

*Integrated managed platform for  
deploying next-generation applications  
with core open source technologies*

## Overview

The Instaclustr Managed Platform delivers a complete foundation for deploying your mission critical applications. We make it easy to deploy and manage the most advanced open source data-related technologies, in combination with the right infrastructure and cloud providers, advanced security and a highly experienced support team.

Our platform delivers a managed and tuned environment for data-layer open source technologies. This platform is a one-stop-shop for deploying, managing and monitoring the foundational open source technologies for the most innovative, reliable and scalable applications - all managed and operated in unison by the same provider.

## Highlights

The Instaclustr Managed Platform and supporting processes have been engineered to deliver the fundamental requirements for deploying next-generation applications and solutions:

- scale
- performance
- availability
- integration
- security

## WHY INSTACLUSTER?

We have unparalleled experience operating these leading open source technologies providing assurance that your environment is maintained and monitored to support delivery of your business objectives. We remove the operational risks associated with deploying and operating complex data layer technologies and solutions.

Our processes and controls have been independently assessed as being compliant with the SOC 2 accredited processes and controls to deliver a highly secure and reliable solution for enterprises.

## Core Open Source Data Technologies

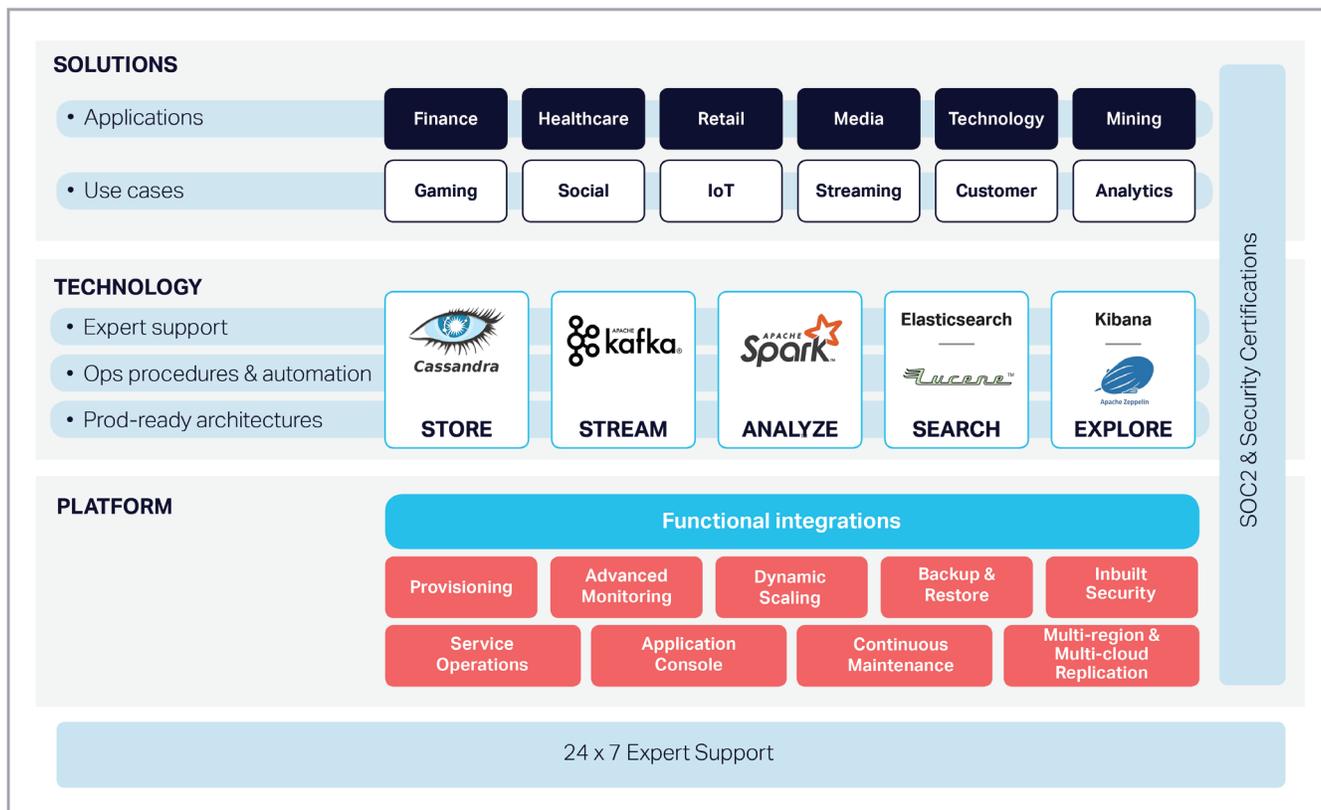
### WHAT ARE CORE OPEN SOURCE DATA TECHNOLOGIES?

