

Digital-Ready IT

An API Platform Story

By Maudrit Martinez

apigee

Table of Contents

Introduction: Digital-Ready IT | 2

Apigee Edge: Operating On-Premises at Cloud Scale | 4

Architectural design principles

On-premises architecture

Topology: a layout for availability, scalability, performance, and resiliency

Multi-tenancy: the flexibility of the cloud in the datacenter

Summary | 10

Introduction: Digital-Ready IT

The success or failure of business strategies in the digital economy depends on the ability of IT to adapt to rapidly changing customer behavior and ever-increasing customer expectations. The challenge for IT organizations is no longer the ROI of an IT or business strategy but how to achieve the speed, agility, and scale needed to survive in a marketplace that is customer-driven.

Businesses structured from the ground up as digital businesses are setting the pace in the marketplace. These businesses are closely followed by those quick enough to react, adapt, and transform to provide rapid customer-centric innovation.

Traditional IT infrastructure and practices at most companies are challenged to keep up with the agility required to become successful digital businesses. The needs of digital strategies and the speed of change and innovation required can even be perceived as a threat to years of investment in efficiency by IT organizations. That said, according to a global survey of IT leaders by Gartner, digital transformation that is based upon the latest and most effective business technology advances has already begun. However, according to that same survey, most CIOs and IT organizations do not feel prepared for this next era.¹

There is a way to maintain and maximize an enterprise's existing investment; a way to enable the stability of operational systems while also fostering the speed and

agility necessary to create powerful new applications and services that serve customers, partners, and employees in the digital economy. Companies that are succeeding on their digital journeys are characterized by a common set of strategic decisions around their IT infrastructure including embracing APIs and cloud technologies, leveraging big data, and implementing agile development and deployment practices.

Whether you think about this digital-ready IT as bimodal IT, two-speed IT, or traditional IT/agile IT, this new approach recognizes the value of both the back-end operational “systems of record” that have traditionally been optimized for stability and reliability and the new “systems of engagement,” which need to be optimized for agility and the fast-changing world of user engagement.



¹<http://www.gartner.com/newsroom/id/2649419>

Introduction: Digital-Ready IT

Cloud computing in all its flavors (IaaS, PaaS, and SaaS) has opened the door for IT organizations to be able to quickly react to changes in business direction, priorities, and increasing customer expectations. Combined with the power and flexibility of APIs, cloud technology is providing IT with the agility that the business needs while being able to protect and rationalize years of investment in existing systems, processes, and technology.

“Self service cloud makes impossible things instant.”

— Adrian Cockcroft

The real value of the cloud for a company’s digital strategy is not in the infrastructure or software services offered by it, but in the fact that it enables a fast and agile rhythm in the IT organization.

The investment and complexity needed to change the root of any IT organization is prohibitive and in most cases not aligned with the true objectives of a real digital strategy. The path forward requires being able to separate the elements needed to support the digital strategy (the API layer) from those passively contributing to it or not associated at all (traditional IT). Companies that embrace API- and big data-based solutions, which allow them to bring agility in-house, are on the path to succeed on their digital journey.

Not too long ago, during the early days of the Internet, success was measured by being able to establish a web presence and over several months or years attracting and retaining a few hundreds of thousands of customers. But the landscape has changed dramatically. Today, establishing a multi-channel presence and attracting only hundreds of thousands customers in a few days or months is considered failure. Success is being able to scale from zero to millions of users and billions of transactions per month in a few days or weeks.

Success requires having the right technology, approach, and processes in place from day one to ensure the business can react to customers’ needs and disruptive forces in the market with the quality, speed, and agility that both the business and customers expect.

This is the expectation regardless of whether your IT approach involves a cloud or on-premises solution. A complete and unified platform with clear visibility and offerings across the digital value chain—including the back-end operational systems, the APIs that expose a company’s data and services, the developers who consume those APIs, the applications they build, and the end users of those applications—is required.

In this eBook, we describe how the Apigee Edge API Platform allows IT to embrace API and big data solutions and enables API management operations on-premises at cloud scale.

Apigee Edge: Operating On-Premises at Cloud Scale

Operating at cloud scale doesn't come automatically by simply using a public or private cloud. After all, most of the benefits of the cloud quickly disappear when the wrong architecture is implemented or the wrong technology is used. For instance, a monolithic application that uses an RDBMS can be deployed on Amazon AWS, but that doesn't make it more scalable than it was when deployed on static hardware in a data center.

The benefits of using the cloud are primarily realized by having the right underlying architecture and technology stack that is used by cloud-capable applications and platforms. Apigee recognizes this and has clearly demonstrated it in the architectural design principles and technology stack that it has selected to build its cloud and on-premises offering.

