SCALAC

Advanced Computing System for Latin America and the Caribbean

High-Performance Computing Infrastructure in Latin America and the Caribbean

2025 annual usage report

Presented September 2026

English edition | Revised 23 September 2026

Authors:

Nicolás WOLOVICK, PhD.

Tomás BAZÁN

Carlos Jaime BARRIOS HERNÁNDEZ, PhD.

Esteban HERNÁNDEZ, PhD.

Reviewer:

Philippe NAVAUX, PhD.

Disclaimer

The authors are responsible for the analysis and interpretations in this report. Findings are based on information submitted by participating centers and do not constitute an independent audit or a complete census of HPC infrastructure in Latin America and the Caribbean.

Usage refers to calendar year 2025. Inventory and support information was collected during March-September 2026, with a submission cutoff of 19 September 2026. The editorial revision of 23 September 2026 does not extend that cutoff or introduce new operational observations. Comparisons with the previous report use calendar-year 2024 usage and the 28 identifiable systems in its published inventory.

Theoretical performance estimates are not measured application throughput. Missing values, inconsistent accounting, shared resources and unconfirmed measurement units limit interpretation. Unvalidated power submissions must not be summed, ranked or converted into energy or emissions estimates. Values and limitations should be cited together when reusing the findings. Responsibility for the conclusions remains with the authors; participation or review does not imply institutional endorsement.

Reading conventions

The historical coverage baseline is 28 identifiable tabulated systems. Downtime and utilization are distinct rates; Table A2 always places downtime first. Table percentages and Figure 2 use two decimals. Commas separate thousands and periods separate decimals. Personnel are full-time / half-time / part-time headcounts; all zeros are numeric digits. Submitted power numbers remain unvalidated until centers confirm their units.

Contents

Summary

This study examines the infrastructure, utilization and operational support of **34 HPC systems in nine Latin American countries**, with 26 institutional affiliations reported by participating centers. The annual usage reference period is **1 January-31 December 2025**. Data were collected through a survey conducted between March and September 2026, and the report is presented in September 2026. Argentina contributes 12 systems and Brazil 11, together accounting for more than two-thirds of the systems studied. The findings describe participating centers; survey coverage does not extend to every regional HPC facility or include Caribbean countries in this edition.

The inventory reported in the 2026 survey has theoretical capacity of **7.236 PFLOP/s in CPU FP64** across 34 systems and **718.741 PFLOP/s in GPU FP16** across 24 systems with available estimates. These are separate measures of computational capability. Santos Dumont accounts for 50.3% of the reported CPU FP64 total, while Santos Dumont and Clementina XXI together account for 80.1% of GPU FP16 capacity. This concentration indicates that a small number of participating systems provide much of the sample's reported arithmetic potential. Outstanding technical validation limits the precision of these estimates. These inventory figures are not a reconstructed snapshot of capacity available throughout 2025; ql-front, for example, reports commissioning in 2026.

Operational performance varies widely. **Twenty-four systems provide numeric accounting for the 2025 usage study**, of which 17 satisfy the study's consistency check. Within that analytical group, weighted utilization of available capacity is **61.98%**, compared with a median of **43.31%**. Serafin records the highest utilization, at **96.18%**. Five systems combine utilization of at least 70% with downtime below 10%, using the screening thresholds retained from the previous study.

Reported support arrangements also differ: five systems have no full-time personnel assigned in their submissions, and 14 report no more than one. Facility efficiency information is available for 27 systems, with reported PUE values ranging from 1.10 to 2.00. Shared teams and facilities require careful interpretation before these observations can be aggregated.

Compared with the 28 identifiable systems in the 2025 edition, reported survey coverage increases to 34 systems (+6; +21.43%). Each edition's top-ten CPU FP64 total rises from 3.246 to 6.529 PFLOP/s, and the top-ten GPU FP16 total from 67.548 to 698.324 PFLOP/s. These groups differ in membership. Among nine provisionally matched systems with consistent current accounting, utilization rises in three and falls in six, comparing calendar-year 2024 with calendar-year 2025. Capacity estimates and operational use therefore require separate interpretation.

The findings point to three areas for regional cooperation: broadening effective access to available capacity, strengthening operational support, and improving comparable measurement across centers. Seven accounting inconsistencies and incomplete usage information for ten systems constrain the operational analysis. The study therefore distinguishes observed differences from conclusions about regional trends or their causes.

1. Introduction and methodology

1.1. Purpose and continuity

High-performance computing supports research and computational applications whose requirements vary across institutions, disciplines and countries. Understanding the region's capabilities requires attention to how systems are used and supported, alongside their theoretical performance.

This study continues the SCALAC HPC Observatory's examination of infrastructure and operational conditions in Latin America and the Caribbean, following its September/October 2025 report [2]. The present report analyzes calendar-year 2025 usage and compares it with calendar-year 2024 usage published in the previous edition. It also examines the inventory, personnel and facility power indicators submitted in the 2026 survey. The analysis identifies differences among centers and priorities for cooperation and subsequent measurement.

The historical comparison uses the **28 systems identifiable in the previous published inventory table** as its single coverage baseline. The current survey includes 34 systems: **six more, or 21.43%**. The previous narrative count of 29 has no identifiable additional tabulated system and is retained only as a source discrepancy, not as a comparison denominator. This expansion reflects survey participation and cannot, by itself, establish growth in installed capacity. Section 4 reconstructs comparisons from the previous study's published tables, separating changes in coverage from provisionally matched system observations.

1.2. Survey design and analytical population

All current inventory statistics come from a survey conducted between March and September 2026, using data submitted by participating HPC centers [1]. The unit of analysis is an individually reported HPC system. A center may contribute more than one system, and systems may share a facility, support team or user community. Institutional affiliations are retained as reported, without assuming that each represents an independent facility. LNCC's SDumont and Santos Dumont are analyzed as separate survey submissions. The name Santos Dumont follows LNCC's institutional spelling [7]; SDumont remains the label of the other submitted record. This editorial normalization does not establish distinct physical assets or resolve possible overlap.

Three dates define the study: **usage is measured for 1 January-31 December 2025**; responses were collected during **March-September 2026**; and this edition is **presented in September 2026**, using submissions available by 19 September. The September cutoff applies to collection and compilation, not to the annual usage period. The corresponding accounting interval starts on 1 January 2025 and ends immediately before 1 January 2026. The previous report measures calendar-year 2024 usage and was published in 2025.

The annual reference period is confirmed by the data provider, but exact accounting extract timestamps are not documented for every system. TMC- FIERRO's accounting text explicitly concerns 2024 and is excluded from the 2025 usage analysis. The submitted information dates ql-front's commissioning to 2026: it is retained in the survey inventory as a post-period system, with no 2025 usage rates. Santos Dumont's commissioning is tentatively placed in mid-2025; its available operating time may therefore cover only part of the reference year.

Inventory, theoretical capacity, staffing and facility indicators describe the information submitted in the 2026 survey. Their reference dates are not uniformly specified, so they do not establish a common 31 December 2025 snapshot or resources available throughout 2025. In particular, the total of 34 submitted systems must not be read as 34 systems operating during all of 2025. Capacity comparisons consequently retain report-edition labels, while utilization comparisons use the measured years 2024 and 2025.

Missing observations are excluded, while reported zeros are retained. ql-front's zero-user count describes a post-period submission, not 2025 operation. CPU and GPU peaks remain submitted estimates, not independently reconstructed hardware results; a missing GPU estimate does not establish absence of GPUs.

1.3. Indicator definitions and analytical method

The study uses two operational indicators:

- **Downtime share:** unavailable capacity, including planned downtime, divided by total reported capacity-time.
- **Utilization of available capacity:** allocated capacity-time divided by reported capacity-time after subtracting downtime and planned downtime.

In the accounting terminology used by participating centers, these are calculated as `(Down + PlannedDown) / Reported` and `Allocated / (Reported - Down - PlannedDown)`, respectively. Rates were independently recalculated from submitted accounting components. All utilization and downtime labels in tables and Figure 2 are rounded from the unrounded ratios to two decimal places. Numbers use commas for thousands and a period for decimals (for example, 2,982.70); personnel counts use the digit zero. Display rounding does not change the underlying calculations.

SchedMD distinguishes PlannedDown from Planned: PlannedDown covers specified maintenance reservations and certain node states, whereas Planned concerns eligible queued work that could not start [3]. These categories are kept separate. Downtime therefore includes planned unavailability, and allocation does not establish that processors or GPUs were fully busy while assigned.

