International Journal of Applied and Advanced Scientific Research (IJAASR)
Impact Factor: 5.655, ISSN (Online): 2456 - 3080
(www.dvpublication.com) Volume 7, Issue 2, 2022

COMPARING CLOUD NETWORK LATENCY ACROSS THREE DATA CENTRE ENDPOINTS LOCATED IN INDIA USING GLOBAL PROBES

Divya Vishwanath* & Paul Johnson Asamoah**
* Janardhan Rai Nagar Rajasthan Vidyapeeth, Udaipur, Rajasthan, India
** Power Life University College, Accra, Ghana, West Africa

Cite This Article: Divya Vishwanath & Paul Johnson Asamoah, "Comparing Cloud Network Latency Across Three Data Centre Endpoints Located in India Using Global Probes", International Journal of Applied and Advanced Scientific Research, Volume 7, Issue 2, Page Number 84-90, 2022.

Abstract:

This analysis compares cloud network latency across three endpoints in India to inform endpoint selection for globally distributed services. Using a secondary dataset of probe measurements, it applies paired mean comparisons to samples of 8,383 to 8,405 probes. Mean recorded latency for AsiaPacificMumbai was 29.07 milliseconds lower than for Chennai-CHE01 and 166.43 milliseconds lower than for asia-south1. Chennai-CHE01 was 137.54 milliseconds lower than asia-south1. All three differences remained statistically significant after adjustment for multiple comparisons. Repeating the analysis with 8,378 probes observed at all three endpoints produced similar mean differences. The ordering did not hold for every probe, indicating that aggregate rankings can conceal variation across measurement points. The findings extend comparative cloud benchmarking through matched endpoint comparisons within one country. They do not establish whether routing, geographic distance or infrastructure caused the observed differences. Cloud users and service operators should assess candidate endpoints from their intended user networks before deployment. The measurement dates and construction of the recorded values require verification before the findings can be generalized across periods.

Key Words: Cloud Benchmarking, Cloud Endpoints, India, Network Latency, Paired Comparisons

1. Introduction:

Network latency shapes the performance of cloud applications that depend on timely communication between users and remote infrastructure. Selecting a cloud endpoint therefore involves more than identifying a data centre within a preferred country. The network path between the user and the endpoint can affect the delay experienced by an application. Comparative evidence is needed to determine whether endpoints in the same national setting provide similar latency to the same measurement locations.

Cloud benchmarking research has established that network performance varies across services and measurement locations. Li et al. (2010) developed a framework for comparing cloud providers on measures relevant to application performance. Palumbo et al. (2021) examined latency between geographically distributed measurement points and cloud services hosted by two providers. These studies show the value of comparing observable performance across cloud infrastructure, while also indicating that the results depend on the destinations and measurement points included in the analysis.

The relationship between an endpoint and recorded latency reflects several connected processes. Signals travel over network paths whose length and configuration can differ across destinations. Access networks, routing decisions and intermediate network segments can also influence the measured delay. Dang et al. (2021) found that geographic distance, access connectivity and the measurement platform affected cloud connectivity in their international analysis. Jin et al. (2019) identified latency degradation in different segments between clients and a cloud provider. These findings support endpoint comparison, but they do not establish which process accounts for a particular observed difference.

The closest source for the present analysis is the measurement research of Charyyev et al. (2020), which compared latency to cloud data centres and edge servers using a large set of probes. Its principal comparison concerned the latency characteristics of cloud and edge destinations. The present analysis uses the associated cloud latency file to examine a narrower question: how recorded latency differs among three selected endpoint labels in India when observations are paired by probe. This comparison is a secondary analysis of existing measurements. Its contribution lies in the specific matched endpoint comparison rather than a claim that the underlying observations are new.

This study aims to compare recorded network latency across the AsiaPacificMumbai, Chennai-CHE01 and asia-south1 endpoints using observations from common probes. The research question is: How does recorded latency differ between each pair of these endpoints for probes with observations for both destinations? The three comparisons examine AsiaPacificMumbai against Chennai-CHE01, AsiaPacificMumbai against asia-south1, and Chennai-CHE01 against asia-south1.

The analysis uses a categorical endpoint measure and the latency value recorded for each available probe and endpoint pair. Pairwise samples contain between 8,383 and 8,405 probes. Matching provides a common basis for comparing destinations, although the file does not identify probe locations, network paths or

the exact dates and construction of the recorded values. The resulting estimates are associational and do not isolate the effects of geographic distance, routing or provider infrastructure.

