

Instaclustr Managed Platform

A Complete and Integrated Managed Platform for Deploying Next-Generation Applications With Core Open Source Technologies

The Instaclustr Managed Platform delivers a complete ecosystem for deploying your mission critical applications. We make it easy to deploy and manage the most advanced open source data layer technologies, in combination with the right infrastructure and cloud providers, advanced security, and a highly experienced support team. All of this 100% open source technology with no proprietary extensions is managed and operated in unison by the same provider without a lock-in agenda or conflicting roadmap.

The Instaclustr Managed Platform provides open source data layer capabilities encompassing **storage, caching, search, observability, stream processing, machine learning, analytics,** and **visualization.** The Instaclustr Managed Platform, and supporting processes, have been engineered to deliver these fundamental requirements for deploying next-generation applications and solutions with the core values of **scale, performance, availability, integration,** and **security.**

Why Instaclustr?

We deal with the complexity of integrating and managing these leading open source technologies and provide assurance that the environment is maintained and monitored to deliver the best performance, in a secure and transparent way. We remove the operational risks and costs associated with deploying and operating complex data layer technologies and solutions. Finally our processes and controls have been independently assessed as being compliant with SOC 2 standards to deliver a highly secure and reliable solution for enterprise.

Instaclustr allows you to enhance your data layer by adopting advanced technologies that normally require a higher initial investment in personnel and infrastructure. You can simultaneously de-risk your data layer of vendor lock-in by adopting pure open source technology. Additionally open source doesn't have per server licencing so the economics only gets better as you scale. Operationally, when you adopt the Instaclustr Managed Platform you also adopt our advanced cybersecurity posture.

Core Open Source Data Technologies

What are core open source data technologies?

Our platform and support environments provide a complete managed ecosystem for leading open source data layer technologies. We consider these to be those that have proven capability to scale in **storing, stream processing, analyzing, searching,** and **managing** the types of data and workloads required for next-generation global-scale applications and solutions.

Our offerings are always pure open source. Typically these technologies have been developed by more than one entity and are managed as a project by the Apache Software Foundation. However, we do not preclude any technology, as long as it has a suitable open source licensing arrangement. These technologies have been configured and tuned to operate efficiently and effectively together, and delivered either via leading cloud provider infrastructures or for private data center deployments.

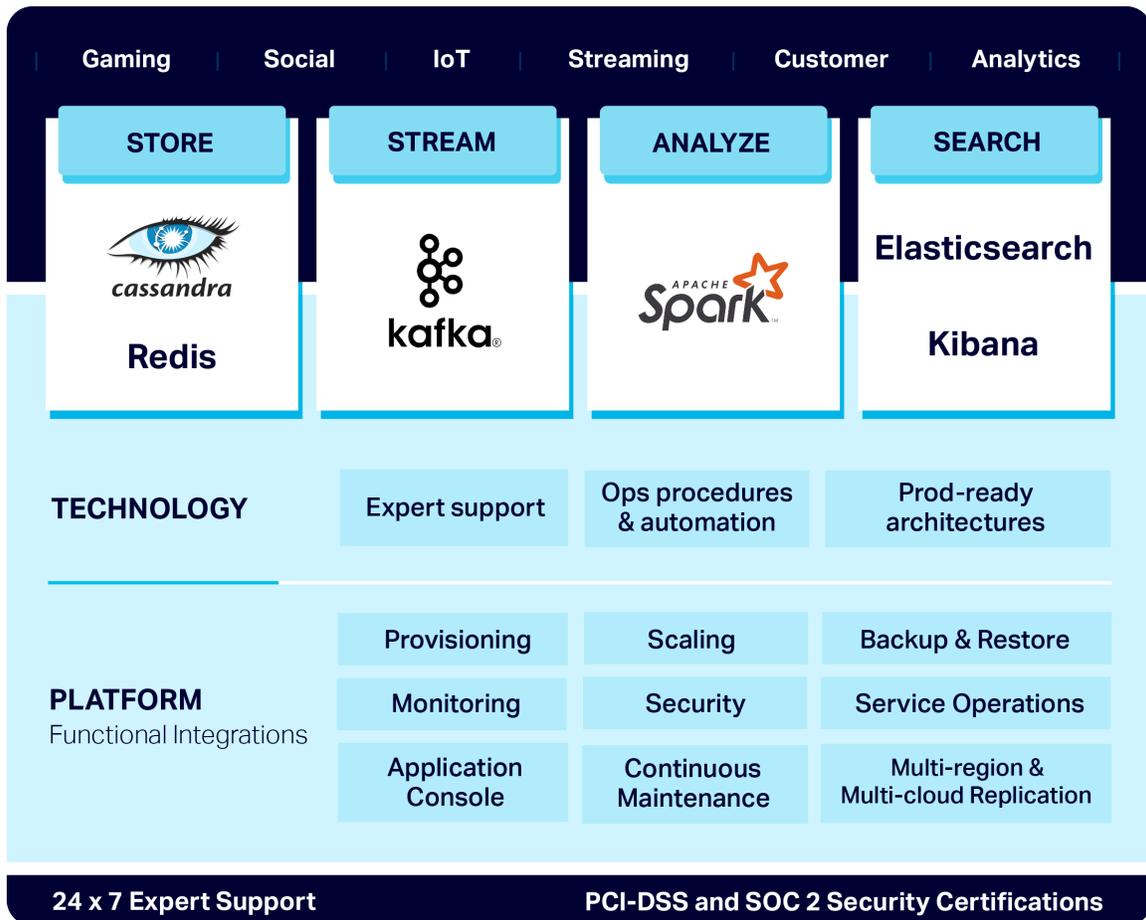
Why have we built for open source?

