



Service Level Agreement for Instacluster by NetApp

Last Modified: September 2022

Terms

This Service Level Agreement (“SLA”) for NetApp’s Instacluster Service (“Service”) is a part of the NetApp Cloud Services-Terms of Service (“Terms”). Capitalized terms used, but not defined in this SLA, will have the meaning assigned to them in the Terms.

If the Service does not achieve the service levels described in this SLA, then you may be eligible for a service credit.

We reserve the right to change the terms of this SLA or discontinue the SLA at our discretion. We will honor the SLA in effect at the outset of your subscription for the duration of your initial Subscription Term. However, if you renew your subscription, the version of this SLA that is in effect the time of renewal will apply throughout your renewal term.

Instacluster by NetApp Services

Instacluster by NetApp’s service standards are tiered based on the size, number and/or class of the cluster that the customer is running. The tiers recognize that larger clusters can support more consistent levels of performance and availability.

Apache Cassandra® Services

Tier ¹	Service Standards ²	Customer Requirements
Starter (Developer nodes)	<ul style="list-style-type: none"> o No guaranteed availability⁴ (99.9% target) o No latency³ SLAs 	<ul style="list-style-type: none"> o Minimum replication factor of 2 on all keyspaces o Add capacity or remove data when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests from the Instacluster product team to modify application for best practice Cassandra usage
Small (5 or less production nodes)	<ul style="list-style-type: none"> o 99.95% availability for LOCAL_QUORUM o No latency SLAs o 20% monthly fees at risk in total; 10% credit for each breach 	<ul style="list-style-type: none"> o Minimum replication factor of 3 on all keyspaces (please ensure that cluster is initially configured with target RF of 3) o Add capacity or remove data when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests to modify the application for best practice Cassandra usage
Enterprise (6+ production nodes)	<ul style="list-style-type: none"> o 99.99% availability for LOCAL_QUORUM consistency operations o 99% of read/write transactions to Instacluster-maintained table in the cluster within specified latency threshold³ o 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA 	<ul style="list-style-type: none"> o Minimum replication factor of 3 on all keyspaces (please ensure that cluster is initially configured with target RF of 3) o Add capacity or remove data when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests to modify the application for best practice Cassandra usage
Critical (12+ production nodes)	<ul style="list-style-type: none"> o 100% availability for LOCAL_QUORUM consistency operations 	<ul style="list-style-type: none"> o Minimum replication factor of 5 on all keyspaces (please ensure that cluster is initially configured with target RF of 5) o Separate testing and production clusters

	<ul style="list-style-type: none"> ○ Custom latency SLA negotiable (or use medium SLA)³ ○ 100% of monthly fees at risk in total; 30% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA 	<ul style="list-style-type: none"> ○ Customer notifies that they wish to receive this SLA, commissions the Instacluster product team to review their application for best practice alignment and actions finding from that review ○ Instacluster product team review prior to deploying changes that may impact latency SLA ○ Add capacity or remove data when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% ○ Comply with reasonable requests to modify the application for best practice Cassandra usage
--	--	--

**For Enterprise and Critical level Cassandra clusters, we also provide Recovery Point Objective SLA: The native replication of data in Cassandra means restoration of data from backups is rarely required. However, we will maintain backups to allow restoration of data with less than 24 hours data loss for standard backups and less than 5 minutes data loss for our Continuous Back-ups option. Should we fail to meet this recovery point objective, you will be eligible for SLA credits at 100% of monthly fees for the relevant cluster. If you have undertaken to restore testing of your cluster in the last 6 months (using our automated restore functionality) and can demonstrate that data loss during an emergency restore is outside target RPO and your verification testing, then you will be eligible for SLA credits at 500% of monthly fees.