At the heart of our capability is the leading core open source data technologies that we have specifically selected to be integrated into our platform. Our platform and support environments provide a complete managed ecosystem for these technologies.

Often, the technologies we provide have been developed and are managed as a project by the Apache Software Foundation, however, we do not preclude any technology, as long as it has an unparalleled track record of delivering reliability at scale, a suitable community development model and 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 open source projects have huge communities, with significant transparency. Open source communities comprise many engineers all developing for the best outcomes without a commercial agenda. There is no technology lock-in and no licensing fees.



## OUR CURRENT PLATFORM AND TECHNOLOGIES

We will always continue to innovate and add new technologies as they emerge, with a selection process that ensures only those projects that have the core characteristics of **scale**, **security**, **reliability** and **performance** are added to our platform.

Our current selected and integrated open source technologies fall into one of the following categories:

### ▶ STORE

Our foundation was built on leading NoSQL database **Apache Cassandra**<sup>®</sup>, which delivers high performance, consistent low latency, tunable consistency, and multi-region replication.

### ▶ STREAM

**Apache Kafka**<sup>®</sup> is the leading streaming and queuing technology for large-scale, always-on applications. Kafka is a system that takes streams of messages from applications (producers), stores them reliably on a central cluster (brokers) and allows those streams to be received by applications that process the messages.

### ▶ ANALYZE

The analytics and tools that can be used to process and analyze your data. The Instaclustr Managed Platform currently provides support for **Apache Spark**<sup>™</sup> for fast, in-memory processing of your application data.

### ▶ SEARCH

The platform currently has two specific search technologies available:

#### **Apache Lucene**<sup>™</sup>

A plugin for Apache Cassandra that extends the secondary index functionality to provide near real time search, built on an Apache Lucene based implementation of Cassandra secondary indexes, where each node of the cluster indexes its own data.

#### **Elasticsearch**

A search engine based on the Apache Lucene foundation. It can be used to search all kinds of documents. It provides scalable, near real-time search and sophisticated analytics functionality over unstructured data.

## ► EXPLORE

Our platform has a range of tools and environments designed to help explore and interact with your data environment:

### **Apache Zeppelin™**

A notebook user interface to allow interactive development and execution of code against both Cassandra and Spark along with data visualization capabilities.

### **Kibana**

A powerful web-interface with capabilities to explore, analyze and visualize data via search.

## CLOUD

Our platform also integrates and manages the lead cloud provider platforms. We provide support and capability for private datacenters and networked environments. The Instacluster Managed Platform is fully multi-cloud enabled supporting AWS, Azure, Google Cloud Platform and IBM Cloud.

## The Platform for Next-Generation Applications

Companies are developing now for the next wave of applications and solutions. There is a realization that the globalization of technology not only delivers immense opportunity, it also delivers a truly global competitive environment and drops the barriers for small agile players.

Now is the time to be building with the fundamental inputs for next-generation applications: integration, scale, performance, availability and security. The Instacluster Managed Platform has been engineered to deliver these core principles for all of our customers.

## INTEGRATION

The Instacluster Managed Platform provides a one-stop-shop for all data layer technologies. These technologies integrate to work efficiently and effectively together.