The best Apache Software Foundation projects have huge communities that form significant and transparent coalitions around valuable shared resources. Open source communities consist of many engineers all developing for the best outcomes without a single dominant vendor or commercial agenda.

The Open Distro for Elasticsearch has built a similar community architecture around Elasticsearch to shield users from the whims of a single vendor. The open source core of Redis is similarly beyond the influence of a single vendor.

Open source means no technology lock-in and no licensing fees which makes scaling up much more economical. The shared resource of the project code means there is a real market, with real competition over prices and customer satisfaction, so customers can select their preferred vendor. With no lock-in, swapping between vendors is possible and forward “roadmap risk” is controlled.

“Open Core” vendors try to blur the distinction between their proprietary features and the open source project you think you are buying. Often the code of their proprietary extensions might be available but is not open source, only “source available”. In terms of vendor lock-in, the “open core” story is very much like proprietary software. It is important to identify which vendors are “open core”, as pursuing an open data data layer strategy with them is the same as using closed source proprietary solutions.



The Platform for Next-Generation Applications

The Instaclustr Managed Platform consists of open source technologies that are built for scale, proven at scale, and provide the core capabilities required at the data layer to support truly ambitious and agile next-generation applications.

- STORE:** Built on the leading distributed NoSQL database, Apache Cassandra, we deliver high performance, consistent low latency, strong consistency, multi-master replication, granular security, and completely automatic self-tuning. For ultra low latency applications we offer robust in-memory storage clusters with Redis.
- SEARCH:** The Open Distro for Elasticsearch brings speed and massive scalability to classical search as well as observability applications like logs, metrics, and security. Available with our Kibana offering you can leverage the power of the ELK stack. Being completely open means you can also choose to adopt other technologies and pipelines without vendor friction.

- **STREAM:** Apache Kafka is the leading data streaming and processing technology for any scale, always-on applications. Open sourced by LinkedIn it has been adopted by major technology players as well as banking and finance. Kafka Connect will be added to this offering in 2020. Our Redis offering also has stream processing capabilities, directly inspired by Kafka, which are available at in-memory real-time speeds.
- **ANALYZE:** Instaclustr offers Apache Spark and the associated toolset for fast, in-memory processing of your application data. Spark allows you to add machine learning capabilities to your data layer that scale with your data and workloads.
- **EXPLORE:** Kibana is a powerful interface with capabilities to explore, analyze, and visualize data. Paired with Elasticsearch you can empower real-time analytics and visualizations to make your business more real-time, responsive, and evidence-driven.