The matched comparisons show lower mean recorded latency for AsiaPacificMumbai than for either of the other endpoints, and lower mean recorded latency for Chennai-CHE01 than for asia-south1. The relative ordering did not hold for every probe. This variation makes the evidence relevant to cloud users beyond India: endpoint selection can benefit from measurements taken from the networks a service intends to reach. Transfer of the observed ranking to other users, periods or endpoints requires further measurement.

2. Theory, Conceptual Framework and Hypotheses Development:

Cloud network latency reflects interactions between a measurement probe, the network path and a cloud endpoint. A request travels through connected networks before reaching its destination, and the response returns through a network path. Propagation distance, routing choices and delays along those paths can affect the measured response time. Comparative benchmarking therefore evaluates endpoints through observations from common probes, allowing differences in network performance to be examined under a shared measurement design. Li and colleagues, 2010, established this comparative approach in their evaluation of public cloud services.

The focal outcome is network latency associated with each probe and endpoint pair. The explanatory construct is cloud endpoint, represented by three observed categories: AsiaPacificMumbai, Chennai-CHE01 and asia-south1. These categories identify destinations rather than separate continuous dimensions. The analytical comparisons use probes with observations for both endpoints in each comparison. The resulting differences represent associations between endpoint and latency; they do not isolate the effects of provider, location or routing.

Variation in endpoint latency can arise because traffic to each destination may follow a different path. Shorter geographic distance may reduce propagation time, while routing and connectivity can produce a different ordering. Palumbo and colleagues, 2021, found that cloud latency varied across measurement locations and cloud services. That evidence supports comparison, but it does not establish which endpoint in the present data should have the lowest latency.

AsiaPacificMumbai and Chennai-CHE01 represent two destination categories in India. Their latency may differ for the same probe because the paths to the endpoints can differ. Charyyev and colleagues, 2020, documented variation in latency across cloud locations using probe measurements. The evidence does not establish a directional ranking for these two endpoints.

H1: Network latency differs between the AsiaPacificMumbai and Chennai-CHE01 endpoints for probes with observations for both endpoints.

AsiaPacificMumbai and asia-south1 provide a second comparison between destinations labelled as being in the Mumbai region. A shared regional label does not guarantee identical latency because network access and routing can vary between endpoints. Li and colleagues, 2011, identified network performance as a relevant basis for comparing cloud services. The expected direction remains unspecified because the available evidence does not establish the relative latency of these endpoints.

H2: Network latency differs between the AsiaPacificMumbai and asia-south1 endpoints for probes with observations for both endpoints.

Chennai-CHE01 and asia-south1 provide a third comparison across the two endpoint categories. Differences may reflect the paths connecting a probe to each destination, while common network segments may reduce the observed difference. Mok and colleagues, 2021, showed that cloud network measurements depend on the locations and connectivity of measurement points. This supports probe-level comparison without assuming that one endpoint will consistently have lower latency.

H3: Network latency differs between the Chennai-CHE01 and asia-south1 endpoints for probes with observations for both endpoints.

The three hypotheses specify pairwise comparisons within a single categorical endpoint variable. They allow evidence of a difference or insufficient evidence of a difference. The available file does not identify probe locations or measure routing mechanisms, so the analysis cannot attribute an observed difference to a particular network cause.

3. Data:

This section establishes data provenance, constructs the endpoint comparison samples and defines the measures. It also documents the checks completed and the information required before hypothesis testing.

3.1 Data Sources, Population and Sample Construction:

The analysis uses the supplied cloud latency file associated with the measurement research of Charyyev et al. (2020). The authors report measurements from RIPE Atlas probes to cloud and edge locations and provide a public research repository. The supplied file contains 8,571 distinct probe identifiers. Its relationship to the 8,456 end users reported in the publication requires clarification before the file can be treated as identical to the published analytical sample.

The target population comprises network paths from measurement probes to the three selected cloud endpoints in India. The accessible sampling frame comprises the probes represented in the supplied file. The unit of analysis is a probe and endpoint pair. The file supplies one latency value per available pair, but it does not identify the dates or frequency underlying those values. It supports a comparison of recorded values rather than a time series analysis.

The three endpoint labels were selected because they match the conceptual framework. The selection was made after examining the supplied file. The exact measurement period, dataset version and applicable licence were not identified in the supplied file. The source publication dates from 2020, but that date alone does not establish the dates of every measurement.