Apache Kafka® Services

Tier ¹	Service Standards ²	Customer Requirements
Starter ^{11,18} (Developer nodes)	<ul style="list-style-type: none"> ○ No guaranteed availability⁴ (99.9% target) ○ No latency³ SLAs 	<ul style="list-style-type: none"> ○ Minimum replication of 2 on all topics ○ Add capacity or adjust retention settings when requested by the Instacluster product team to maintain disk usage in normal operations as less than 80% ○ Comply with reasonable requests to modify application for best practice Kafka usage
Small ¹⁷ (5 or less production nodes)	<ul style="list-style-type: none"> ○ 99.95% availability for writes with 2 replica consistency requirement and all reads ○ No latency SLAs ○ 20% monthly fees at risk in total; 10% credit for each breach ○ PrivateLink availability of 99.99% 	<ul style="list-style-type: none"> ○ Minimum replication of 3 on all topics ○ Add capacity or adjust retention settings when requested by the Instacluster product team to maintain disk usage in normal operations as less than 80% ○ Comply with reasonable requests to modify application for best practice Kafka usage
Enterprise ¹⁷ (6+ production nodes)	<ul style="list-style-type: none"> ○ 99.99% availability for writes with 2 replica consistency requirement and all reads ○ Availability SLA is increased to 99.999% when dedicated ZooKeeper nodes are used ○ 99% of read/write transactions to Instacluster-maintained topic in the cluster within specified latency threshold³ ○ 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA ○ PrivateLink availability of 99.99% 	<ul style="list-style-type: none"> ○ Minimum replication factor of 3 on all topics ○ Add capacity or adjust retention settings when requested by the Instacluster product team to maintain disk usage in normal operations at less than 80% ○ Comply with reasonable requests to modify the application for best practice Kafka usage
Critical ¹⁷ (12+ production nodes)	<ul style="list-style-type: none"> ○ 99.999% availability for writes with 2 replica consistency requirement and all reads 	<ul style="list-style-type: none"> ○ Minimum replication factor of 3 on all topics ○ Separate testing and production clusters

	<ul style="list-style-type: none"> ○ 99% of read/write transactions to Instaclustr-maintained topic in the cluster within specified latency threshold³ ○ 100% of monthly fees at risk in total; 30% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA ○ PrivateLink availability of 99.99% 	<ul style="list-style-type: none"> ○ Must use dedicated ZooKeeper nodes for production ○ Customer notifies that they wish to receive this SLA, commissions the Instaclustr product team to review their application for best practice alignment and actions findings from that review ○ Instaclustr product team review prior to deploying changes that may impact latency SLA ○ Add capacity or remove data when requested by the Instaclustr product team to maintain disk usage in normal operations at less than 80% ○ Comply with reasonable requests to modify application for best practice Kafka usage
--	---	---

Kafka® Connect Services

Tier ¹	Service Standards ²	Customer Requirements
Starter ¹¹ (Developer nodes)	<ul style="list-style-type: none"> ○ No guaranteed availability⁸ (99.9% target) 	<ul style="list-style-type: none"> ○ Add capacity as reasonably requested by the Instaclustr product team to manage operational loads ○ Comply with reasonable requests to modify the application for best practice Kafka Connect usage
Production Nodes	<ul style="list-style-type: none"> ○ 99.99% availability⁸ to NetApp Instaclustr maintained synthetic transaction connector ○ 20% monthly fees at risk in total; 10% credit for each breach 	<ul style="list-style-type: none"> ○ Minimum of 3 nodes ○ Add capacity as reasonably requested by the Instaclustr product team to manage operational loads ○ Comply with reasonable requests to modify the application for best practice Kafka Connect usage

Redis™ Services

Tier ¹	Service Standards ²	Customer Requirements
Starter (Developer nodes)	<ul style="list-style-type: none"> ○ No guaranteed availability (99.9% target) ○ No latency³ SLAs 	<ul style="list-style-type: none"> ○ Add capacity as reasonably requested by the Instaclustr product team to manage operational loads ○ Comply with reasonable requests to modify the application for best practice Redis usage
Enterprise (6+ production nodes)	<ul style="list-style-type: none"> ○ 99.99% availability⁹ to NetApp Instaclustr maintained synthetic transaction ○ 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA 	<ul style="list-style-type: none"> ○ The number of Redis Replica Nodes is greater than or equal to the number of Redis Master Nodes ○ Add capacity as reasonably requested by the Instaclustr product team to manage operational loads ○ Comply with reasonable requests to modify the application for best practice Redis usage

Apache ZooKeeper™ Services

Tier ¹	Service Standards ²	Customer Requirements
Starter (Developer nodes)	<ul style="list-style-type: none"> ○ No guaranteed availability (99.9% target) ○ No latency³ SLAs 	<ul style="list-style-type: none"> ○ Add capacity as reasonably requested by the Instaclustr product team to manage operational loads

		<ul style="list-style-type: none"> o Comply with reasonable requests to modify the application for best practice ZooKeeper usage
Enterprise (3+ production nodes)	<ul style="list-style-type: none"> o 99.99% availability¹⁰ o No latency³ SLAs o 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA 	<ul style="list-style-type: none"> o Add capacity as reasonably requested by the Instacluster product team to manage operational loads. o Comply with reasonable to modify the application for best practice ZooKeeper usage