INTEGRATION

With these technologies the Instaclustr Managed Platform is a one-stop-shop for all data layer capability needs. The technologies available can be integrated effectively because they are open source, with many users and contributors.

With a robust open data layer strategy in place your choice of future actions, directions and integrations are not constrained. As a provider we will also recommend the right tool for the job, e.g. whether to use Kafka or Redis for your business case that needs stream processing, whereas “open core” providers will suggest their tool is appropriate for every use case.

The Instaclustr Managed Platform still has the benefit of a vendor with a big picture overview that can assist you in your support needs across the whole business need, rather than just a piece of it. Beyond our extensive 24x7x365 support we can provide you access to experts in each of our technology offerings, capable of patching upstream projects or conducting research and development.

SCALE

Modern applications need to be able to scale efficiently and rapidly when required. Architecting the supporting fundamental data layer capabilities to enable your business to do this from scratch is a titanic task for even a single capability, let alone multiple ones. However, adopting a proprietary solution locks a fundamental aspect of your business to a single vendor, a fact that sometimes only reveals itself once you scale.

Our platform enables you to adopt core data layer technologies that are proven to scale across technology leaders like Facebook, Apple, Netflix, Twitter, LinkedIn, and much of the Fortune 500. With petabytes of customer data and tens of millions of node hours, Instacluster offers these capabilities with true economical scaling. The costs of scaling with Instacluster are clear and predictable with online pricing calculators that let you forecast scaling costs accurately.

PERFORMANCE

Open source is a wide category spanning enterprise grade projects through to student assignments posted to GitHub. Instacluster only offers open source best-of-breed technologies that are high performance and proven to scale.

In addition to deploying by default with optimized performance parameters, our teams specialize in tuning and scaling performance of these technologies to your business requirements. For seriously ambitious multi-technology plays our expert consultants can offer you advice on your architecture, or design one for you, so that it has the best performance characteristics.

CONTINUOUS AVAILABILITY

Downtime is simply not an option for modern data layer technologies. For many of our customers any downtime on their cluster or application means a significant loss of revenue. In a global 24x7x365 world not only is the “weekend upgrade” unacceptable but any outage is unacceptable. Our systems are designed for continuous performance and our industry leading SLAs are written to match. 99.9% looks good on paper but it means around 40 minutes of downtime a month. This is why we offer SLAs with 99.999% uptime guarantees and pair that with around the clock responsive support.

