



# Fundamental Principles of Building Next-Generation Applications

## Overview

Fueled by data, the next generation of applications and solutions is being built for cloud first. Scalable, adaptable, flexible, and constantly transforming, these applications are changing the way businesses are being built and are impacting the way the customers interact and use technology in their daily lives.

Netflix, Instagram, Uber, and several others have pioneered new business models by using big data technologies to create applications that are truly next-gen. Built on a foundation of open source technologies, these applications are highly scalable, extremely reliable, and continuously available.

**Ever wonder how these types of organizations manage to create these next-gen applications? How do you establish a foundation for building a next-gen application that is secure, always on, robust, and provides the required high levels of performance? This paper discusses the core principles of building next-generation applications.**

# Principle 1: Connected and Contextual

**Businesses need to deliver personalized contextual information for each customer, resulting in hyper-connected applications that rely on data-rich technologies.**

Next-gen applications have to be built with the key ingredients of success. The most successful applications and global solutions provide contextual experience and a rich personalized experience to each individual customer. Customer interactions are customized to an individual's specific requirements, using the right mix of user, time, device, context, and message to provide better customer experiences. Getting this personalization correct at scale often leverages machine learning and artificial intelligence, which require a robust and up to date data infrastructure.

The most successful applications are continually connected, ensuring seamless user experience across all channels and touch-points. They are built to be continuously available — with downtime not being an option. In the eye of the customer, these applications are mission-critical.

- **Read our [Case Study about BlackBerry Radar IoT solution](#), designed to provide continuous visibility into the assets of an organization's transportation fleet. A perfect example of a new age connected and contextual application providing real-time information.**

# Principle 2: Accessing Multiple Data Sources in Real-Time

**Next-generation apps are consuming data from multiple and varied sources in real-time.**

Rich-data applications have complex requirements. They are built to provide a unique user experience and often must consume and process data from multiple sources—legacy applications, sensors, third parties, different databases and many more. However, it is not only the data that is critical—an app's ability to act on the data in real-time is just as important. Near real-time access to and processing of data are essential for any next-generation application.

Take for example a logistic management application built on an IoT framework. Tracking the location of goods—stored in warehouses or being transported in a complex logistics environment—needs the tracking to be constant using real-time location information for the application to be considered of value. So, while managing data is a requirement,

ensuring the real-time movement may determine the success and adoption of the solution.

- Read our [Blog Series](#) discussing Apache Kafka® Streams that can help manage the messaging stream at the warehouse.

## Principle 3: Develop a Scalable Data Layer

**Next-generation applications require a scalable data layer, providing a data-centric method for delivering data-intensive applications.**

It is all about the data. Delivering next-generation applications presents a major challenge to the traditional platforms that are rigid and not designed to handle the massive volume with the agility required for real-time data movements.

The modern application architecture needs to be capable of storing, searching, analyzing, streaming, and exploring data, all from one integrated and cohesive system. The term data fabric has emerged as a way of describing the necessity of these technologies working in unison.

The concept of building out the data fabric describes not only the data architecture operating in unison, but also the optimization of management, operations, cost, and performance—providing an architectural framework for enterprises to unify data and transform data into intelligence.

While the idea of fabric-connecting computing resources has been around since 1990, data fabric is a more recent concept. Data fabric also supports the integration of multiple data types on a vast scale, spanning multiple regions, on-premises or on-cloud, and providing applications a single and secure view of data across an organization's various applications.

The data fabric architecture for the next-gen applications incorporates both security and privacy considerations, with the data at the core of the decision-making process for access controls and sensitivity.

- The [Instaclustr Managed Platform](#) can help customers' transitioning from legacy environments to durable, secure, and optimized open source data-layer technology solutions—all delivered through the cloud and as-a-service.

# Principle 4: Provide Enhanced Data Security

**Next-generation applications are built to answer data security-related challenges associated with the rise of cloud and mobile technologies.**

Applications must be built around a security framework that provides confidence that relevant data security standards and statutory obligations associated with data protection and security are adequately addressed and are compliant.