Table 1: Sample construction

Stage	Probes	Probe and Endpoint Observations	Rule
Supplied file	8,571	Not established	Distinct probe identifiers in the original file
Selected endpoints	8,571	25,352	Retain recorded values for the three selected endpoint labels
H1 comparison	8,401	16,802	Require values for AsiaPacificMumbai and Chennai-CHE01
H2 comparison	8,405	16,810	Require values for AsiaPacificMumbai and asia-south1
H3 comparison	8,383	16,766	Require values for Chennai-CHE01 and asia-south1
Common three endpoint sample	8,378	25,134	Require a value for each selected endpoint

The comparison samples overlap. Their observations must not be added to obtain a total sample size. The table reports data availability for the hypotheses; it does not report completed model estimation.

3.2 Data Integration, Cleaning and Quality Assurance:

The supplied file stores each probe on one row and places endpoint labels and corresponding latency values in paired fields. We identified the three selected labels across those fields and associated each value with its probe identifier and endpoint label. The resulting key is the combination of probe and endpoint. No external dataset was merged.

The file contains no duplicate probe identifiers. The extracted records contain no duplicate probe and endpoint pairs, and all 25,352 extracted latency values are positive. The selected endpoint values are absent for 54 probes for AsiaPacificMumbai, 153 for Chennai-CHE01 and 154 for asia-south1. These absences were not replaced with zero or imputed.

Pairwise comparisons require both endpoint values for the same probe. A separate complete sample requires all three values. We did not apply an outlier threshold or remove recorded extremes. The supplied file does not explain whether absent values reflect failed measurements, unavailable endpoints or another cause. Probe locations and network paths are also unavailable, limiting assessment of selection effects and the reasons for latency differences.

3.3 Variable Measurement, Construction and Validation:

Endpoint is an observed categorical variable. Network latency is the recorded numeric value paired with each selected endpoint label. The source publication describes latency in milliseconds, but the supplied file does not independently document the calculation used to produce each value (Charyyev et al., 2020). The interpretation of the values as the source authors' latency summaries should therefore be checked against the original measurement documentation before submission.

Table 2: Variable dictionary

Variable	Analytical Role	Definition and Construction	Source Identifier	Unit and Frequency	Missing Value Rule	Coverage and Limitation
Probe identifier	Matching identifier	Original identifier used to associate observations across endpoints	prb	One identifier per file row	No missing identifiers observed	Probe location is not provided
Cloud endpoint	Explanatory category	Endpoint label matched to its recorded value	Endpoint label fields	Three selected categories	Retain only selected labels	Labels do not independently identify routing or provider effects
Network latency	Outcome	Numeric value paired with a selected endpoint label	Corresponding value fields	Milliseconds according to the source study; observation frequency not established	Exclude unavailable values from the relevant comparison	Underlying aggregation and measurement dates require verification

Table 3 reports the available recorded values before restricting the data to matched comparisons. Higher values indicate longer recorded latency. No latent index or factor score was constructed.

Table 3: Descriptive statistics for available endpoint values

Endpoint	Valid observations	Mean	Standard Deviation	Median	Minimum	Maximum
AsiaPacificMumbai	8,517	175.50	71.59	141.29	1.67	827.07
Chennai-CHE01	8,418	204.03	67.72	185.16	1.61	862.34
asia-south1	8,417	341.35	80.55	363.26	1.66	996.97

The values in Table 3 describe different available samples. They do not constitute matched estimates of the three hypothesis comparisons. The checks establish identifier uniqueness, endpoint matching and positive recorded values; they do not establish measurement accuracy through an independent source.

3.4 Analytical Framework, Identification and Reproducibility:

The intended analysis is associational. Each hypothesis compares latency between two endpoints using values recorded for the same probe. Probe matching accounts for the pairing in the supplied data and permits each comparison to use the shared probes identified in Table 1.

The endpoint categories combine differences that may arise from location, network routing and provider infrastructure. The file does not separately measure these mechanisms. A difference in recorded latency therefore cannot be attributed to a specific mechanism or interpreted as a causal effect. Probe coverage may also limit generalization beyond the available measurement network, consistent with the measurement limitations discussed by Mok et al. (2021).