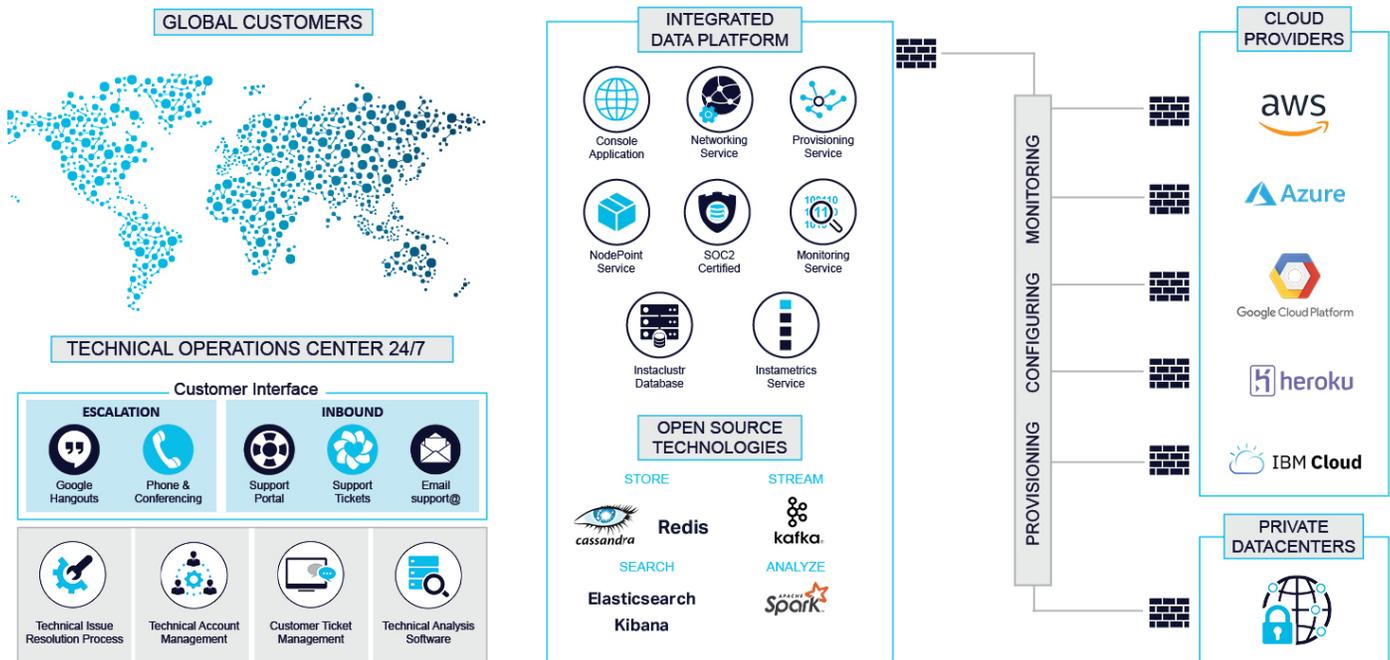
SECURITY

The data layer is important in the cybersecurity equation and the Instacluster Managed Platform allows you to take a significant step towards a more solid cybersecurity posture. “Security is in our DNA” has become a marketing phrase but at Instacluster it really is the case: three of the Instacluster founders have a background in cybersecurity. This means from our beginnings in 2012 through to the present they have baked security, monitoring, and auditing into every business process and software architecture at the company.

This means that the Instacluster platform has PCI-DSS compliant options and was the first to achieve SOC 2 compliance for our Cassandra offering. Because of this we are used by many financial and banking institutions. HIPAA compliance has also been implemented for customers, ask our representatives for our HIPAA whitepaper if this applies to your business needs.

A Managed Platform for Every Business

We offer battle tested enterprise-grade open source solutions with 24x7x365 support, security certifications, and industry leading SLAs.



Proven at Scale

At Instaclustr our managed platform alone has delivered over 70 million node hours of continual management for our customers; add our support customers and you can almost double that number. We manage thousands of nodes, hundreds of clusters, and petabytes of data for our customers across critical industries and deployment types.

Guaranteed SLAs

Our uptime guarantees and latency percentile guarantees for our customers are industry leading. We stand by our zero downtime mantra and have SLA and contractual arrangements. The 40 minutes of downtime a month that sits behind “three nines” is not good enough for our customers, we go further.

24x7x365 Expert Support

We have hands-on, real world experience with our technologies, and a depth of expertise that runs from R&D and patching upstream projects through to our dedicated 24x7 technical operations center ready to deliver expert support when required.

Core Features and Capabilities

Application Console

Instaclustr provides an easy to use console where you can access important information about your cluster, view real-time status, and performance monitors for every node and manage the configuration of your cluster.

Multi-Region and Multi-Cloud Replication

We automate adding data centers to your cluster so you can achieve geographical redundancy, or regulatory compliance, at the click of a button. Our system supports hybrid and multi-cloud clusters for applications that require systemic redundancy.

Backups and Restoration

We offer continuous and daily backup options. Restoring a backup is as simple as a few clicks in our console or a single API call.

Dynamic Scaling

Our managed solutions support dynamic scaling, achieving up to an 8x processing capacity increase in 30 minutes. Permanent horizontal scaling is also a base feature of our solution.

Continuous Maintenance

Our team is aware of the status of every single node in every cluster we manage through instrumentation that will alert on changes in performance or service levels. We actively monitor and maintain the health of all nodes and clusters to ensure that you can use the capabilities of our platform without the overheads of worrying about monitoring, tuning, preventative maintenance, or upgrades.

Advanced Monitoring and Alerting

We collect and monitor real time metrics from every node we manage. That lets us know exactly what the availability and performance of your cluster is at any time, and allows for alerting at different levels of performance degradation.

Inbuilt Deployment Security

Our inbuilt deployment security features enable data encryption at rest and in transit, user security with role-based access control, and two factor authentication. Private Network clusters are also available as an option.