instaclustr

Instaclustr On-Premises Managed Platform Overview and Initial Questionnaire

Introduction

Instaclustr's Managed Platform is available through multiple public cloud providers (AWS, GCP, Azure, and IBM Cloud) and also in on-premises and other bespoke data center environments (on-premises from here in). To help potential customers decide if Instaclustr's on-premises managed platform is the right fit for them, this document provides an overview of Instaclustr's technical and service delivery approach for on-premises managed platform and some initial questions to help start the process of further qualification and scoping.

Managed Platform Scope

At a high level, moving to a Managed Platform means Instaclustr takes full responsibility for ensuring your Cassandra or Kafka is properly managed to achieve availability and latency SLAs. Your responsibility under this model is to provide and manage the compute and network infrastructure that is used by the cluster.

To fulfill our responsibilities, Instaclustr staff require a full administrative access to the system under management. This access is subject to Instaclustr's security controls audited under our SOC 2 regime. Instaclustr's system design allows customers to access specified services (Cassandra or Kafka APIs) and does not provide for customer SSH access to nodes.

Our management services also rely on the use of our cloud hosted (AWS), multi-tenant management, and monitoring infrastructure. This infrastructure receives detailed monitoring information from clusters under management, including those in on-premises data centers, and alerts our operations team of issues that require investigation. The central infrastructure also provides automated admin tooling and access control, and logging mechanisms for our operations staff.

The Instaclustr Managed Services Responsibility Matrix (supplied with this document) provides a detailed view of the division of responsibility with the Instaclustr Managed Service.

Technical Overview

The simplest approach for on-premises solutions is what we refer to as "manual provisioning" (although it is better described as semi-manual). The key features of this model are:

- The Instaclustr console and management system are still used as the sole source of configuration information for the cluster and central point for monitoring.
- Instances required for the cluster, up to the install of Debian, are provisioned by the customer's infrastructure team according to the requirements of the cluster. Required networking configuration is also undertaken by the customer's infrastructure team. Instaclustr provides detailed specifications (rack layout, vm capacity, firewall rules, etc.)
- Once the OS is installed, the instance is handed over to Instaclustr, and a script is run which writes configuration files and pulls the docker image (from our repository in Quay.io) for Instaclustr's node-agent node management application.
- Node-agent then contacts the central management system to determine required configuration, pulls Cassandra/Kafka and other docker images, and starts Cassandra/Kafka.
- The Instametrics monitoring agent is one of the installed components and begins sending metrics to the central management system for processing by the monitoring agent.
- Instaclustr operations engineers will provide ongoing management of the cluster including responding to alerts, undertaking proactive maintenance including patching and upgrades, and other cluster maintenance upgrades such as scaling.

For customers with private cloud environments (e.g. VMware, OpenStack), it may be possible to integrate the provisioning APIs of those environments with Instaclustr's management system for fully automated provisioning of clusters (i.e. click a button/ call an API and get a new cluster without any administrator intervention like the public cloud managed platform).

Pricing

- An initial on-boarding engagement is required to finalize detailed technical design and service management arrangements for all on-premises managed platform deployments. This may be undertaken as a stand-alone engagement or built in to a service commitment.
- On-ongoing management fees are aligned with Instaclustr's standard rate card for Run In Your Own Account managed service, but with a higher minimum monthly fee reflecting the fixed overhead of maintaining the bespoke service delivery model.

Other Key Assumptions

The following are some key design assumptions of the Instaclustr Managed Platform. While it is possible to vary them for specific requirements it is likely to require customization work and impact pricing:

- Customer is responsible for all infrastructure and network configuration according to Instaclustr requirements.
- Moving to the managed platform would require all nodes to be re-provisioned to Instaclustr standards. Instaclustr can normally manage this process without downtime.
- Instaclustr support engineers would have exclusive SSH access to all nodes under management (from 2 access management servers we maintain). The simplest approach is for all nodes to have public IPs, however we realize this is not suitable for many scenarios so we can also support use of a gateway server. Any customer-specific access solutions will need to be assessed for fit with Instaclustr's delivery model and toolset.
- All nodes require outbound public Internet access (this is necessary for connecting to Instaclustr monitoring service, quay.io and other similar services). Instaclustr has current planned work to tighten this requirement to specified services.
- Debian will be the operating system used, Instaclustr's standard, prepackaged configuraton will be used, and no customer-specific software will be installed.