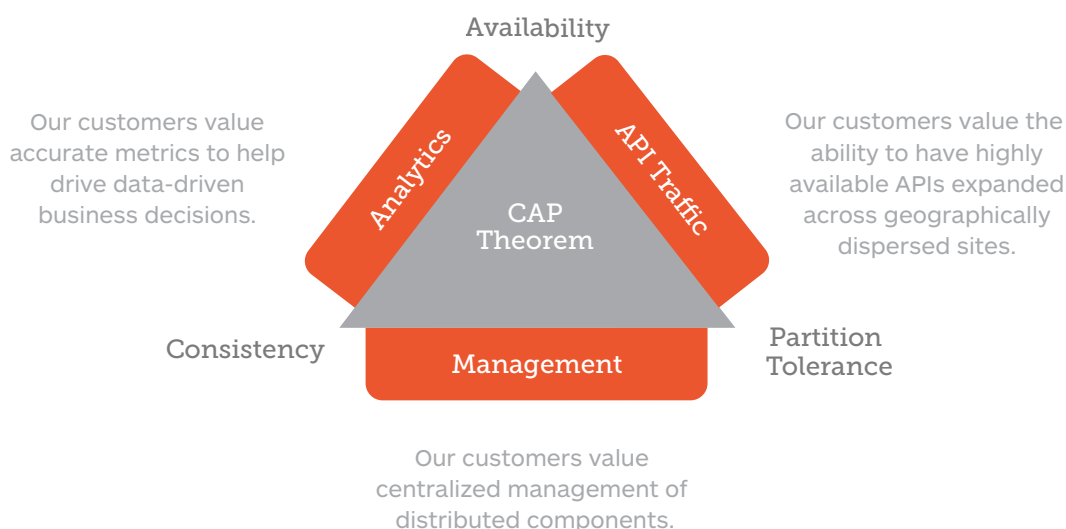
Architectural design principles

A. APIs everywhere – Components should communicate via APIs. Interactions between components and management capabilities from the command line and graphical user interfaces (GUIs) should all leverage APIs.

B. Dependencies matter – Platform components should be able to function and scale independently from each other, minimizing and controlling dependencies and using loosely coupled interfaces.

C. Eventual consistency – The platform should embrace availability and partitioning tolerance across all critical components.

D. Asynchronous execution – If something can be done asynchronously, it must be.



Apigee Edge: Operating On-Premises at Cloud Scale

E. Execution partitioning – Logically separate execution to facilitate scalability, processing, and multi-tenancy.

F. Data Partitioning – Logically separate data to facilitate scalability, processing, and multi-tenancy.

CAP theorem

In theoretical computer science, the CAP theorem, also known as Brewer's theorem, states that it is impossible for a distributed computer system to simultaneously provide all three of the following guarantees:

- **Consistency** – all nodes see the same data at the same time.
- **Availability** – a guarantee that every request receives a response about whether it succeeded or failed.
- **Partition tolerance** – the system continues to operate despite arbitrary message loss or failure of part of the system.

G. Everything Fails – The platform must keep operating even if part of it is unavailable. Apigee architecture favors graceful degradation, resiliency, and self-healing.

H. Capture and Expose Analytics and Metrics – Measure, analyze, improve. The platform must provide visibility for a business across the entire digital value chain.

I. Everything runs in the cloud and on-premises – Cloud and on-premises deployment are equally important.

J. Multi-tenancy by design – Our solution is designed to ensure that each user is isolated from other users and that each user gets guaranteed service levels.

Our understanding of our customer's needs and the current limitations around distributed systems design drives Apigee's technology selections.

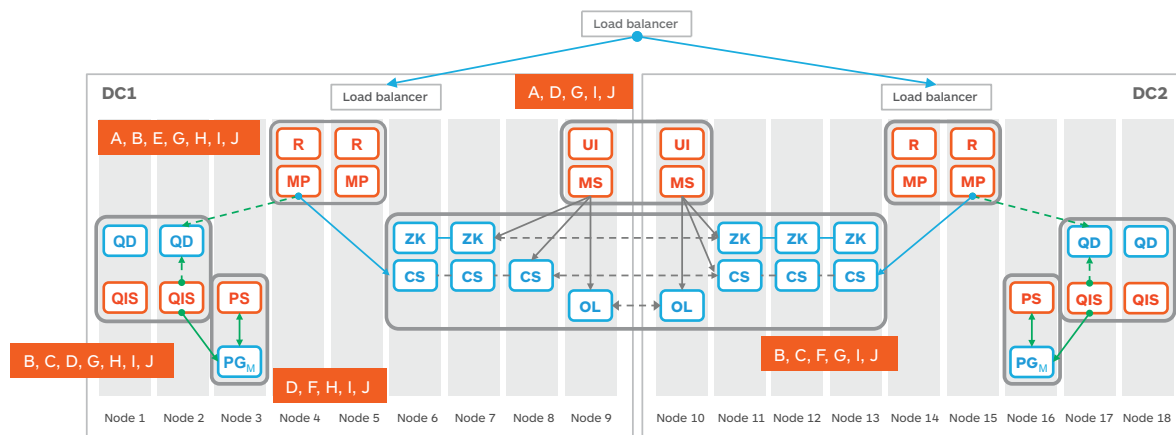
Apigee Edge: Operating On-Premises at Cloud Scale