## SCALE

Your application must be capable of scaling efficiently and rapidly as required, and from the start. There is no longer any point to developing a solution with a technology that wasn't intended to scale. Our platform enables you to develop and stage your solution in an environment that is easily replicated into full scale production.

Our clusters and management environment can scale incrementally and dynamically, ensuring efficiency and cost effectiveness. There is no specific limit on the amount that your solution can scale with us, we currently deploy hundreds of clusters with thousands of nodes and manage petabytes of data for customers.

## **PERFORMANCE**

Having an application that services millions of transactions per second is no excuse for delay or slowness. You can't simply tell your users that they need to wait because you have too many transactions occurring. This is not an option in today's global scale application environment.

**Our clusters, nodes and networked environment are all tuned for optimal performance.**

## **CONTINUOUS AVAILABILITY**

Downtime is simply not an option, for many of our customers any downtime on their cluster or application often means lost revenue, and in some cases their whole business being unavailable to customers. Not to mention the loss of confidence and damage to brand reputation that downtime has on your user base.

**Our technologies and solutions have all been selected to ensure continuous availability.**

## **SECURITY**

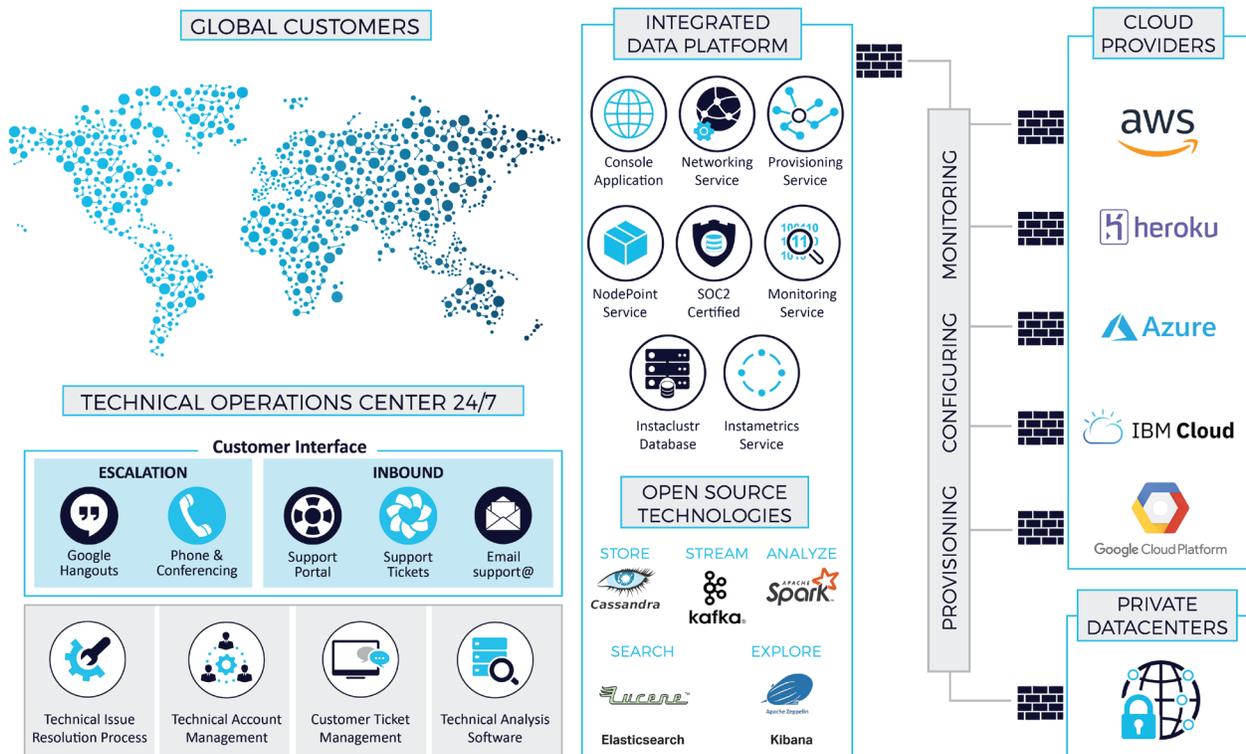
It is not an option to deploy an application or solution without assurance that cyber security requirements are being effectively and diligently managed.