Key Questions

- What standard arrangements are in place for external access to on-premises systems? Is it possible to vary these arrangements to use Intaclustr's standard approach or will Instaclustr need to fit with the current enterprise approach?
- Are they any enterprise requirements for auditing or other software to be installed on the servers?
- What other enterprise security and other compliance standards would apply? Is there an enterprise security checklist/review that will need to be completed?
- Who will be responsible for provisioning and management of the infrastructure for the managed clusters?
- What is the standard infrastructure used (virtualization, SAN, etc.)?

About Instaclustr

Instaclustr is the open source-as-a-service company delivering reliability at scale through our integrated data platform for technologies such as Apache Cassandra®, Apache Kafka®, Apache Spark™, Elasticsearch™, Redis™, and PostgreSQL®.

Our expertise stems from delivering more than 100 million node hours under management. We provide a range of consulting, enablement, and integration, and support services relating to 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.

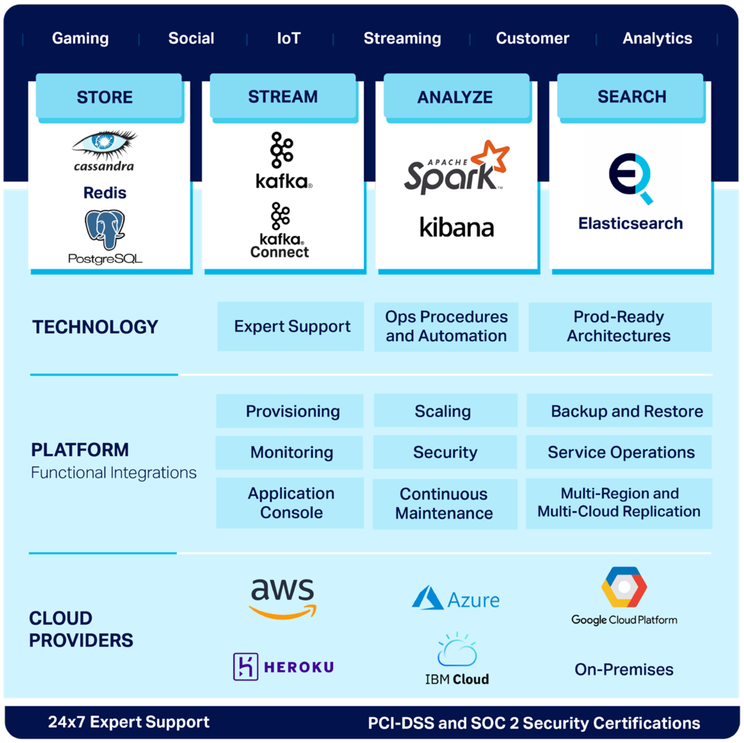
Gaming
Social
IoT
Streaming
Customer
Analytics
STORE
STREAM
ANALYZE
SEARCH
cassandra
Redis
PostgreSQL
kafka
kafka Connect
APACHE Spark
kibana
Elasticsearch
TECHNOLOGY
Expert Support
Ops Procedures and Automation
Prod-Ready Architectures
PLATFORM
Functional Integrations
Provisioning
Scaling
Backup and Restore
Monitoring
Security
Service Operations
Application Console
Continuous Maintenance
Multi-Region and Multi-Cloud Replication
CLOUD PROVIDERS
aws
Azure
Google Cloud Platform
HEROKU
IBM Cloud
On-Premises
24x7 Expert Support
PCI-DSS and SOC 2 Security Certifications