On-premises architecture

Apigee Edge architecture brings the benefits of Apigee Edge cloud to on-premises installations. The architecture is conformed by a collection of components carefully organized around functional capabilities that align with the architecture design principles described above.

As with any other high performance system, Edge capabilities have been grouped by related functional

responsibilities and mapped to maximize the availability, scalability, and resiliency of the platform. To achieve balance, emphasis is given to the key design principles across different areas. The following figure is a conceptual representation of an Apigee Edge installation on two data centers and illustrates the influence of the design principles across different areas of the on-premises architecture.



Legend:

Router	Zookeeper	Round robin
Message Processor	Cassandra	Call direction:
Enterprise UI	ApacheDS	API call flow
Management Server	Openldap	Analytics flow
Postgres Server	PostgreSQL	Management flow
Qpid / Ingest Server	Apache Qpid	Eventual consistency/ async data replication

Architecture Design Principles:

- | | | |
|-------------------------|---------------------------|---|
| A. APIs Everywhere | D. Asynchronous Execution | G. Everything Fails |
| B. Dependencies Matter | E. Execution Partitioning | H. Analytics and Metrics |
| C. Eventual Consistency | F. Data Partitioning | I. Everything must run on the cloud and on-premises |
| | | J. Multi-tenancy |

Apigee Edge: Operating On-Premises at Cloud Scale

Apigee customers that deploy Edge on-premises usually deploy in multi-datacenter configurations with from two to six (or more) datacenters. This drives the need for partition tolerance, high availability, performance, and reliability.

Meeting these requirements demands a careful balance, and achieving it means applying the principles judiciously and thoughtfully to each of the software components.

When it comes to availability, performance, and scalability, software is just one piece of the puzzle. Achieving the benefits of the architecture design principles takes more than good software design; software, network and hardware must align on a comprehensive design. After all, as mentioned previously about a monolithic application, all the benefits of a highly scalable, available, and resilient software system quickly disappear when all of its components are deployed to a single server instance.

Topology: A layout for availability, scalability, performance, and resiliency


Edge includes API Services and Analytics Services components and can be categorized into three base functional areas: API traffic, analytics, and management. A well-designed topology makes appropriate use of available hardware and network infrastructure resources in order to fulfill non-functional requirements, such as availability, scalability, performance, and resiliency for each area.

- ▶ **API traffic:** Edge's ability to handle API traffic is primarily driven by the availability and scalability of routers and message processors. Scaling API traffic is key to scaling Edge, and therefore we recommend that customers concentrate hardware investment in this area and maximize the resources available for the platform to use.
- ▶ **Analytics:** The ability to provide visibility across the digital value chain—from back-end systems to user-facing apps—is a key factor for the success of any digital strategy. At the same time, fulfilling this need cannot compromise the availability and performance of live API traffic transactions. To this end, Apigee makes use of asynchronous fail-tolerant components dedicated to capturing, processing, and visualizing analytics data. Downtime or failures across analytics components do not affect live API traffic.
- ▶ **Management:** Being able to scale components horizontally across multiple datacenters and cloud regions requires a flexible approach to configuration and management. The ability to dynamically add and remove components and perform zero-downtime deployment is important.

Apigee Edge: Operating On-Premises at Cloud Scale

Topology


API Traffic



Analytics



Management



Legend:

Router	Zookeeper
Message Processor	Cassandra
Enterprise UI	ApacheDS
Management Server	Openldap
Postgres Server	PostgreSQL
Qpid / Ingest Server	Apache Qpid

Apigee Edge achieves this by leveraging the flexibility and distributed nature of technologies such as Apache Zookeeper and Cassandra, which, combined with components designed to expose and consume APIs, offer unparalleled flexibility. Downtime or failures across management components do not affect live API traffic.