We don't expect our customers to take our word for it, we have taken the steps of having our platform and processes independently assessed against the SOC2 certification, our platform incorporates a range of additional security measures and capabilities.

# A Managed Platform for Enterprises

To meet the demanding and unique needs of the large scale enterprises, you have to have proven your capability at scale and be able to provide your customers with a high level of assurance that your solution will deliver.

A truly enterprise focused solution provider will have independently assessed cyber security and will also have real experts providing round the clock dedicated support.



## PROVEN AT SCALE

At Instaclustr, our managed platform alone has delivered over 25+ 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 and hundreds of clusters for our customers across so many different industries and deployment types.

## GUARANTEED SLAs

We provide uptime guarantees and also latency percentile guarantees for our customers. We stand by our zero downtime mantra with suitable SLA and contractual arrangements.

## **24/7 EXPERT SUPPORT**

We have seen it all. The good, the bad and the ugly. Small, medium, large and extra extra large. Quiet, noisy and in between. We have hands-on, real-world experience. We operate a dedicated 24/7 technical operations centre ready to deliver expert support when required.

## **Core Features & Capabilities**

### **PROVISIONING**

You don't have to think about any of the tricky configuration of the cluster, we have done all the hard work so that you spin up a production ready cluster within minutes.

### **ADVANCED MONITORING**

We collect and monitor over one thousand different metrics from every node we manage. That lets us know exactly what the availability and performance of your cluster is at any time.

### **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.

### **BACKUP & RESTORE**

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.

### **INBUILT SECURITY**

Our inbuilt deployment security feature, enables data encryption at rest and in transit. Private Network clusters are available as an option.

### **SERVICE OPERATIONS**

We ensure that incidents follow the process and are resolved effectively and efficiently in the agreed resolution time.

### **APPLICATION CONSOLE**

The Instacluster console is 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.

## CONTINUOUS MAINTENANCE

Our team is aware of the status of every single node in every cluster we manage. We constantly maintain the health and all nodes and clusters to ensure that problems don't appear in the future.

## MULTI-REGION & MULTI-CLOUD REPLICATION

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

## Innovation & Evolution

Our solutions deliver the ability to deploy, manage and operate open source NoSQL storage, search, streaming, analytics and data exploration technologies – all integrated within the one single fully managed platform.

## HOLISTIC & STILL FLEXIBLE

With Instacluster, you can operate your development and staging environments, right next to both your analytical and operational workloads.

We have taken a holistic and flexible approach that enables our customers to simply use one managed technology, or have a complete environment that manages the entire data layer for their application environment.

## MORE TECHNOLOGIES & INTEGRATION

We continue to integrate additional open source technologies to our platform and provide suitable integrations and efficiencies. We select the best open source technologies based on the core characteristics of scale, availability and performance at the data layer.

## About Instacluster

Instacluster delivers reliability at scale through our integrated data platform of open source technologies such as [Apache Cassandra®](#), [Apache Kafka®](#), [Apache Spark™](#) and [Elasticsearch](#). Our expertise stems from delivering more than 25+ million node hours under management, allowing us to run the world's most powerful data technologies effortlessly.

We provide a range of managed, consulting and support services to help our customers develop and deploy solutions around open source technologies. Our integrated data platform, built on open source technologies, powers mission critical, highly available applications for our customers and help them achieve scalability, reliability and performance for their applications.

Apache Cassandra®, Apache Spark™, Apache Kafka®, Apache Lucene Core®, Apache Zeppelin™ are trademark of the Apache Software Foundation in the United States and/or other countries. Elasticsearch and Kibana are trademarks for Elasticsearch BV, registered in U.S and in other countries.