PostgreSQL® Services

Tier	Service Standards	Customer Requirements
Starter (Developer nodes)	<ul style="list-style-type: none"> o No guaranteed availability (targeting 99.95%¹²) o No latency SLAs 	<ul style="list-style-type: none"> o Add capacity as reasonably requested by the Instacluster product team to manage operational loads. o Configure cluster replication and query settings which are appropriate for their data loss tolerance¹³ o Comply with reasonable requests from the Instacluster product team to modify the application for best practice PostgreSQL usage o Add capacity or remove data when requested to maintain disk usage in normal operations at less than 70%
Enterprise (2+ production nodes, or 3+ if utilizing synchronous mode strict)	<ul style="list-style-type: none"> o 99.99%¹² availability for read and write operations o 99.9%³ of read/write transactions to Instacluster-maintained table in the cluster within specified latency threshold o 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA 	<ul style="list-style-type: none"> o Add capacity as reasonably requested by the Instacluster product team to manage operational loads o Comply with reasonable requests from the Instacluster product team to modify the application for best practice PostgreSQL usage o Configure cluster replication and query settings which are appropriate for their data loss tolerance¹³ o Add capacity or remove data when requested to maintain disk usage in normal operations at less than 70%

**Availability uptime is not inclusive of one 30-minute scheduled maintenance window per month. Customers will be notified of scheduled maintenance at least 7 days in advance, and the Instacluster product team will make all reasonable attempts to minimize the impact to your availability.

OpenSearch Services¹⁴

Tier ¹	Service Standards ²	Customer Requirements
Starter (Developer nodes)	<ul style="list-style-type: none"> o No guaranteed availability (99.9% target) o No latency³ SLAs 	<ul style="list-style-type: none"> o Minimum of one replica shard on all indices o Add capacity or adjust retention settings when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests to modify application and index settings for best practice OpenSearch usage
Small (5 or less production nodes)	<ul style="list-style-type: none"> o 99.95% availability¹⁵ for search and index operations, where 	<ul style="list-style-type: none"> o Minimum of 2 replica shards on all indices o Add capacity or adjust retention settings when requested by the Instacluster product

	<ul style="list-style-type: none"> wait_for_active_shards is 2 or less o No latency SLAs o 20% monthly fees at risk in total; 10% credit for each breach 	<ul style="list-style-type: none"> team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests to modify the application and index settings for best practice OpenSearch usage
Enterprise (6+ production nodes)	<ul style="list-style-type: none"> o 99.99% availability¹⁵ for search and index operations, where wait_for_active_shards is 2 or less o 95% of index/search operations to Instacluster-maintained Index in the cluster within specified latency threshold³ o 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing breach of latency SLA 	<ul style="list-style-type: none"> o Minimum of 2 replica shards on all indices o Use dedicated masters o Add capacity or adjust retention settings when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests to modify the application and index settings for best practice OpenSearch usage
Critical (12+ production nodes)	<ul style="list-style-type: none"> o 99.999% availability¹⁵ for search and index operations, where wait_for_active_shards is 2 or less o 99% of index/search operations to Instacluster-maintained index in the cluster within specified latency threshold³ o 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing a breach of latency SLA 	<ul style="list-style-type: none"> o Minimum of 2 replica shards on all indices o Use dedicated masters o Separate testing and production clusters o Customer notifies that they wish to receive this SLA, commissions the Instacluster product team to review their application and index settings for best practice alignment and actions findings from that review o Instacluster product team review prior to deploying changes that may impact latency SLA o Add capacity or remove data when requested by the Instacluster product team to maintain disk usage in normal operations at less than 70% o Comply with reasonable requests to modify the application for best practice OpenSearch usage

Cadence Services

Tier	Service Standards	Customer Requirements
Starter (Developer Size Nodes)	<ul style="list-style-type: none"> o No guaranteed availability (99.9% target) o No latency SLAs 	<ul style="list-style-type: none"> o Add capacity as reasonably requested by the Instacluster product team to manage operational loads o Comply with reasonable requests to modify the application for best practice Cadence usage o Comply with requirements for Starter level SLAs for dependency services (Cassandra, Kafka, OpenSearch)
Production (3+ Production size nodes for each of Cadence and its dependency service clusters)	<ul style="list-style-type: none"> o 99.95% availability as measured by NetApp Instacluster synthetic transaction monitoring o 30% of monthly fees at risk in total; 15% credit for each incident causing breach of availability SLA and 10% credit for each incident causing a breach of latency SLA 	<ul style="list-style-type: none"> o Add capacity as reasonably requested by the Instacluster product team to manage operational loads o Comply with reasonable requests to modify the application for best practice Cadence usage o Comply with requirements for at least Small level SLAs for dependency services (Cassandra, Kafka, OpenSearch)



¹SLA tier is per-cluster and based on the number of nodes in the cluster. Customer credits are calculated based on the fees payable for the cluster or clusters impacted by the incident. SLAs credits apply to production clusters only.