Apigee Edge: Operating On Premises at Cloud Scale

Multi-tenancy: the flexibility of the cloud in the datacenter

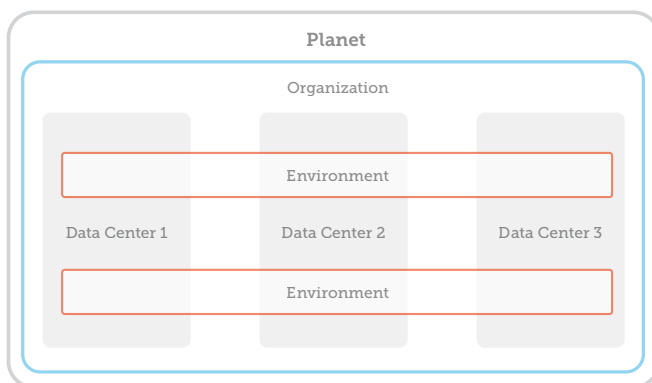
Apigee Edge has been architected from the ground up to support multi-tenancy. Multi-tenancy opens the door for on-premises customers to bring the flexibility and cost efficiency offered by Apigee Cloud into their datacenters. Edge multi-tenancy mechanisms allows operations teams to achieve API runtime processing isolation and data partitioning, while maximizing resources allocated to the platform.

To achieve multi-tenancy, Apigee has defined a number of logical boundaries such as “planets,” “organizations,” and “environments” to help customers logically slice infrastructure and platform resources.

In Apigee terminology:

- ▶ **A planet** represents a collection of resources (server instances and components) expanding across one or more datacenters.
- ▶ **An organization** represents a logical grouping of resources belonging to a tenant. Using Apigee Edge, you can create one or more organizations on top of the resources associated to a Planet.
- ▶ **An environment** represents a horizontal logical cut across an organization. Each organization can contain one or more environments. environments data is logically partitioned and distinct API traffic can be re-directed to any environment for a given organization.

A combination of organizations and environments can be used to provide API services to a number of independent tenants while sharing infrastructure resources.



Summary

The growth and strong positioning of digital native companies that leverage the power of API-based platforms has significantly changed the business landscape and demands that all companies adapt fast in order to compete.

The speed of innovation needed by the lines of business to respond to increasing customer expectations is often not aligned with traditional IT operating models. This creates the need for a dual focus, with one eye on agile innovation to meet the demands of customers and new business models, and the other on maintaining and protecting existing investments in IT infrastructure and processes, which may not need to be impacted by business transformation.

A strong customer focus, our understanding of the forces behind digital business transformation, and proven architectural design principles all drive Apigee innovation and the choices for our technology stack. Our philosophy is simple: offer innovative software solutions across core capabilities and complement with best-of-class open source technologies in those areas where it exists. This approach enables Apigee to accelerate innovation and provide our customers with proven and reliable enterprise-grade solutions.

Digital-ready IT requires a software platform that reflects the unique challenges of digital business today with room for the inevitable growth and change of tomorrow.

Apigee leads the industry with the only purpose-built architecture for digital business and the track record of customer deployments that exemplify the scale, availability, performance, and analytics visibility critical for success.

About Apigee Edge

Edge is a purpose-built platform that provides everything an enterprise needs to manage the digital value chain from API exposure to API consumption, and to measure the success of your API program with end-to-end analytics. Edge is offered as a cloud or on-premises deployment model on a single code base and includes the following components: API Services, Developer Services, and Analytics Services. API Services provides rock-solid reliability and enterprise-grade flexibility. Analytics Services enable end-to-end visibility across the digital value chain with the unified operational, developer, app performance, and business metrics that are required to monitor, measure, and manage success. Developer Services enable a developer and community experience that simplifies API discovery and learning for developers and accelerates API adoption for API providers.

About Apigee

Apigee is a leading platform for digital acceleration. Apigee empowers enterprises to gain the speed, scale, insight, and agility required to become a digital business. Through Apigee Edge API platform and Apigee Insights predictive big data analytics, Apigee helps businesses move at the new pace and scale of digital, while predicting and continuously adapting to change. Used together, APIs and predictive analytics create a powerful adaptive cycle of continuous improvement—and the faster an enterprise goes through this cycle, the faster it accelerates to become a digital business. Many of the world's leading businesses, including 20 percent of the Fortune 100, use Apigee for digital acceleration. Apigee customers include global enterprises such as Walgreens, eBay, Shell, Live Nation, Kaiser Permanente, and Sears.

For more information, visit apigee.com.

Share this eBook

