

Managed Platform Subscription: Apache Kafka® and Kafka® Connect

Versions

Any of the generally supported versions available through the Instaclustr Managed Platform as mutually agreed between the parties.

Services Provided

Instaclustr will provide the Customer with a 24x7 managed platform on supported public cloud infrastructure for the open source technologies detailed above and related add-on technologies and services that form part of Instaclustr's standard product catalogue as at the date of execution of this SOW. Under the managed platform, Instaclustr shall be responsible for ensuring Customer's open source technology clusters (and add-on services) are operating effectively without impacting Customer's service (in consideration of the limitations of Customer's data model and use case), which generally includes:

- **24x7 Support**

Instaclustr will provide a 24x7 support capability for responding to incidents related to Customer's open source nodes (and related technologies) in the managed platform scope set out in this Agreement.

- **Monitoring and Alerting on the Open Source Nodes**

Instaclustr will monitor performance and availability of Customer's open source nodes. Alerts are sent to Instaclustr's 24x7 Support team, who will investigate and action alerts in accordance with Instaclustr's defined Incident Severity Levels as defined in Instaclustr's Support Policy (see [Service Support](#)). Monitoring metrics can be accessed on the managed platform dashboard and also through an API. Where Instaclustr identifies issues caused by Customer's use of the cluster Instaclustr will provide guidance on solutions.

- **Backups and Snapshots**

Instaclustr takes daily backups of ZooKeeper™ data (topics and other configuration). Should

Customer require a backup of data written to Kafka, Customer will be responsible for configuring and operating a Kafka consumer to maintain that back up outside of Customer's managed cluster.

- **Updates and Patch Management**

Instaclustr will update operating systems, product software, and Instaclustr's software as required, generally on a monthly basis, or as/if when issues are identified (such as a critical security vulnerability is released).

- **Health Checks and Repairs**

Instaclustr will perform daily cluster health checks covering disk usage, indicators of data model issues, and verify that repairs and backups are successfully completing.

- **Dashboard and API With Key Metrics of the Cluster**

Instaclustr will provide Customer with access to live and historical performance and capacity metrics at the per node and at table level via its managed platform dashboard and API.

- **Service Level Management and Reporting**

Instaclustr will manage and monitor the Service on an ongoing basis and provide a mechanism for Customer to verify that SLAs are being met (see [Service Levels](#)).

- **Capacity Management**

Instaclustr will assist Customer in managing capacity of Customer's open source clusters through constant monitoring and reporting on cluster performance and health. Customer shall be entirely responsible for approval of capacity increases, however Instaclustr may temporarily add capacity to your cluster without prior permission where there is an imminent risk of cluster stability issues. Instaclustr will support Customer with this by advising when the appropriate time is to upscale or downscale the cluster for predetermined events. Customer can initiate via the managed platforms platform or via a support channel a request to upscale or downscale the cluster at any time by either adding additional nodes or increasing the size (capacity) of existing nodes depending upon the product types used to implement Customer's clusters in the managed platform.

Customer will not install any additional software on managed nodes. Customer will have access to the cluster via the standard open source technology (and related technology) APIs to manage and configure pre-built and custom connectors. Besides software implemented as Kafka connectors, Customer will not install any additional software components within the managed platform.

Open Source Software

Customer acknowledges and agrees to Instacluster's use of open source software under a variety of licenses in the provision of the managed platform and that Customer is not required to directly license any open source software to use our managed platform. The open source software Instacluster use includes the following permissive open source software:

- Apache Cassandra® (under the Apache 2.0 license)
- Apache Kafka® (under the Apache 2.0 license)
- Apache Spark™ (under the Apache 2.0 license)
- Cadence® (MIT)
- Docker (under the Apache 2.0 license)
- Debian GNU/Linux (under GNU General Public License v2 and compatible licences)
- OpenSearch® Version 2.0 (under the Apache 2.0 license)
- PostgreSQL® (under the PostgreSQL License)
- Redis™ (under BSD 3-Clause "New" or "Revised" License)
- Apache ZooKeeper™ (under the Apache 2.0 license)
- A number of enabling packages including Java

Required Third Party Services

Zendesk: Instacluster uses Zendesk as its support ticketing system. Customer will be required to establish a Zendesk user account (no additional fees apply).

Service Support

Support Summary

The specific support capabilities and service levels associated with each level of support are described in the following table. In addition to what is listed there, all support tiers include:

- **24x7x365 Expert Support**

We provide around the clock expert support in English with initial response time as defined in the table below for customers contacting us via our self-service portal **support.instacluster.com** (our preferred method) or emailing us at **support@instacluster.com**. You can also use the live chat at **support.instacluster.com** to ask for assistance, but please note that our live chat is only

monitored 24 hours per day on workdays. We may be able to provide support in other languages; please contact us to find out more.

- **Instaclustr Console and API Access**

Customers may securely manage their Instaclustr services and resources through the Instaclustr console and API. Customers can view detailed live and historical performance metrics, change firewall rules, expand a cluster, and add and delete clusters.