With GDPR now in force, enterprises are giving consideration to the security and privacy impacts of an application's architectural design. A [surprising range of data now falls under this regulation](#), making enterprises focus more than ever on the data-layer and its architecture and how that data is stored and managed within the infrastructure. Apart from making changes in the architectural design, regular risk assessments and reviews can help identify the potential threats and close the gaps.

Creating an autonomous data layer with in-built security provides assurance that the environment is capable of detecting and mitigating threats. Relying on a single managed platform with the necessary controls and certifications in place can significantly improve and simplify your security and compliance risks.

- **Our Managed Platform is SOC 2 certified, ensuring that your data integrity is always protected and meets the highest possible industry standards. Read more about [how we manage security](#).**

# Principle 5: Unifying operations at scale for the data layer

**A unified and aligned approach to managing the core data technologies delivers storage, search, stream, analytics, and exploration capabilities.**

Moving into an era of personalization, machine learning, and artificial intelligence only increases the focus and requirements for a wide range of integrated and unified data technologies—delivering storage, search, stream, analysis, and exploration. Next-generation applications rely on a unified approach to managing these technologies and capabilities to deliver continuous availability, high performance, and seamless scale.

These core data-centric technologies need to be managed in unison with the application and aligned with the requirements and specifications of the system owner. Having

competing priorities with different methods for monitoring, scaling, and managing these environments can be a distraction and cause generation-operational problems.

- Read our blog [Pick'n'Mix](#) that focuses on achieving operation at scale while managing the core data layer technologies.

## Principle 6: Built on Open Source Technologies

**One fundamental that binds some of the most globally successful and viral applications is that they are built on open-source technologies.**

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 there are no licensing fees. Migrating to an open source ecosystem platform and tools not only saves costs but it can also be deployed on any operating system. Also, the development time is shorter for applications developed on open source technology.

Open source delivers on any number of parameters—including flexibility, cost-effectiveness, portability, freedom to use and reuse, security, reliability, ease of deployment, community support, and innovation. Whether it's modernizing existing applications and cloud-based infrastructure or building new, next-generation applications, open source technologies are at the core of the data layer.

- Our [infographic](#) provides the top 10 advantages of open source technologies for enterprise.

## The Instaclustr Managed Platform

**A single managed environment for open source technologies that deliver storage, search, streaming, analytics, and exploration.**

The Instaclustr Managed Platform delivers a unified management environment for the core open-source technologies that deliver storage, search, streaming, analytics, and exploration services. Our platform entails a single vendor operating, monitoring, orchestrating these technologies in alignment with your requirements and to your needs of scale, performance and availability.

The Instaclustr Managed Platform provides a foundation for deploying mission-critical applications with guaranteed SLAs and 24x7 support. Developers can rely on our platform to provide a central orchestration environment with a simple user interface that can take care of provisioning, monitoring, scaling, backup and restore, multi-region and multi-cloud replication, continuous maintenance, and more.

Developers can build and operate development and staging environments, right next to both the analytical and operational workloads. Our complete managed environment enables a company's developers and operations teams to focus on solving problems and delivering for their customers, rather than managing the complexity of the data layer.

■ Keen to understand the **Benefits of Instaclustr Managed Platform?**

Get intouch with our **Sales team.**

## ■ About ■ Instaclustr

Instaclustr helps organizations deliver applications at scale through its managed platform for open source technologies such as **Apache Cassandra®**, **Apache Kafka®**, **Apache Spark™**, **Redis™**, **OpenSearch®**, **PostgreSQL®**, and **Cadence**.

Instaclustr combines a complete data infrastructure environment with hands-on technology expertise to ensure ongoing performance and optimization. By removing the infrastructure complexity, we enable companies to focus internal development and operational resources on building cutting edge customer-facing applications at lower cost. Instaclustr customers include some of the largest and most innovative Fortune 500 companies.

© 2021 Instaclustr. 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 Instaclustr Pty Limited is for referential purposes only and does not indicate any sponsorship, endorsement or affiliation between Redis and Instaclustr Pty Limited. All product and service names used in this website are for identification purposes only and do not imply endorsement.