Consistency was assessed by comparing the sum of allocated, down, planned-down, idle and planned capacity-time with the reported total. The absolute difference was allowed to reach five reported accounting units to accommodate rounding across the five components. In practice, the 17 consistent submissions differ by no more than one unit, whereas the seven flagged submissions differ by thousands of units or more. This check establishes arithmetic consistency; it does not independently verify extract dates, time in service or accounting configurations.

Operational summaries use the 17 consistent submissions. Weighted utilization is total allocated capacity-time divided by total available capacity-time; weighted downtime uses total reported capacity-time as its denominator. Medians give equal weight to each system. Aggregation assumes compatible accounting units. The previous collection protocol specified CPU-hours, but unit documentation in the current submissions is incomplete, so absolute discrepancies are expressed in reported accounting units. GPU utilization is not measured separately.

Active users are analyzed per system against the 2025 annual reference period. Counts are retained as submitted where individual dates remain unconfirmed; the TMC- FIERRO count cannot be independently assigned to 2025 because its accompanying usage text refers to 2024. The previous study used a threshold above 1,024 CPU-hours, but a common threshold and cross-account deduplication were not confirmed for all current submissions. Staffing is described using the reported full-time, half-time and part-time categories. Shared personnel are not summed into a regional headcount, and no full-time-equivalent conversion is assumed.

1.4. Institutional context and study scope

SCALAC provides a regional framework for cooperation in advanced computing, involving national research and education networks and RedCLARA [2]. The HPC Observatory's continuing analysis can support discussion of access, operational expertise and infrastructure planning across participating centers.

This study is descriptive and based on participating centers' submissions. It does not establish complete regional coverage, current TOP500 or Green500 membership, measured HPL performance, or comprehensive private-sector capacity. Connectivity and renewable-energy shares are also outside the available evidence. The performance rankings therefore compare systems within the study population.

2. HPC infrastructure inventory and system peak performance

2.1. Geographic coverage

Table 1. Coverage and reported theoretical capacity by country, as submitted in the 2026 survey. CPU FP64 and GPU FP16 values are in TFLOP/s; they are separate measures. A dash means no numeric GPU estimate is available for that country's systems.

Country	Systems	Systems reporting usage	CPU FP64 TFLOP/s	GPU FP16 TFLOP/s
Argentina	12	7	961.46	258,571.75
Brazil	11	9	5,318.82	430,164.72
Chile	1	1	383.19	2,620.00
Colombia	3	3	93.91	1,774.80
Costa Rica	1	0	33.50	2,982.70
Ecuador	1	1	13.82	7,488.00
Mexico	3	1	149.84	2,299.68
Peru	1	1	20.56	-
Uruguay	1	1	260.81	12,839.40
Total	34	24	7,235.92	718,741.05

Source: participating centers, SCALAC HPC survey 2026 [1]. Brazil accounts for 73.5% of reported CPU FP64 capacity and 59.8% of reported GPU FP16 capacity in the sample. Differences in participation and unresolved estimates limit wider regional interpretation.

2.2. CPU FP64 theoretical peak

Table 2. Ten highest reported CPU FP64 estimates. This ranks CPU-only theoretical performance. It is not a TOP500 ranking or an HPL Rmax result [4].

Rank	Cluster	Country	CPU FP64 TFLOP/s
1	Santos Dumont	Brazil	3,637.12
2	coaraci	Brazil	589.21
3	Clementina XXI	Argentina	446.46
4	leftraru	Chile	383.19
5	NPAD-UFRN	Brazil	370.53
6	Ada Lovelace	Brazil	352.59
7	ClusterUY	Uruguay	260.81
8	SDumont	Brazil	175.31
9	serafin	Argentina	157.29
10	pirayu	Argentina	156.62

Source: participating centers, SCALAC HPC survey 2026 [1]. Santos Dumont's submitted performance estimates require further technical validation; its leading position and the resulting concentration statistics are provisional. Different CPU generations and package configurations are retained as reported. A high theoretical peak does not establish achieved application throughput.

2.3. GPU FP16 theoretical peak and AI workloads

Table 3. Ten highest reported GPU FP16 estimates. The survey's FP16 estimates include a tensor-core performance basis. However, the accompanying technical information does not consistently distinguish dense from sparse operations or include measured AI benchmarks.

Rank	Cluster	Country	GPU FP16 TFLOP/s
1	Santos Dumont	Brazil	327,035.20
2	Clementina XXI	Argentina	248,344.00
3	SDumont	Brazil	45,772.80
4	Abaporu	Brazil	25,184.40
5	ClusterUY	Uruguay	12,839.40
6	Recod	Brazil	9,245.42
7	NPAD-UFRN	Brazil	9,206.40
8	cedia	Ecuador	7,488.00
9	coaraci	Brazil	6,934.20
10	mendieta	Argentina	6,273.80

Source: participating centers, SCALAC HPC survey 2026 [1]. Twenty-four systems have reported GPU FP16 estimates. The remaining ten have no estimate available and are excluded from the GPU total rather than assigned zero capacity.

These figures characterize potential arithmetic throughput. They do not establish whether a system can train a particular model or meet an inference latency target. Memory capacity, interconnect, software, allocation policy and workload shape are not captured adequately for that conclusion. Statements of FP8 support describe hardware capabilities and are not added to FP16 totals.

CPU FP64: ten highest reported estimates

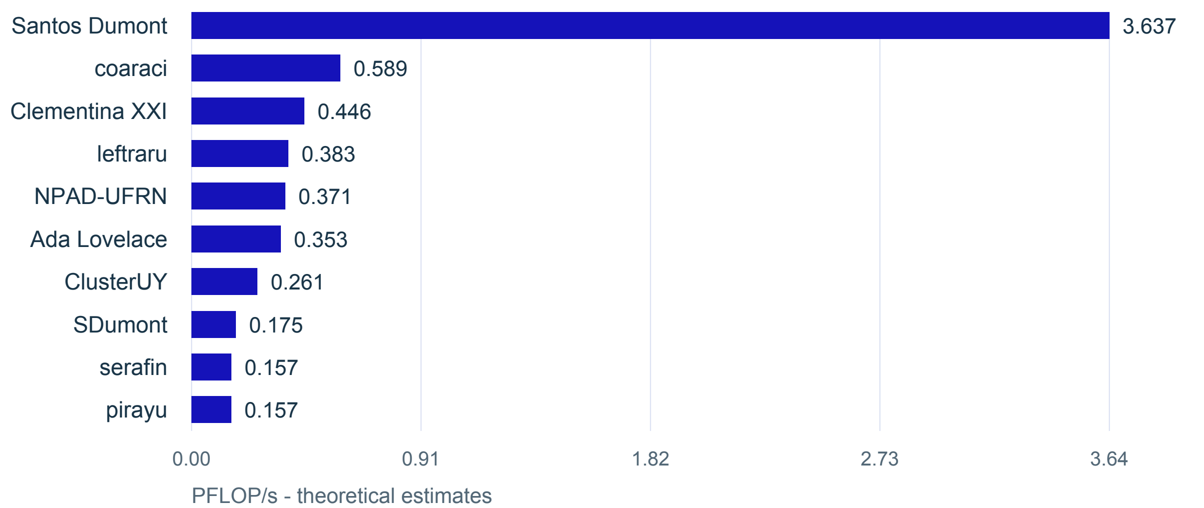

GPU FP16: ten highest reported estimates

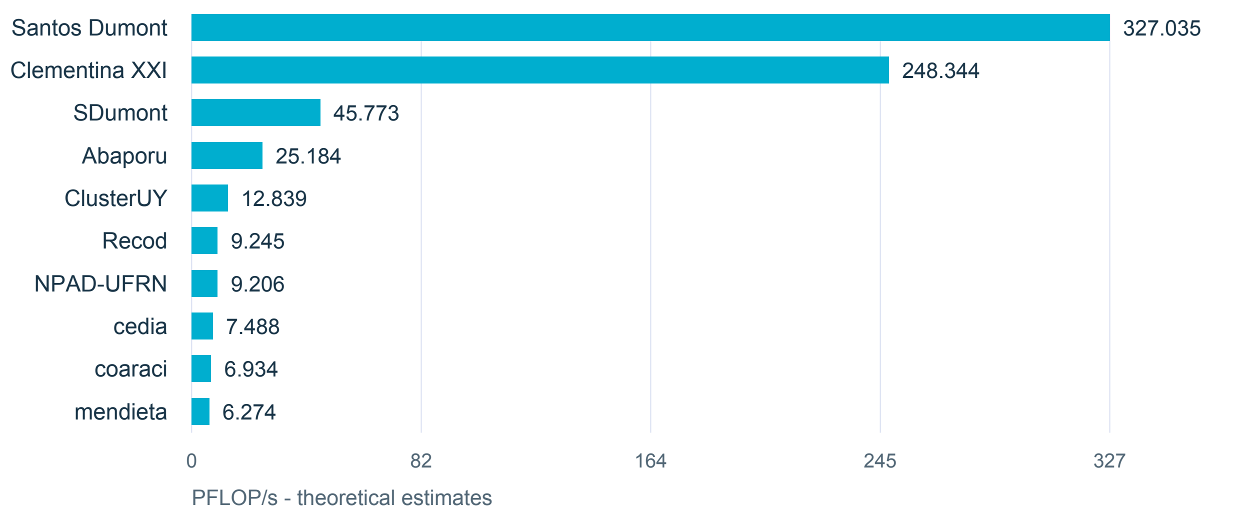

Figure 1. Separate theoretical performance rankings in PFLOP/s. One PFLOP/s equals 1,000 TFLOP/s. Source: [1], participating centers' performance submissions. Estimates awaiting technical validation remain provisional.

