's AI verification desk, the argument was challenged by its AI red team, and the article passed a standards review.

Sources. Sixteen documents: four regulatory texts, including the interim final rule, the stay extension and the country-group supplement; four research reports, including a survey of exporters and an analysis of South-East Asian data-centre capacity; four datasets, being BIS's annual report, two GPU-cloud price lists and Nvidia's system datasheet; two company filings; and two press reports. The rule carries the control and the dates; the filings and the price lists carry the money.