²Service levels are measured on a monthly basis based on Instacluster by NetApp's monitoring systems. All service levels exclude outages caused by non-availability of service at the underlying cloud provide region level or availability zone level in regions which only support two availability zones.

³Latency is measured at a minimum rate of one read/write pair per node per 20 second period. Latency SLA excludes incidents where the cause is determined to be changes to a customer's application or unusually high loads on the cluster.

⁴Availability is measured by Instacluster's synthetic monitoring at a minimum rate of one read/write pair per node per 20 second period. A cluster is considered to be unavailable where read/write operations fail for a majority of nodes in the cluster in a given check-in period.

⁵Where a customer meets requirements for a tier based on cluster size but does not meet other requirements for a tier, the highest level of SLA where all requirements are met will apply.

⁶All SLAs exclude issues caused by customer actions including but not limited to attempting to operate a cluster beyond available processing or storage capacity.

⁷SLA credits must be claimed by customers within 14 days of the end of the relevant calendar month.

⁸For Kafka Connect, availability is measured by Instacluster's synthetic monitoring at a minimum rate of one Connector read or write operation per cluster per 20 second period. Excludes issues caused by BYO Kafka Connect Connectors due to the potential impact of user code on the availability of these environments.

⁹For Redis, availability is measured by Instacluster's synthetic monitoring at a minimum rate of one read or write operation per cluster per 20 second period. Excludes latency issues caused by the use of integrated Lua scripting (EVAL and EVALSHA). Excludes issues caused by customers executing commands marked as "dangerous" by the Redis project (turning on authentication will restrict access to these commands). Details of these commands can be [found here](#).

¹⁰For ZooKeeper, availability is measured by establishing a connection with the ZooKeeper server on each node using a local ZooKeeper client, on a per node per 20 second basis.

¹¹For preview versions of Kafka and Kafka Connect, only their respective "Starter" tier SLAs are valid. Production usage may be brought under an agreed SLA for the GA version after joint testing. Please [contact us](#) if you wish to discuss this option.

¹²PostgreSQL availability is measured by Instacluster's synthetic monitoring at a minimum rate of one read/write pair per node per 20 second period. A cluster is considered to be unavailable where read/write operations fail for all nodes in the cluster in a given check-in period. PostgreSQL SLAs exclude issues caused by customer actions including but not limited to; attempting to operate a cluster beyond available processing or storage capacity, modifications to application configuration, or customer initiated reloads or resizes. PostgreSQL uptime is not inclusive of one 30-minute scheduled maintenance window per month. Customers will be notified of scheduled maintenance at least 7 days in advance, and the Instacluster product team will make all reasonable attempts to minimize the impact to your availability.

¹³Client PostgreSQL applications should be configured in order to maintain high availability and reestablish connections in the event of a master replica failure. For more information see [Replication and High Availability](#)

¹⁴OpenSearch SLAs also apply to legacy OpenDistro for Elasticsearch clusters.

¹⁵The KNN plugin will use additional off heap memory. The default cache and selected node size may be inappropriate depending on the specific use of the plugin combined with other OpenSearch activities. This may result in cluster instability and customers need to be aware this could impact high availability of the cluster.

¹⁶Clusters created as "Bundled Use Only" are covered by SLAs only when used purely as a supporting service for another Instacluster by NetApp offering (i.e., no direct access).

¹⁷Enterprise Add-ons for Kafka (Schema Registry, REST Proxy, Karapace Schema Registry and Karapace REST Proxy) are excluded from availability and latency SLAs.

¹⁸For preview versions of Enterprise Features for Kafka, only the "Starter" tier SLAs are valid. Production usage may be brought under an agreed SLA for the GA version after joint testing. Please [contact us](#) if you wish to discuss this option.

*Apache®, Apache Cassandra®, Apache Kafka®, Apache Spark™, and Apache ZooKeeper™ are trademarks of The Apache Software Foundation. OpenSearch is a registered trademark of Amazon Web Services. Postgres®, PostgreSQL®, and the Slonik Logo are trademarks or registered trademarks of the PostgreSQL Community Association of Canada and used with their permission. Redis™ is a trademark of Redis Labs Ltd. *Any rights therein are reserved to Redis Labs Ltd. Any use by Instacluster Pty Limited is for referential purposes only and does not indicate any sponsorship, endorsement, or affiliation between Redis and Instacluster Pty Limited.*