3. Usage, users, personnel and energy indicators

3.1. Accounting coverage and integrity

The 24 complete accounting records cover 70.6% of the inventory. Seventeen records, or 50.0% of all systems, pass the arithmetic reconciliation screen. The seven flagged records are yuca, Clemente, Santos Dumont, Abaporu, Recod, Sorgan and leftraru. Their original figures and computed rates remain visible in Appendix A with a warning marker; they are excluded from the reconciled-subset summary below.

The ten systems without numeric accounting are ql-front, Sersic, Mirta, Ada Lovelace, atocatl, tochtli, FIUNER, Kabré, TMC- FIERRO and Lobocarneiro. Their gaps have different causes. Ada Lovelace and Lobocarneiro provide free-text usage totals, atocatl and tochtli report unsupported accounting configuration, FIUNER reports that slurmdbd is not implemented, and Kabré labels its metrics inconsistent. TMC- FIERRO provides a historical SGE extract for 2024. The free-text totals are not mixed with the standardized accounting observations.

3.2. Utilization and downtime

Table 4. Descriptive usage indicators for 2025, for the 17 reconciled systems. These characterize the available records for the stated reporting period, not all regional systems. Weighting uses the reported accounting volumes and assumes consistent units across those records.

Indicator	Value	Calculation or scope
Weighted use of available capacity	61.98%	Sum Allocated / sum available capacity
Median per-cluster utilization	43.31%	Unweighted median across 17 systems
Weighted downtime	10.23%	Sum (Down + PLND Down) / sum Reported
Systems below 70% utilization	11	Descriptive screening threshold
Systems above 10% downtime	6	Includes planned downtime
Systems meeting both screens	5	Downtime <10% and utilization ≥70%

Source: participating centers, SCALAC HPC survey 2026 [1]. The weighted utilization of 61.98% exceeds the median of 43.31%, illustrating how high-volume systems affect the aggregate. These percentages do not establish a regional service standard or a trend relative to 2024.

Serafin records 96.18% utilization and 2.57% downtime. SDumont records 84.70% utilization and 4.12% downtime. Together with mendieta, eulogia and lamcad, they meet both screening thresholds. These cutoffs retain the prior report's analytical framing; they are not contractual availability requirements.

Pirayu combines 78.87% use of available capacity with 33.04% downtime. Its survey submission describes shutdowns when electricity is expensive, identifying an operational constraint that merits follow-up. PACCA records 0.67% utilization and 50.79% downtime. Guane records 68.04% utilization and 33.40% downtime. These patterns warrant examination of access, maintenance, power and accounting practices rather than an assumed common cause.

Leftraru's calculated utilization is 95.02%, but its component total exceeds Reported by 16,036,453 reported accounting units. Clemente's calculated utilization is 9.29%, with a negative reconciliation residual of 2,504,399 units. Both require correction or explanation before comparison with the reconciled group. Its current accounting discrepancy prevents a reliable utilization comparison with the earlier observation.

Use of available capacity

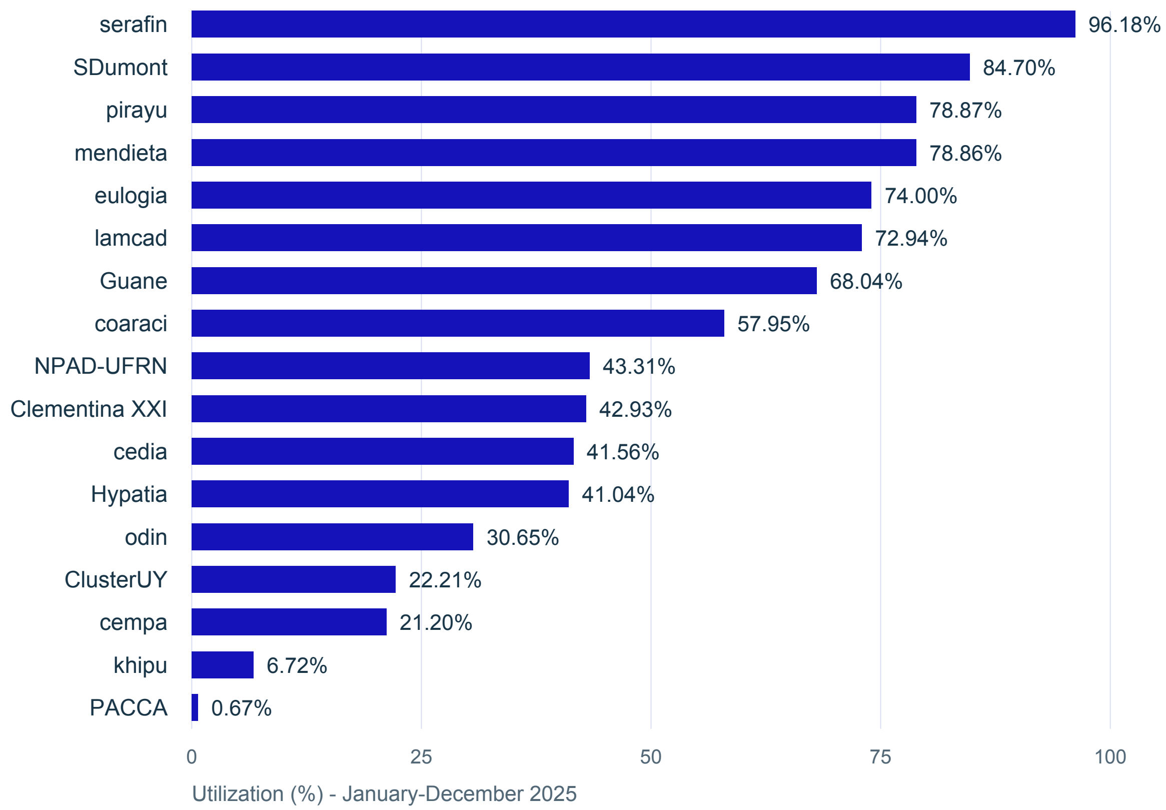

Figure 2. January-December 2025 utilization calculated from submitted accounting data for the 17 records passing the arithmetic screen. All labels use two decimal places, matching Table A2. Missing and flagged records are excluded. This screen does not independently verify extract dates, time in service or scheduler configurations.

3.3. Active users

Active-user counts are present for 32 systems. Their arithmetic sum is **3,286 cluster-level user counts**, which may include the same person on several clusters or accounts. The result must not be presented as 3,286 unique people across the region.

The largest reported counts are Ada Lovelace (502), SDumont (405), leftraru (367), NPAD-UFRN (219), Santos Dumont (217), coaraci (192) and serafin (188). These rankings concern submitted counts; they do not establish a common activity threshold or prove comparable demand. Sersic and Mirta have missing counts. ql-front and Recod explicitly report zero, a condition to retain and clarify rather than automatically replace with a missing value.

For future comparable counts, each site should identify the resource and threshold, consolidate multiple accounts for each user before applying the threshold, and provide reporting start and end dates. Counting user-account combinations can count a person more than once, consistent with the structure of Slurm's UserUtilizationByAccount output [3].

3.4. Personnel and support capacity

All 34 systems provide personnel information by time commitment. Five list no full-time personnel: cedia, ClusterUY, khipu, PACCA and odin. Fourteen systems list no more than one full-time person. Six list more than three: yuca, SDumont, Santos Dumont, Ada Lovelace, leftraru and coaraci.

These are system-level staffing observations. Repeated staffing profiles within LNCC, UNC Supercómputo, LAMOD and other institutions may describe shared teams, so summing them would risk double-counting personnel. The submitted information does not consistently specify how much of each person's time is assigned to a particular system. No regional headcount or full-time-equivalent total is therefore reported.