4. Methods:

4.1 Research Design and Identification:

The analysis uses a matched observational design to compare recorded network latency across three cloud endpoints in India. The unit of analysis is a probe and endpoint pair. For each comparison, the estimand is the mean difference in recorded latency between two endpoints among probes with values for both.

Matching uses variation across endpoints observed by the same probe. This controls for characteristics shared by the probe within a comparison, but it does not isolate the effects of location, provider infrastructure or routing. Endpoint selection was not randomized, and the data do not identify network paths or probe locations. The estimates are therefore associations for the available probes rather than causal effects. Comparative measurement of cloud network performance provides a basis for endpoint evaluation, subject to the coverage of the measurement points used (Palumbo et al., 2021).

4.2 Specification and Estimation:

For each endpoint pair, the analysis calculates the difference between the two recorded values for each matched probe. Equation 1 defines the estimated mean difference:

$$\widehat{\Delta}_{ab} = \frac{1}{n_{ab}} \sum_{i \in S_{ab}} (Y_{ia} - Y_{ib})$$

In Equation 1, Y_{ia} and Y_{ib} are the recorded latency values for probe iand endpoints aand b. S_{ab} is the set of probes with values for both endpoints, andn_{ab} is the number of matched probes. A negative estimate indicates that the first endpoint has lower mean recorded latency than the second.

The three comparisons correspond to the three hypotheses: AsiaPacificMumbai against Chennai-CHE01, AsiaPacificMumbai against asia-south1, and Chennai-CHE01 against asia-south1. Each hypothesis was assessed using a two-sided paired mean test against a null mean difference of zero. The analysis calculated a 95% confidence interval using the standard error of the probe level differences and a studenttdistribution with $n_{ab} - 1$degrees of freedom.

The individual confidence intervals are unadjusted. The three tests were also assessed using the Bonferroni criterion, with a familywise significance level of 0.05 divided by three. The procedure treats probe level differences as independent across probes. It does not account for dependence arising from shared networks or geographic clustering.

4.3 Sample and Measurement Alignment:

The target population, accessible frame, selection rules and source are described in the data section. The estimation samples contain 8,401 matched probes for H1, 8,405 for H2 and 8,383 for H3. The common sample contains 8,378 probes with values for all three endpoints.

The endpoint labels define observed comparison categories. The outcome is the recorded latency value paired with each selected label, as specified in the variable dictionary. The source study reports latency in milliseconds (Charyyev et al., 2020). The supplied file does not establish the exact measurement dates or the procedure used to aggregate the recorded values. These limits constrain interpretation of the estimates and their transfer to other periods.

A large number of probes does not establish representative coverage of users or networks. The available file does not identify the reasons for missing endpoint values, so selection into each matched comparison remains a possible source of bias.

4.4 Processing, Inference and Sensitivity:

The analysis read the supplied file and matched each endpoint label with its associated numeric value and probe identifier. It then restricted the records to the three selected endpoints and arranged them by probe. For each hypothesis, it retained probes with both values, calculated the paired differences and estimated their mean, standard error, confidence interval and test statistic. Missing values were excluded from the relevant pairwise comparison; they were not imputed or replaced with zero.

The extracted records contained no duplicate probe and endpoint pairs, and all selected recorded values were positive. No outlier threshold was applied and no valid extreme value was removed. The analysis did not apply weights or transformations to latency values.

As an exploratory sensitivity check, the three mean differences were recalculated using the common sample of probes with values for every endpoint. This check assesses whether differences in pairwise sample composition materially change the mean contrasts. It does not establish that missing observations are random or eliminate selection bias. The comparisons and common sample check were selected after the data had been examined.

The calculations used Python 3.12.14, pandas 2.2.3, NumPy 2.3.5 and SciPy 1.17.0. SciPy supplied the paired mean test and the distribution used for confidence intervals; its statistical functions are documented by Virtanen et al. (2020). The analysis used no random sampling or simulation. The source file and the endpoint labels, matching rules, equations and software versions reported here provide the basis for replication; a public repository for the analysis script has not been established.

5. Findings:

This analysis compares recorded network latency across three cloud endpoints in India. It uses matched observations from the same probes to assess the three endpoint comparisons. The findings describe differences in the available measurements and do not establish their causes.

The analytical unit is a probe and endpoint pair. The supplied file contains 8,571 distinct probes, of which 8,378 have recorded values for all three selected endpoints. Pairwise samples range from 8,383 to 8,405 probes. The measurement dates and the precise construction of the recorded latency values are not established in the supplied file. The source study reports latency in milliseconds (Charyyev et al., 2020).