- **Technical Documentation**

Instaclustr provides access to technical documentation and guides to assist customers in maximizing the value of their Instaclustr service. Technical documentation and guides are published and available through our support portal.

- **Security Reviews and Fixes**

In addition to performing security reviews of our features before they are implemented, we perform regular scans and tests of our environment. Any issues that arise from the scans are resolved as soon as practical. More detail on our security program can be found here: <https://www.instacluster.com/support/security/>

- **Baseline Maintenance**

Instaclustr ensures our clusters are running a recent, stable version of the core technology (Cassandra, Spark, Kafka, etc.) and underlying operating system, with the latest version of our monitoring software. We will advise of any upgrade to the version before changing the existing configuration. For application versions which are in [General Availability or Deprecated lifecycle states](#), customers may request that Instaclustr defer a minor version upgrade by up to 3 months of the proposed date. Note that for “Basic” tier support, we replace development class clusters without attempting to recover the existing cluster which may result in data loss.

- **Disaster Recovery and Backup (as applicable to the relevant service)**

The way we provide disaster recovery and data backups for our various offerings differs, as listed below:

- **Apache Cassandra**

Snapshot backups are taken once daily for all Cassandra clusters by default and retained for 7 days. Other optional backup options are provided to customers, such as the ability to trigger manual backups. More information is available [here](#).

- **Apache Kafka**
We do not provide any backup services for Kafka, aside from backing up ZooKeeper data directory once daily. Kafka's architecture inherently allows for replication between nodes.
- **Apache ZooKeeper**
The entire ZooKeeper data directory is backed up once daily by us by default and retained for 7 days. You can read more about it [here](#).
- **Cadence**
We do not carry out any backups for Cadence itself. However, for customers using Advanced Visibility, any backup settings configured for their Cassandra, Kafka, and OpenSearch clusters apply.
- **Elasticsearch**
Snapshot backups are triggered once every hour by default and retained for 7 days. More information on this is available [here](#).
- **OpenSearch**
Snapshot backups are triggered once every hour by default and retained for 7 days. Information on snapshot backups for OpenSearch clusters can be found [here](#).
- **PostgreSQL**
Full daily backups are taken by default and retained for 7 days. In addition, WAL (Walk Ahead Log) archiving is also carried out by default. Information on backups for PostgreSQL clusters can be found [here](#).
- **Redis**
By default, once daily backups are taken of your Redis cluster data using AOF (Append Only File) and retained for 7 days. Customers do have the option of configuring their own frequency for backups. More information can be found [here](#).

Support Service	Basic	Baseline	Enhanced	Premier
Definition	Applicable for all non-production node size clusters.	Applicable for all production node size clusters. ¹	Applicable for all production node size clusters, where total customer spend for the previous calendar month was: Run-In-Instaclustr-Account between 3,000 USD and 17,000 USD per month, or Run-In-Your-Own-Account between 2,500 USD and 10,000 USD per month ¹	Applicable for all production node size clusters, where total customer spend for the previous calendar month was: Run-In-Instaclustr-Account larger than 17,000 USD per month, or Run-In-Your-Own-Account larger than 10,000 USD per month
Initial response time	Best Effort	20 Minutes	15 Minutes	10 Minutes
Health and performance monitoring ²	Reduced	✓	✓	✓
Extended maintenance ³		✓	✓	✓
Named contacts ⁴	1	2	5	25
Direct access to assigned account executive ⁵				✓

Ticket Severity Levels

Support tickets are assigned a severity level that helps us prioritize issues for remediation. Urgent issues are worked 24x7 until they are resolved. High tickets are worked continuously with updates provided every 4-6 hours until they are resolved, or severity reduced. Normal and Low priority issues are worked during the working hours of the assigned Support team member.

Severity Level	Classification	Description
1	Urgent	<ul style="list-style-type: none"> • Total cluster outage or an outage having a significant impact on a customer’s business. • An event causing significant degradation of a cluster’s service with a significant impact on customer’s business. • An event identified as a significant and immediate risk of causing a total cluster outage or significant degradation in service. • A vulnerability that is currently impacting the confidentiality, integrity or availability of customer data. • An unmitigated vulnerability with an identified exploit that would impact the confidentiality, integrity, or availability of customer data. • Incidents impacting clusters with non-production SLA tier will not be classified Urgent.
2	High	<ul style="list-style-type: none"> • An event causing significant reduced availability of Instaclustr management features but not impacting customer cluster functionality. • Intermittent (but not persistent or current) significant degradation of cluster service. • A vulnerability, without an identified exploit, that materially increases the risk of impact to the confidentiality, integrity, or availability of customer data • Proposed maintenance that, if not performed in a timely manner, will cause degradation of a customer’s service. • Incidents impacting clusters with non-production SLA tier will not be classified High.
3	Normal	<ul style="list-style-type: none"> • An incident with effects that are inconvenient though not impacting a customer’s business. • Normal is the maximum priority that will be assigned to clusters with a non-production SLA tier.
4	Low	<ul style="list-style-type: none"> • Requests for information or planned changes. • Incidents or events not covered in any of the aforementioned severity levels.