The data identify systems where support coverage deserves discussion. They do not establish that staffing caused low utilization or high downtime. A stronger analysis would collect team identifiers, responsibility boundaries, support hours, vacancies and the number of systems served by each team.

3.5. Power and environmental reporting

Participating centers report facility maximum power for 28 systems, PUE for 27, and cluster peak consumption for 22. The reported PUE values range from 1.10 to 2.00, with a median of 1.80. PUE is the ratio of total data-center energy to ICT equipment energy [5]. These are submitted observations, not independently validated annual ratios or a capacity-weighted regional efficiency measure. Repeated facility values may refer to a shared data center.

Power units are absent from the submissions and cannot be inferred from magnitude. Table A3 retains the submitted numbers solely as a validation register: they are excluded from cross-center comparisons and regional totals. Each center must confirm the unit, quantity, measurement boundary and observation date before these values can support analysis. ql-front has a reported peak-consumption value of zero; this is not evidence that the operating cluster consumes no electricity.

Power is measured in watts (W) and energy in joules (J); unit prefixes distinguish kilowatts from megawatts [6]. Peak power does not establish energy consumed over the reporting period. A value expressed as apparent power (kVA) cannot be treated as active power (kW) without the corresponding power factor [8]. No regional energy, carbon-emissions or water-consumption total is calculated. Such an assessment requires confirmed power units, facility and IT measurement boundaries, time-integrated energy, measurement dates and appropriate emissions factors. PUE observations would also need a consistent measurement basis before comparing facilities.

4. Comparison with the previous report

4.1. Comparison basis and survey coverage

The comparison uses the 2025 study's published system, CPU FP64 and GPU FP16 tables [2] and the 2026 submissions [1]. Capacity comparisons refer to the inventories presented in the two editions. Operational comparisons concern **calendar-year 2024 usage reported in the 2025 edition** and **calendar-year 2025 usage presented in September 2026**. Both usage reference periods are complete, consecutive calendar years. The submission dates do not define the usage period. Inventory estimates retain their report-edition dates because a common year-end hardware snapshot is not established.

Two views are used: changes in survey-wide coverage and rankings, and changes for systems provisionally matched across editions. System names and institutional associations support those matches, but permanent identifiers and identical hardware boundaries are not available for every case. The results describe changes in reported estimates and observations, rather than independently verified hardware additions.

Table 5. Principal changes between editions. The two top-ten totals cover each edition's ten highest estimates, whose membership changes. Historical usage rates cannot undergo the same component-level reconciliation because only the published percentages are available.

Indicator	2025 edition	September 2026 edition	Change / interpretation
Systems in the comparison inventory	28	34	+6 (+21.43%); survey coverage
Countries represented in tabulated inventory	6	9	+3; Ecuador, Costa Rica and Peru added to coverage
Top-ten CPU FP64 estimates (PFLOP/s)	3.246	6.529	+101.16%; different ranking membership
Top-ten GPU FP16 estimates (PFLOP/s)	67.548	698.324	+933.82%; different ranking membership
Systems with both usage rates available	26 of 28 (92.9%)	24 of 34 (70.6%)	-22.3 percentage points in reporting coverage
Systems in current reconciled usage analysis	Not assessed on historical components	17	A stricter current analytical subset
Sum of reported system-level active-user counts	3,966 (27 systems)	3,286 (32 systems)	-680; participation and user definitions differ
Systems reporting more than three full-time staff	3, according to prior narrative	6	+3 system-level reports; shared teams may recur

Source: participating centers' submissions [1] and the 2025 study, Tables 2-4 and accompanying discussion [2]. All coverage comparisons use the 28 systems identifiable in historical Table 2. The narrative mentions 29, but no additional system can be identified from that table; this unresolved source discrepancy is not used as a second denominator. The older study also describes eight incomplete submissions, a broader condition than the absence of usage rates.

Table 6. Geographic coverage of the tabulated systems. Absence from an edition means absence from its comparison sample, not the absence or retirement of national HPC infrastructure.

Country	2025 report systems	2026 report systems	Difference
Argentina	11	12	+1
Brazil	6	11	+5
Chile	2	1	-1
Colombia	4	3	-1
Costa Rica	0	1	+1
Ecuador	0	1	+1
Mexico	4	3	-1
Peru	0	1	+1
Uruguay	1	1	+0

Brazil's representation rises from six to 11 systems, while Argentina rises from 11 to 12. Chile, Colombia and Mexico each have one fewer system represented. Survey expansion therefore combines broader country coverage with changes in participation within countries. Counts alone do not establish commissioning or decommissioning.

4.2. Changes in reported computational capacity

The top-ten CPU FP64 total increases from **3.246 to 6.529 PFLOP/s**, a **+101.16%** change. The corresponding GPU FP16 total increases from **67.548 to 698.324 PFLOP/s**, or **10.34 times** the earlier value. These comparisons keep the precision and number of ranked systems consistent, but include changes in membership, reported configurations and estimation methods.

For the reference systems compared here, the 2025 study publishes CPU FP64 and GPU FP16 estimates in top-ten lists, rather than for every surveyed system. A complete historical capacity total for this survey population is therefore unavailable. The current full-sample totals of 7.236 PFLOP/s CPU FP64 and 718.741 PFLOP/s GPU FP16 must not be compared with those historical top-ten sums as if both covered the entire survey.

Santos Dumont and Clementina XXI together account for 575.379 PFLOP/s, or 82.4% of the current top-ten GPU total. Clementina XXI is absent from the previous top-ten table, which does not establish that it was newly commissioned between studies. LNCC previously appeared as one Santos Dumont entry and now submits SDumont and Santos Dumont separately. Their system boundaries require confirmation before interpreting that change as an upgrade to a single matched machine.

Table 7. CPU FP64 estimates for seven provisionally matched systems. Values are TFLOP/s; historical values retain published precision.

System (2026 name)	2025 report estimate	2026 report estimate	Difference (TFLOP/s)	Change
coaraci	589.20	589.21	+0.01	0.00%
leftraru	373.10	383.19	+10.09	+2.71%
Ada Lovelace	314.40	352.59	+38.19	+12.15%
serafin	157.30	157.29	-0.01	-0.01%
pirayu	116.70	156.62	+39.92	+34.21%
ClusterUY	90.40	260.81	+170.41	+188.51%
eulogia	85.20	85.20	0.00	0.00%

Source: survey submissions [1] and historical Table 3 [2]. For these seven matched systems, the combined estimate rises from 1,726.30 to 1,984.91 TFLOP/s (+14.98%). ClusterUY shows the largest relative increase, at 188.51%, followed by pirayu at 34.21% and Ada Lovelace at 12.15%. Coaraci, serafin and eulogia are effectively unchanged at the historical reporting precision; differences below 0.02% should not be interpreted as hardware changes.

Table 8. GPU FP16 estimates for six provisionally matched systems. These remain theoretical estimates, subject to consistent tensor-core and dense/sparse conventions.

System (2026 name)	2025 report (TFLOP/s)	2026 report (TFLOP/s)	Difference (TFLOP/s)	Change
coaraci	6,934.20	6,934.20	0.00	0.00%
mendieta	6,273.80	6,273.80	0.00	0.00%
Ada Lovelace	3,120.00	3,876.40	+756.40	+24.24%
leftraru	2,989.20	2,620.00	-369.20	-12.35%
ClusterUY	1,783.80	12,839.40	+11,055.60	+619.78%
Guane	543.00	543.00	0.00	0.00%

Source: survey submissions [1] and historical Table 4 [2]. The six-system total rises from 21,644.00 to 33,086.80 TFLOP/s (+52.87%). ClusterUY contributes the largest increase, from 1,783.80 to 12,839.40 TFLOP/s; its submission reports newer GPUs installed in late 2025. The resulting inventory estimate must not be treated as capacity available for the whole of 2025. Ada Lovelace rises by 24.24%, while coaraci, mendieta and Guane retain the same estimates. Leftraru's reported FP16 estimate is 12.35% lower, a result requiring examination of hardware coverage and estimation assumptions before attributing it to reduced capability.

The much larger rise in the changing top-ten GPU group than in the six matched systems shows how strongly the headline result depends on which systems and configurations are represented. None of these estimates measures achieved application throughput. Comparable HPL Rmax, AI training or inference benchmarks are not available for the two editions, so a change in measured computational performance cannot be quantified [4].