Table 4 presents matched comparisons. Each mean difference subtracts the second endpoint value from the first for the same probe. Confidence intervals and tests use paired mean differences and treat probes as independent observations. The comparisons and the complete sample sensitivity check were selected after examining the available data.

Table 4: Matched comparisons of network latency

Hypothesis	Endpoint Comparison	Matched Probes	Mean Difference in Milliseconds	95% Confidence Interval	Decision
H1	AsiaPacificMumbai minus Chennai-CHE01	8,401	−29.07	−30.42 to −27.73	Supported
H2	AsiaPacificMumbai minus asia-south1	8,405	−166.43	−168.70 to −164.17	Supported
H3	Chennai-CHE01 minus asia-south1	8,383	−137.54	−139.33 to −135.76	Supported

Note. Differences and confidence intervals are based on paired observations. Each two-sided paired test produced a p value below 0.001. All three comparisons also remain statistically significant after adjustment for three comparisons using the Bonferroni method.

For H1, the mean recorded latency was 174.93 milliseconds for AsiaPacificMumbai and 204.01 milliseconds for Chennai-CHE01 among the 8,401 matched probes. AsiaPacificMumbai was lower by 29.07 milliseconds on average. It had a lower recorded value for 76.90% of matched probes. The evidence supports H1, while showing that the ordering did not hold for every probe.

For H2, the mean recorded latency was 174.89 milliseconds for AsiaPacificMumbai and 341.32 milliseconds for asia-south1 among 8,405 matched probes. AsiaPacificMumbai was lower by 166.43 milliseconds on average and had a lower recorded value for 89.86% of those probes. The evidence supports H2.

For H3, the mean recorded latency was 203.78 milliseconds for Chennai-CHE01 and 341.32 milliseconds for asia-south1 among 8,383 matched probes. Chennai-CHE01 was lower by 137.54 milliseconds on average and had a lower recorded value for 95.67% of those probes. The evidence supports H3.

The estimated differences have practical meaning beyond their statistical significance. The two comparisons involving asia-south1 show larger mean gaps than the comparison between AsiaPacificMumbai and Chennai-CHE01. Restricting all comparisons to the 8,378 probes with values for every endpoint produced

mean differences of −28.99, −166.53 and −137.54 milliseconds for H1, H2 and H3 respectively. This check indicates that the differing pairwise sample sizes have little effect on the observed mean differences. It does not resolve potential selection bias from probes missing one or more endpoint values.

The findings are consistent with prior evidence that measured cloud latency varies across destinations and measurement points (Palumbo et al., 2021). They do not identify whether the observed differences arise from geographic distance, routing, provider infrastructure or other network conditions. The confidence intervals also depend on the assumption that probes provide independent observations. Shared networks or geographic clustering could reduce their effective independence.

6. Discussion:

Cloud endpoints within the same country can present substantially different network latency to the same measurement probes. The matched comparisons in Table 1 show that the gap involving asia-south1 was larger than the gap between AsiaPacificMumbai and Chennai-CHE01. This analysis adds an endpoint specific comparison to cloud benchmarking by showing how the relative performance of destinations in India varies across common probes. The contribution is a more precise account of the recorded endpoints, not a claim that all services in either city or all services from a provider share the same performance.

The observed differences are consistent with network paths contributing to latency, but the data do not identify the contribution of any individual path component. Geographic distance, access networks and routing may operate together, and their effects cannot be separated from the available endpoint values. Jin et al. (2019) demonstrated that latency degradation can arise in different segments between users and cloud networks. The present comparisons extend the question to matched destinations in India while leaving the responsible segments unresolved.

The asymmetry across comparisons suggests that a country or regional label alone may be insufficient for selecting a low latency cloud endpoint. However, the relatively large gap involving asia-south1 should not be interpreted as the effect of location or provider infrastructure. Dang et al. (2021) found that geographic distance, measurement platform and access networks can each affect cloud connectivity. Their findings provide competing explanations for the pattern in Table 1, rather than evidence that any one explanation caused it.

The results also show why comparative performance should be assessed at probe level. AsiaPacificMumbai had lower recorded latency than Chennai-CHE01 for most matched probes, but the ordering was not universal. An endpoint with a lower overall mean may therefore perform differently for particular network locations. This distinction matters for cloud benchmarking because an average comparison can conceal variation across users. Uhlir et al. (2016) similarly emphasize the need for latency measurements suited to the cloud service being evaluated.