Support Service Details

1. **Upgrading to Enhanced or Premier Support Tiers**

Customers on the support tier Baseline can be supported under the Enhanced tier under certain conditions. Similarly, customers qualifying for the Enhanced support tier, wishing to upgrade to the Premier support tier can be accommodated, under certain conditions. Please reach out to your customer success representative to find out how.

2. **Health and Performance Monitoring**

We actively monitor health and performance benchmarks for all our clusters to provide early detection and remediation of problems with a cluster. Where issues are identified which affect customer applications, we will notify customers through the relevant [account support contact](#).

3. **Extended Maintenance**

Extended Maintenance includes our pre-emptive intervention services to ensure clusters are running within expected limits. We allow time for manual intervention to ensure operations such as node clean-up are successful. We also provide initial assessment of and recommended resolution of high latency, high disk usage, and high CPU Usage events.

4. **Named Contacts**

Named contacts are the primary point of contact registered with the Instaclustr support team. They have the right to raise issues within the Instaclustr ticket management system. The Instaclustr Support team notifies the support contract address or emergency support contact address (as appropriate) nominated in the Instaclustr console via email where a maintenance issue is identified and for resolution of support tickets.

5. **Direct Access to Assigned Account Executive**

Premier support customers will have direct contact details for their Instaclustr account executive during business hours to escalate support issues and raise service concerns.

Support Exclusions

The following types of activities are excluded:

- Application architecture, design, and implementation
- Open source technology training

- Attendance at the Customer site
- Activities or products that are listed under Instacluster's consulting services including, but not limited to, broader advice on open source technology strategies and implementations.

For Customer-supplied connectors used with Kafka Connect:

- Any issue with the functions of the connector itself are excluded from support
- Issues that the connector causes with the general functioning of the cluster are excluded from SLA's
- Customer must ensure they have an appropriate licence for use of the connector on the Instacluster Managed cluster.

Additional Support

Services falling outside included support scope may be provided under a "fee for service" basis. If you have a request which does not fall under the defined scope of support as described above, please contact Instacluster support or your customer success representative to discuss your requirement. Services provided under this arrangement include, but are not limited to:

- Data model design and review focused on your application's requirements;
- Cluster reconfiguration;
- Specific cluster performance tuning;
- Requests for data loading in lower environments;
- BCP/DR simulations; and
- Actions from us to bring your cluster back from a catastrophic failure (not caused by our advice or actions), including:
 - Actions required due to exhausted disk space;
 - Actions required to recover from accidental deletion of data by you; and
 - Cluster recovery due to inappropriate data models or configurations.

Conditions of Running in Your Account

When running in Customer's account, Customer will have significant access to change configuration. It is a condition of Instacluster's SLAs and support for Customer's cluster that Customer does not make any changes to any object created by Instacluster's provisioning system without first agreeing those changes with Instacluster.

Instacluster will agree to any reasonable requested changes that do not pose a risk to the security or reliability of Customer's cluster or Instacluster's management systems. Some changes may result in additional charges if they will increase the management load for Customer's cluster.

Any issues caused by unauthorized changes will not be subject to SLAs and Instacluster may directly charge for support time required to resolve such issues.

Customer's Responsibilities

Maintaining the health of your cluster(s) is a cooperative effort between our customers and Instacluster. Your obligations in this cooperative effort are:

1. Maintain node disk usage below requirements for the managed application (see SLA document for details).
2. Action reasonable requests from Instacluster support to prevent issues arising and troubleshoot existing issues.
3. Use our self-service portal **support.instacluster.com** (our preferred method) or email us at **support@instacluster.com** to raise support requests.
4. Ensure that **support@instacluster.com** is whitelisted in your email service.
5. Choose a secure password for your Instacluster account that is commensurate with the sensitivity and value of your data.
6. Ensure that the logical security within your cluster (e.g., authentication, encryption) is commensurate with the sensitivity and value of your data.
7. Where you have specific requirements for backup, security, or operations, verify with Instacluster support that our offerings meet your requirements.
8. Plan for significant changes on operational load or data volumes in your cluster and provision capacity in anticipation of those loads.
9. Remove all sensitive information or data from any logs, query results, code snippets, screen shares, or any other correspondence before sending to Instacluster via support channels.

Support is not provided to customers with accounts over 30 days in arrears.

© 2023 Instacluster Copyright | Apache®, Apache Cassandra®, Apache Kafka®, Apache Spark™, and Apache ZooKeeper™ are trademarks of The Apache Software Foundation. Elasticsearch™ and Kibana™ are trademarks for Elasticsearch BV. Kubernetes® is a registered trademark of the Linux 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. Cadence is a trademark of Uber Technologies, Inc. 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. All product and service names used in this website are for identification purposes only and do not imply endorsement.