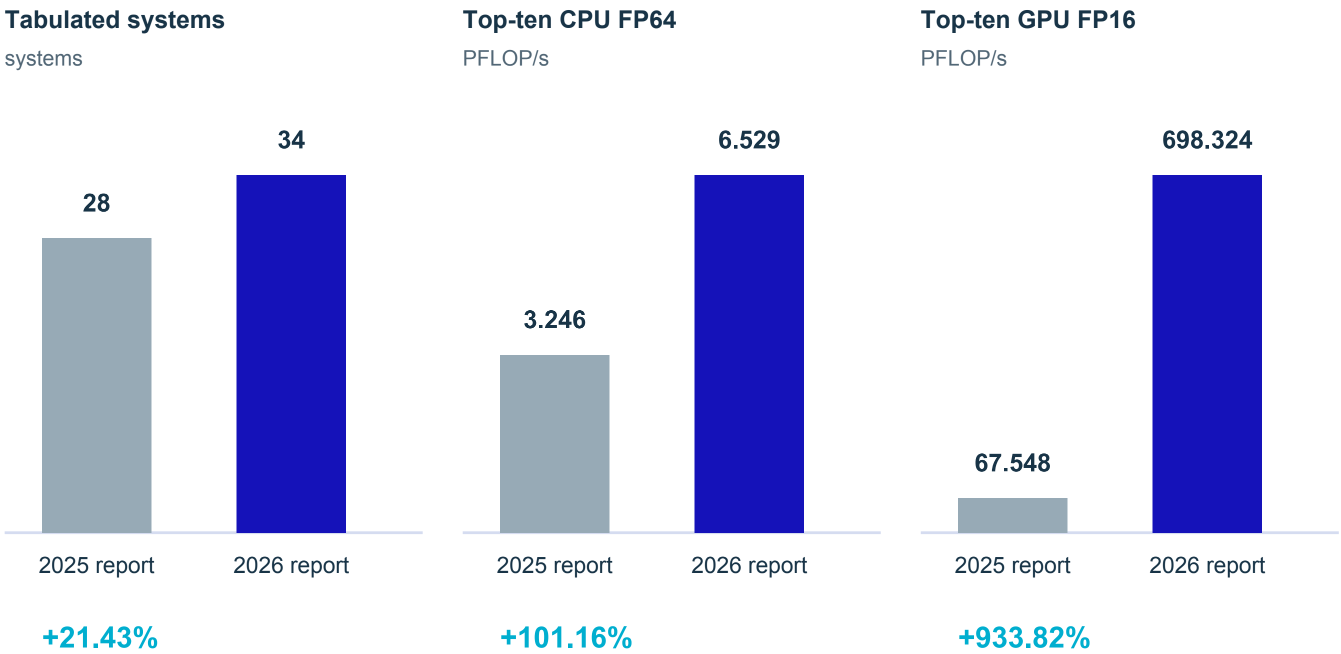

Figure 3. Reported system count and the sums of each edition's ten highest CPU FP64 and GPU FP16 estimates. The historical baseline is the 28 systems identifiable in the published inventory table. The ranked systems differ between editions. Sources: participating centers [1] and the 2025 study [2].

4.3. Changes in operational performance

Table 9. Usage indicators for nine provisionally matched systems with reconciled 2025 accounting. Historical rates cover January-December 2024; current rates cover January-December 2025. Differences are percentage points (pp), not percentage growth. The current consistency check cannot retrospectively validate the older percentages.

System	2024 utilization	2025 utilization	Change (pp)	2024 downtime	2025 downtime	Change (pp)
pirayu	83.95%	78.87%	-5.08	43.54%	33.04%	-10.50
mendieta	75.26%	78.86%	+3.60	5.09%	3.60%	-1.49
eulogia	93.02%	74.00%	-19.02	4.75%	7.05%	+2.30
serafin	92.34%	96.18%	+3.84	10.71%	2.57%	-8.14
coaraci	49.45%	57.95%	+8.50	2.12%	0.37%	-1.75
Guane	77.88%	68.04%	-9.84	26.13%	33.40%	+7.27
Hypatia	57.35%	41.04%	-16.31	4.80%	4.10%	-0.70
PACCA	1.36%	0.67%	-0.69	25.65%	50.79%	+25.14
ClusterUY	43.92%	22.21%	-21.71	1.38%	1.86%	+0.48

Source: historical Table 2 [2] and current submissions [1]. Utilization rises in three of these nine systems and falls in six. The median declines from **75.26% to 68.04%**. This is a year-over-year comparison of the 2024 and 2025 annual reference periods. Changes in configuration, time in service and accounting coverage still limit attribution of the differences to operational improvement or deterioration.

Serafin combines a 3.84 pp utilization increase with an 8.14 pp downtime decrease. Mendieta also improves both indicators, while coaraci's utilization rises by 8.50 pp and downtime falls by 1.75 pp. Pirayu reports 10.50 pp less downtime but 5.08 pp less utilization of available capacity, illustrating that availability and allocation can move differently.

Eulogia and ClusterUY show utilization decreases of 19.02 pp and 21.71 pp, respectively. PACCA combines a 0.69 pp utilization decrease with a 25.14 pp downtime increase. These observations identify specific operational questions for centers; they do not isolate the effects of demand, access policy, maintenance or resource configuration.

Five systems in the earlier published table meet utilization ≥70% and downtime <10%, compared with five in the present reconciled group. The older narrative names four, but its table also places Clemente within the thresholds. Because current Clemente accounting needs reconciliation and other systems lack current numeric usage, an unchanged count is not evidence of stable performance across a fixed population.

4.4. Active users, personnel and facility indicators

Table 10. Submitted active-user counts for selected provisionally matched systems. The annual reference periods are January-December 2024 and January-December 2025. The current activity threshold and deduplication procedures are not uniformly confirmed; differences must not be interpreted directly as growth or decline in research demand.

System	2024 users	2025 users	Difference
Ada Lovelace	420	502	+82
leftraru	354	367	+13
ClusterUY	118	117	-1
serafin	180	188	+8
coaraci	157	192	+35
pirayu	39	57	+18
mendieta	85	86	+1
eulogia	66	53	-13
Hypatia	108	113	+5
PACCA	3	1	-2
Guane	20	16	-4
Lobocarneiro	117	53	-64

Source: historical Table 2 [2] and current submissions [1]. Ada Lovelace reports 82 more users, coaraci 35 more and pirayu 18 more; Lobocarneiro reports 64 fewer. Submitted counts sum to 3,286, versus 3,966 previously. Participation, user definitions and unconfirmed individual dates limit comparison; users may recur across systems. The older LNCC count of 1,740 also requires aligned system boundaries before comparison with its two current submissions.

Three systems in the previous report and six in the current survey report more than three full-time staff. Shared teams and changing participation prevent interpreting this as a doubling of the regional workforce. These are survey-edition observations, not annual-average staffing for 2024 and 2025.

PUE and power observations lack a harmonized historical series, so changes in energy efficiency or consumption cannot be quantified. Future comparisons should align dates, units and facility boundaries.

5. Sustainability analysis

5.1. Framework, provenance and comparison

The environmental framework, site-level reference values and multidimensional presentation approach are based on Barrios H. et al. [9], particularly Section 3, equations (1)-(7), Table 4 and Figures 1-4 of that paper. The framework considers altitude, thermal stress, humidity, electricity sources and operational use. This report contrasts its conclusions with the Observatory's center-submitted evidence. Usage covers January-December 2025; inventory, staffing and facility indicators were submitted during March-September 2026 and do not share a confirmed observation date.

The paper reports usability by site or city. This study calculates utilization by system as allocated capacity-time divided by available capacity-time. A common definition, accounting period and one-to-one site-to-system correspondence have not been established. Table 11 therefore juxtaposes the values for consistency assessment; it does not measure changes over time or independently validate the paper's modeled indices.

Table 11. Site-level reference values from the sustainability paper and candidate system-level observations from the present study. An asterisk identifies current accounting that requires reconciliation. Correspondences are provisional; city-level observations must not be substituted for individual system records.

Site in the paper	Paper use (%)	Candidate system: 2025 utilization	Assessment
Cartagena	40	PACCA: 0.67%	Low use in both accounts; environmental causes untested. Downtime is 50.79%.
Piedecuesta	70	Guane: 68.04%	Similar use; 33.40% downtime warrants investigation. Renewable supply is not verified here.
San José	80	Kabré: unavailable	No comparable numeric accounting to test the operational conclusion.
Santiago	90	leftraru: 95.02% *	Apparent high use remains provisional; accounting residual is +14.09%.
Toluca	90	No confirmed match	No system-level comparison established.
Córdoba	90	serafin: 96.18%; mendieta: 78.86%; eulogia: 74.00%	These UNC examples show variation within one institutional setting; they are not a city-wide aggregate.
Bogotá	60	Hypatia: 41.04%	Downtime is only 4.10%; the data do not establish altitude as a cause of lower use.
Petrópolis	90	SDumont: 84.70%; Santos Dumont: 37.10% *	Distinct submissions; system identity, accounting and possible overlap need resolution.
Porto Alegre	70	No confirmed match	No system-level comparison established.