The international contribution lies in the method of comparing the same probes across specific destinations, which can be applied to endpoint selection in other markets. Decision makers should use matched measurements from their intended user networks before choosing among endpoints, rather than relying only on regional labels or aggregate averages. Further research can connect probe locations, network paths and repeated measurements to endpoint comparisons. Such evidence would help determine whether the observed gaps persist across users and time and which network conditions account for them.

7. Conclusion and Implications:

Cloud services connect users across borders, but their network performance cannot be assumed from the location of a data centre alone. The analysis shows that recorded latency differed across the three selected endpoints when measured from common probes. The comparisons also show that an endpoint with a lower average latency did not have a lower value for every probe. This finding adds a specific matched comparison to cloud benchmarking and reinforces the need to consider the relationship between a destination and the networks accessing it. It does not establish why the differences occurred.

For cloud users and service operators, the practical implication is to compare candidate endpoints from the networks they intend to serve before making deployment decisions. Such testing can reveal whether an aggregate ranking reflects the experience of their intended users. Providers can use comparable measurements to identify destinations that merit closer investigation. For public institutions procuring cloud services, requirements for transparent, context specific performance testing could support more informed choices. These actions may help organizations select services better suited to latency sensitive applications, although the present analysis does not measure application performance or social outcomes directly.

The recorded values do not identify probe locations, network paths, measurement dates or the causes of delay. Shared network conditions may also affect the independence of observations. Future research can combine repeated measurements with probe location and routing information to examine whether the observed rankings persist and which network segments contribute to differences. Comparisons across countries and measurement platforms would help assess whether the approach and findings transfer to other settings.

References:

1. Charyyev, B., Arslan, E., & Gunes, M. H. (2020). Latency comparison of cloud datacenters and edge servers. In Proceedings of the IEEE Global Communications Conference (pp. 1-6). IEEE. https://doi.org/10.1109/GLOBECOM42002.2020.9322406
2. Dang, T. K., Mohan, N., Corneo, L., Zavodovski, A., Ott, J., & Kangasharju, J. (2021). Cloudy with a chance of short RTTs: Analyzing cloud connectivity in the internet. In Proceedings of the ACM Internet Measurement Conference (pp. 62-79). Association for Computing Machinery. https://doi.org/10.1145/3487552.3487854
3. Jin, Y., Renganathan, S., Ananthanarayanan, G., Jiang, J., Padmanabhan, V. N., Schröder, M., Calder, M., & Krishnamurthy, A. (2019). Zooming in on wide-area latencies to a global cloud provider. In Proceedings of the ACM Special Interest Group on Data Communication. Association for Computing Machinery. https://doi.org/10.1145/3341302.3342073
4. Li, A., Yang, X., Kandula, S., & Zhang, M. (2010). CloudCmp: Comparing public cloud providers. In Proceedings of the ACM SIGCOMM Internet Measurement Conference (pp. 1-14). Association for Computing Machinery. https://doi.org/10.1145/1879141.1879143
5. Li, A., Yang, X., Kandula, S., & Zhang, M. (2011). Comparing public cloud providers. IEEE Internet Computing, 15(2), 50-53. https://doi.org/10.1109/MIC.2011.36
6. Mok, R. K. P., Zou, H., Yang, R., Koch, T., Katz-Bassett, E., & Claffy, K. C. (2021). Measuring the network performance of Google Cloud Platform. In Proceedings of the ACM Internet Measurement Conference (pp. 54-61). Association for Computing Machinery. https://doi.org/10.1145/3487552.3487862
7. Palumbo, F., Aceto, G., Botta, A., Ciuonzo, D., Persico, V., &Pescapé, A. (2021). Characterization and analysis of cloud-to-user latency: The case of Azure and AWS. Computer Networks, 184, Article 107693. https://doi.org/10.1016/j.comnet.2020.107693
8. Uhlir, V., Tomanek, O., & Kencl, L. (2016). Latency-based benchmarking of cloud service providers. In Proceedings of the 9th International Conference on Utility and Cloud Computing (pp. 263-268). Association for Computing Machinery. https://doi.org/10.1145/2996890.3007870
9. Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., … SciPy 1.0 Contributors. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. Nature Methods, 17(3), 261-272. https://doi.org/10.1038/s41592-019-0686-2