Source: Barrios H. et al. [9], Table 4 (p. 11), and this report's Table A2. Accounting qualifications are documented in Table B1. The paper's Table 4 is different from Table 4 of this report, which summarizes the 17 reconciled systems.

5.2. Supplied multidimensional visualization

Figure 4 presents the supplied radar as a discussion illustration. Its six-axis layout adapts the paper's multidimensional approach; the radar and its 0-10 scores are not reproduced from the paper. They combine reported indicators with environmental assumptions and transformations that have not been independently validated.

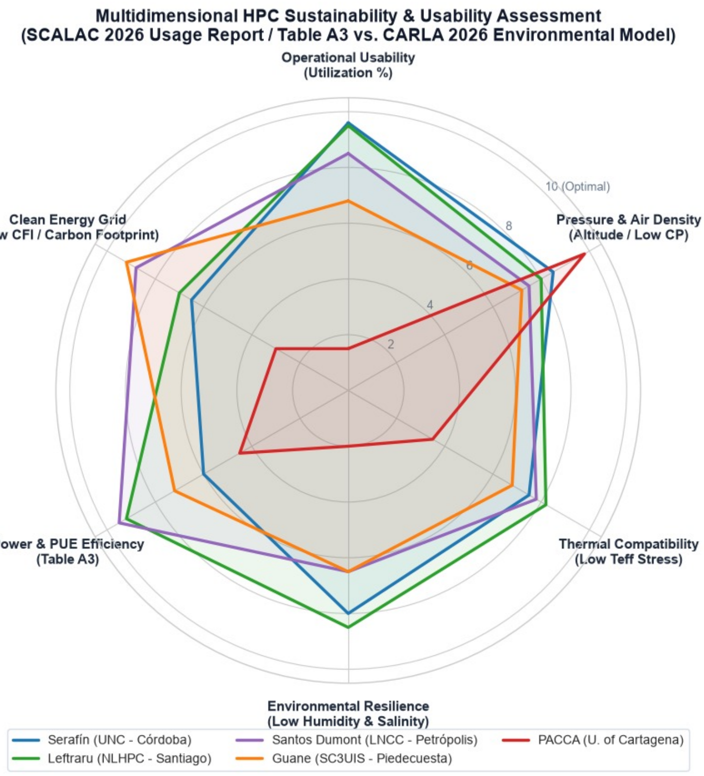

Figure 4. Supplied multidimensional sustainability and usability illustration, reproduced unchanged. Conceptual basis: Barrios H. et al. [9]; operational and facility reference data: Tables A2 and A3 of this report. This is not a validated ranking. The embedded title's "2026 Usage Report" denotes the presentation edition; measured usage in this study is for 2025. Table A3 alone does not support all six axes. Read the qualifications below before interpreting the profiles.

Interpretation qualifications. The Santos Dumont profile uses approximately 84.7% utilization, which belongs to the separate SDumont record; Santos Dumont's own calculated rate is 37.10% and requires reconciliation. Leftraru's 95.02% also requires reconciliation. PACCA's displayed utilization score of approximately 1.5/10 is inconsistent with a linear conversion of its 0.67% utilization, which would yield 0.067/10.

The normalization of environmental and energy scores is not documented sufficiently to reproduce them. Guane and PACCA have no reported PUE, and submitted power numbers have unconfirmed units. The radar therefore cannot establish relative carbon footprints, measured corrosion, hardware degradation, energy efficiency or overall sustainability. These qualifications preserve the distinction between the supplied illustration and the study's verified arithmetic.

5.3. Findings, limits and assessment priorities

The comparison supports considering sustainability beyond peak performance and utilization. It provides only partial operational agreement with the paper's site profiles. PACCA's very low use is consistent with the paper's qualitative description of low use in Cartagena, but does not establish its proposed environmental causes. Guane's utilization is close to the Piedecuesta reference value; its downtime remains an operational issue to investigate. The different rates among UNC systems and between the LNCC submissions prevent a single city-level value from representing every system. Missing or inconsistent accounting limits the remaining comparisons.

The paper attributes environmental penalties or advantages to climate and electricity sources. The Observatory did not collect the facility-level environmental measurements, annual energy totals, electricity-emission factors or failure histories needed to confirm those conclusions. A site's position in the model is therefore a result attributed to the paper, not an observed environmental classification from this survey. Similarity in utilization cannot validate an electricity-source assumption or establish that climate caused downtime. The two sources also concern SCALAC sites and must not be assumed to provide independent observations.

The paper defines a Sustainability Index using usability in its numerator [9, equation (7)]. A relationship between that index and use is consequently part of its construction, rather than independent evidence that higher utilization improves sustainability. Reproduction requires the complete input data, energy-category weights, carbon factors, thermal parameters, humidity model and normalization. The paper itself treats explicit corrosion assessment as an extension of the current model. Neither that model nor the supplied radar supports a measured lifecycle or emissions ranking with the present data.

Our observations identify concrete measurement priorities. Seven of 24 numeric usage records fail reconciliation. Fourteen of 34 systems (41.18%) report no more than one full-time person, while five report none; shared teams prevent treating these as independent facility staffing totals. PUE is reported for 27 systems, ranging from 1.10 to 2.00 with a median of 1.80, but common measurement boundaries and periods remain unconfirmed. These findings justify improving operational and sustainability evidence; they do not demonstrate burnout, environmental damage or their causes.

A reproducible next assessment should collect stable site and system identifiers; harmonized accounting periods and definitions; confirmed power units and facility boundaries; time-integrated facility and IT energy; documented electricity sourcing and emission factors; indoor environmental and corrosion measurements; failure histories; and information on shared support teams and workload. Lifecycle and social-impact assessments also need explicit scope, time horizons and outcome measures. Predictive models and digital twins can then be evaluated against observations.

These requirements provide an evidence agenda for the proposed HPC SCALAC Certification in Section 6.1. Certification remains a prospective recommendation: the study does not establish compliance, validate the supplied scores or demonstrate environmental benefits from certification.

6. Final remarks

The comparison shows expanded survey coverage and substantially higher reported top-ten theoretical capacities, especially for GPU FP16. Changes among matched systems are more varied, and utilization does not rise uniformly alongside capacity. These results support treating computational potential, operational availability and effective use as separate dimensions of regional HPC development.

The participating centers illustrate a heterogeneous HPC landscape. Reported computational potential is concentrated in a few systems, while utilization, support arrangements and facility conditions vary substantially. The separate CPU FP64 and GPU FP16 results also show why regional capacity should be assessed through more than one performance measure.

The difference between weighted utilization and the median indicates that larger accounting volumes materially influence the aggregate. Several centers combine high use of available capacity with comparatively low downtime, while others report considerable unavailability or low allocation. Understanding these differences requires attention to operating schedules, access policies, workload demand, maintenance and power constraints. The present study cannot isolate their individual effects.

Support capacity is a further area for cooperation. Systems reporting few full-time staff may benefit from shared expertise, training and coordinated services, although the responsibilities of shared teams must be established before estimating staffing needs. Similarly, facility indicators offer a starting point for environmental analysis, but consistent units, measurement boundaries and time-integrated energy data are necessary before comparing consumption or emissions.

Three priorities follow from these findings. First, participating centers can investigate access and operational constraints that limit effective use of existing capacity. Second, regional collaboration can strengthen support practices and the exchange of technical expertise. Third, a common measurement framework can make future studies more comparable, including consistent reporting periods, active-user definitions, hardware estimation methods and identifiers for shared facilities and teams.

Resolving the seven accounting discrepancies and completing missing observations will strengthen the evidence for these decisions. Future assessments should follow matched systems over time, distinguishing changes in participation from changes in infrastructure and operational performance.

6.1. Suggested recommendation: HPC SCALAC Certification

To help close the gap between ambitious computing goals and long-term sustainability in Latin America and the Caribbean, the authors recommend that regional institutions and funding organizations adopt **HPC SCALAC Certification as a key requirement for cyberinfrastructure investments, upgrades and shared systems**. Assessment should extend beyond peak FLOP/s and nominal PUE to the conditions that sustain reliable services and meaningful scientific and social outcomes.

For this recommendation, the certification framework should establish a transparent, auditable baseline covering environmental risk, effective utilization, management practices, technological and data sovereignty, scientific and human impact, and regional cooperation. Environmental assessments should address site-specific concerns such as altitude-related cooling constraints, thermal degradation, humidity and galvanic corrosion. Clean-energy sourcing, documented measurement boundaries and transparent reporting should be explicit criteria.

Dedicated technical personnel and sustained operating support should be central requirements. Investment plans should fund ongoing training, monitoring, preventive maintenance and shared-resource management, alongside staffing continuity and manageable workloads. These provisions should reduce the risks of staff burnout, loss of operational knowledge and avoidable hardware failures. Certification should assess how facilities support their people and users as well as their equipment.

Implementation should define published criteria, evidence requirements, responsibilities, periodic reassessment and feasible improvement plans. In this way, SCALAC could help institutions connect investment decisions with accountable operations, regional cooperation and long-term service continuity, strengthening a resilient and productive scientific community.

This is a prospective recommendation by the authors. The criteria described here are proposed expectations for implementation; this study does not assess centers against a certification standard or demonstrate that certification has already produced these outcomes.

7. Acknowledgments

The authors thank the platform administrators, system operators and technical teams at all participating HPC centers who collected and submitted the infrastructure, usage and operational support data during the March-September 2026 survey. Their work maintaining the systems and preparing the accounting information made this study possible.

Special thanks to Tomás Bazán ("Carpincho"), a scholarship holder at UNC Supercómputo. The authors also thank Philippe Navaux for reviewing the report and identifying issues of presentation and interpretation.

It also recognizes the groundwork documented in the preceding study, which acknowledged UNC Supercómputo staff Marcos Mazzini, Carlos Bederián and Alejandro Silva; Carla Osthoff and Micaella Coelho Valente de Paula; Luis Alejandro Torres Niño; Antonia Moreno and Rodrigo Oportot from CENIA; and Lizette Robles Dueñas from CADS-UDG [2].

8. References and bibliography

1. **SCALAC HPC Observatory.** Survey of HPC infrastructure, usage and operational support in Latin America, 2026. Data submitted by participating centers between March and September 2026. Usage reference period: 1 January-31 December 2025, as confirmed by the data provider. Unpublished survey compilation; submission cutoff for this edition: 19 September 2026.

2. **SCALAC HPC Observatory.** High-Performance Computing "Big" Robust Systems Report and Analysis in Latin America and the Caribbean, version 3.1. 2025 edition, revision dated 23 October 2025. Table 2 (printed p. 13) and Tables 3-4 (printed p. 14) provide the historical comparison values. Usage reporting period: calendar year 2024.

3. **SchedMD.** Slurm Workload Manager: sreport. Consulted 23 September 2026 for accounting terminology and the structure of user-account reports.

4. **TOP500.** The Linpack Benchmark. Methodological reference for measured benchmark performance and its distinction from theoretical peak. Accessed 23 September 2026.

5. **The Green Grid.** Power Usage Effectiveness (PUE). Definition of the ratio of total data-center energy to ICT equipment energy. Accessed 23 September 2026.

6. **National Institute of Standards and Technology (NIST).** Guide to the SI, Chapter 4: The Two Classes of SI Units and the SI Prefixes. Reference for power, energy and unit prefixes. Accessed 23 September 2026.

7. **Laboratório Nacional de Computação Científica (LNCC).** Supercomputador Santos Dumont. Institutional reference for system naming only; online capacity figures do not replace the submitted survey estimates. Accessed 23 September 2026.

8. **Schneider Electric.** What is the formula to determine the kVA on a transformer?. Technical reference for the power-factor requirement when relating kW and kVA. Accessed 23 September 2026.

9. **Barrios H., Carlos J.; Meneses, Esteban; Wolovick, Nicolás; Guerrero, Ginés; Navaux, Philippe; Gitler, Isidoro; Montoya, Javier; and Peña, Dario.** *Sustainability-Driven HPC Deployment Analysis: Life-Cycle Sustainability in Large-Scale HPC Infrastructure Implementations*. 2026. CARLA 2026 proposal/manuscript supplied for this study, 17 pp. Methodological basis: Section 3, equations (1)-(7); Table 4; Figures 1-4; and Sections 4-6. Publication details and a DOI are not established from the supplied copy. Its corrosion discussion cites entries [22]-[26], which are absent from that copy's reference list; those citations require completion before reuse of their detailed technical claims.

Appendix A. Complete cluster indicators

Table A1. Inventory and reported capacity submitted in the 2026 survey. All 34 surveyed systems are included. CPU FP64 and GPU FP16 are theoretical TFLOP/s values. A dash denotes an unavailable observation. Institution and cluster are separate fields: IATE is the institution and Clemente is its cluster. The submitted identifiers ql-front and atocatl are used consistently.

Country	Institution	Cluster	CPU FP64	GPU FP16
Argentina	UNMDP	**ql-front**	58.56	2,496.00
Mexico	ACARUS	**yuca**	58.75	-
Colombia	UAndes	**Hypatia**	47.72	919.80
Argentina	IATE	**Clemente**	4.99	-
Argentina	IATE	**Sersic**	2.05	-
Argentina	IATE	**Mirta**	2.05	-
Brazil	LNCC	**SDumont**	175.31	45,772.80
Brazil	LNCC	**Santos Dumont**	3,637.12	327,035.20
Ecuador	CEDIA	**cedia**	13.82	7,488.00
Argentina	CIMEC	**pirayu**	156.62	-
Argentina	UNC Supercómputo	**serafin**	157.29	-
Argentina	UNC Supercómputo	**mendieta**	8.51	6,273.80
Argentina	UNC Supercómputo	**eulogia**	85.20	-
Brazil	CENAPAD-SP	**Ada Lovelace**	352.59	3,876.40
Mexico	LAMOD	**atocatl**	25.93	1,086.30
Mexico	LAMOD	**tochtli**	65.16	1,213.38
Brazil	Núcleo de Processamento de Alto Desempenho	**NPAD-UFRN**	370.53	9,206.40
Uruguay	Centro Nacional de Supercomputación	**ClusterUY**	260.81	12,839.40
Brazil	Recod.ai	**Abaporu**	20.17	25,184.40
Brazil	Recod.ai	**Recod**	19.39	9,245.42
Brazil	Unicamp	**Sorgan**	5.38	349.10
Peru	Khipu	**khipu**	20.56	-
Argentina	Clementina XXI	**Clementina XXI**	446.46	248,344.00
Argentina	HPC-FIUNER	**FIUNER**	7.68	-
Costa Rica	Kabré	**Kabré**	33.50	2,982.70
Argentina	Gerencia Física, Centro Atómico Bariloche	**TMC- FIERRO**	11.66	1,085.53
Colombia	SC3UIS	**Guane**	12.91	543.00
Colombia	Universidad de Cartagena	**PACCA**	33.28	312.00
Brazil	Nacad/Coppe/UFRJ	**Lobocarneiro**	51.46	-
Chile	Laboratorio nacional de computación de alto rendimiento	**leftraru**	383.19	2,620.00
Brazil	LAMCAD	**lamcad**	53.56	1,048.00
Brazil	LAMCAD	**cempa**	44.11	1,512.80
Brazil	CCJDR	**coaraci**	589.21	6,934.20
Argentina	Centro de Computación de Alto Rendimiento (CeCAR)	**odin**	20.40	372.42

Source: participating centers, SCALAC HPC survey 2026 [1]. Estimates are retained as submitted. Outstanding technical validation questions are summarized in Appendix B.

Table A2. 2025 usage and submitted user and personnel counts. "Reconciled" means only that the arithmetic check passed. An asterisk marks rates based on accounting that needs reconciliation. A dash denotes missing data. Personnel are shown as full-time / half-time / part-time headcounts, as reported for each system. Personnel reflect the 2026 survey submissions, rather than a confirmed 2025 headcount. User reference dates are not independently verified for every system; ql-front is a post-period system and TMC- FIERRO's usage text refers to 2024.

Cluster	Downtime (%)	Utilization (%)	Active users	Personnel full/half/part	Accounting
ql-front	-	-	0	1/0/0	Unavailable
yuca	9.28% *	11.76% *	56	4/5/2	Review *
Hypatia	4.10%	41.04%	113	2/0/0	Reconciled
Clemente	0.04% *	9.29% *	14	1/0/0	Review *
Sersic	-	-	-	1/0/0	Unavailable
Mirta	-	-	-	1/0/0	Unavailable
SDumont	4.12%	84.70%	405	6/0/0	Reconciled
Santos Dumont	25.57% *	37.10% *	217	6/0/0	Review *
cedia	0.34%	41.56%	94	0/0/1	Reconciled
pirayu	33.04%	78.87%	57	3/0/0	Reconciled
serafin	2.57%	96.18%	188	2/2/0	Reconciled
mendieta	3.60%	78.86%	86	2/2/0	Reconciled
eulogia	7.05%	74.00%	53	2/2/0	Reconciled
Ada Lovelace	-	-	502	7/1/0	Unavailable
atocatl	-	-	30	3/0/0	Unavailable
tochtli	-	-	60	3/0/0	Unavailable
NPAD-UFRN	20.45%	43.31%	219	2/2/1	Reconciled
ClusterUY	1.86%	22.21%	117	0/2/1	Reconciled
Abaporu	5.27% *	10.94% *	2	1/0/0	Review *
Recod	6.39% *	15.90% *	0	1/0/0	Review *
Sorgan	0.62% *	18.26% *	12	3/0/0	Review *
khipu	3.07%	6.72%	26	0/1/0	Reconciled
Clementina XXI	10.68%	42.93%	116	3/5/3	Reconciled
FIUNER	-	-	3	1/0/0	Unavailable
Kabré	-	-	101	2/0/0	Unavailable
TMC- FIERRO	-	-	80	1/1/0	Unavailable
Guane	33.40%	68.04%	16	1/0/0	Reconciled
PACCA	50.79%	0.67%	1	0/0/1	Reconciled
Lobocarneiro	-	-	53	2/0/0	Unavailable
leftraru	16.55% *	95.02% *	367	8/0/0	Review *
lamcad	2.14%	72.94%	70	3/0/0	Reconciled
cempa	1.01%	21.20%	2	3/0/0	Reconciled
coaraci	0.37%	57.95%	192	10/0/0	Reconciled
odin	12.40%	30.65%	34	0/0/2	Reconciled

Source: participating centers, SCALAC HPC survey 2026 [1]. Rates are calculated from submitted accounting data at full precision and displayed to two decimals. Active-user counts are not deduplicated across clusters. TMC- FIERRO's historical 2024 text is excluded from 2025 usage rates.

Table A3. Unvalidated power submissions and reported PUE. Power numbers are retained only for confirmation by the submitting centers; units and measurement boundaries are unconfirmed. Do not compare, sum or convert these power numbers. PUE is dimensionless [5].

Cluster	Facility max power (unit unconfirmed)	PUE	Cluster peak power (unit unconfirmed)
ql-front	15.00	1.80	0.00
yuca	315.00	1.80	56.77
Hypatia	-	-	-
Clemente	4.20	1.80	0.88
Sersic	4.20	1.80	-
Mirta	4.20	1.80	-
SDumont	1,000.00	1.10	-
Santos Dumont	1,000.00	1.10	-
cedia	350.00	1.90	5.00
pirayu	49.00	2.00	49.00
serafin	200.00	1.72	45.00
mendieta	200.00	1.72	-
eulogia	200.00	1.72	-
Ada Lovelace	350.00	1.43	58.00
atocatl	176.00	1.80	21.00
tochtli	176.00	1.80	45.00
NPAD-UFRN	440.00	1.90	132.00
ClusterUY	-	-	-
Abaporu	47.00	1.40	36.00
Recod	47.00	1.40	15.00
Sorgan	-	-	-
khipu	-	-	-
Clementina XXI	-	-	-
FIUNER	5.60	1.80	2.50
Kabré	65.12	1.29	-
TMC- FIERRO	100.00	2.00	20.00
Guane	40.00	-	32.70
PACCA	-	-	8.20
Lobocarneiro	60.00	1.50	60.00
leftraru	178.00	1.12	88.55
lamcad	45.00	1.80	18.00
cempa	45.00	1.80	17.00
coaraci	350.00	1.57	130.00
odin	35.00	1.60	35.00

Source: participating centers, SCALAC HPC survey 2026 [1]. Zero is retained where explicitly supplied. Missing values are not estimated.

Appendix B. Data validation and follow-up

B1. Accounting reconciliation

Table B1. Material accounting residuals. Residual is the sum of allocated, down, planned-down, idle and planned capacity-time minus total reported capacity-time. Positive values mean the component sum exceeds Reported. Values are in reported accounting units. The original figures have not been corrected.

Cluster	Component sum	Reported	Residual	Residual / Reported
yuca	13,950,690	13,658,023	+292,667	+2.14%
Clemente	491,521	2,995,920	-2,504,399	-83.59%
Santos Dumont	230,643,503	202,582,480	+28,061,023	+13.85%
Abaporu	2,808,933	2,778,101	+30,832	+1.11%
Recod	3,543,394	3,392,984	+150,410	+4.43%
Sorgan	2,347,934	2,342,082	+5,852	+0.25%
leftraru	129,889,902	113,853,449	+16,036,453	+14.09%

Source: participating centers, SCALAC HPC survey 2026 [1]. Some residuals are close to the PlannedDown values, but this pattern alone does not establish a duplicated category or justify adjusting the submitted observations. Operators should reconcile the original exports and their scheduler accounting semantics.

Downtime and utilization were independently recalculated for all 24 systems with complete accounting submissions. Seventeen submissions reconcile within one reported accounting unit, seven exceed the five-unit tolerance, and ten systems lack the numeric information needed for this analysis. These checks establish computational consistency without independently verifying the underlying measurements.

B2. Technical questions requiring validation

The survey compilation identifies the following technical questions. Their resolution has not been confirmed for this study. They are retained as validation priorities rather than treated as evidence that every associated estimate is incorrect.

System or topic	Issue to resolve
GPU FP16 methodology	Document tensor-core basis and dense versus sparse performance consistently.
Santos Dumont	Validate theoretical performance estimates and confirm commissioning date.
pirayu	Confirm node/chassis description; note reports shutdowns when electricity is expensive.
atocatl	Complete processor specifications and validate the performance estimates.
tochtli	Confirm the correspondence between the reported GPU hardware and FP16 estimate.
ClusterUY	Validate the GPU estimate and account for newer GPUs reported as installed in late 2025; this configuration was not necessarily available throughout the year.
Recod	Validate the GPU estimate against the submitted hardware configuration.
TMC- FIERRO	Processor models need clarification; usage text explicitly refers to 2024.
Guane	Review GPU estimate against heterogeneous hardware inventory.

System or topic	Issue to resolve
Lobocarneiro	Confirm processor specifications used in the performance estimate.
lamcad and cempa	Confirm that CPU capacity includes compute nodes only, excluding head and storage nodes.
coaraci	Confirm whether the submitted GPU estimate resolves the correction identified during compilation.
odin	Confirm the hardware configuration with the operating center.
ql-front	Commissioned in 2026, after the 2025 usage period; retained as post-period inventory information. Zero users and zero peak consumption do not describe 2025 operation.
Power units	Confirm each power unit and whether the quantity is active or apparent power, plus facility boundaries and measurement dates. No regional power or energy calculation uses these unvalidated values.

B3. Comparability over time

The historical comparison uses the published system list and the two top-ten performance tables. The 28 identifiable tabulated systems are the sole baseline for coverage calculations, including Figure 3. The prior narrative count of 29 is retained only as an unresolved source discrepancy. Its narrative identifies four systems meeting the operational thresholds, while five satisfy them in the published table. The comparison identifies both discrepancies and uses table values for system-level calculations.

Provisional name matches include Pirayú-pirayu, Serafín-serafin, Eulogia-eulogia, Lovelace-Ada Lovelace, Lobo Carneiro-Lobocarneiro, MendietaF2-mendieta, Guane-1-Guane, Coaraci-coaraci and Leftraru-leftraru. These match reported system lineages rather than certify unchanged hardware. Coaraci's institutional attribution differs between editions and requires confirmation. Hypatia, PACCA and ClusterUY retain their names. CeCAR-odin is not assumed to be a confirmed system match. LNCC's single historical Santos Dumont entry is not mapped one-to-one to either of its two current submissions.

Capacity changes are calculated from the published estimates; small differences can reflect rounding or method revisions. Usage-rate differences compare the complete calendar-year reference periods 2024 and 2025. Actual time in service, accounting extract dates and hardware changes still require system-level verification. No complete historical capacity total or measured-benchmark change is inferred from the two top-ten lists.

For the next collection, each record should include stable cluster and facility identifiers, extract start/end timestamps, commissioning or retirement dates, scheduler and version, accounting units, active-user definition, hardware-estimate method and shared-team identifiers. Retaining those fields will make a matched comparison possible without confusing coverage changes with operational progress.