

Blueprint Architecture for Modern Commerce

A overview of the templates for deploying the commercetools API-based commerce platform

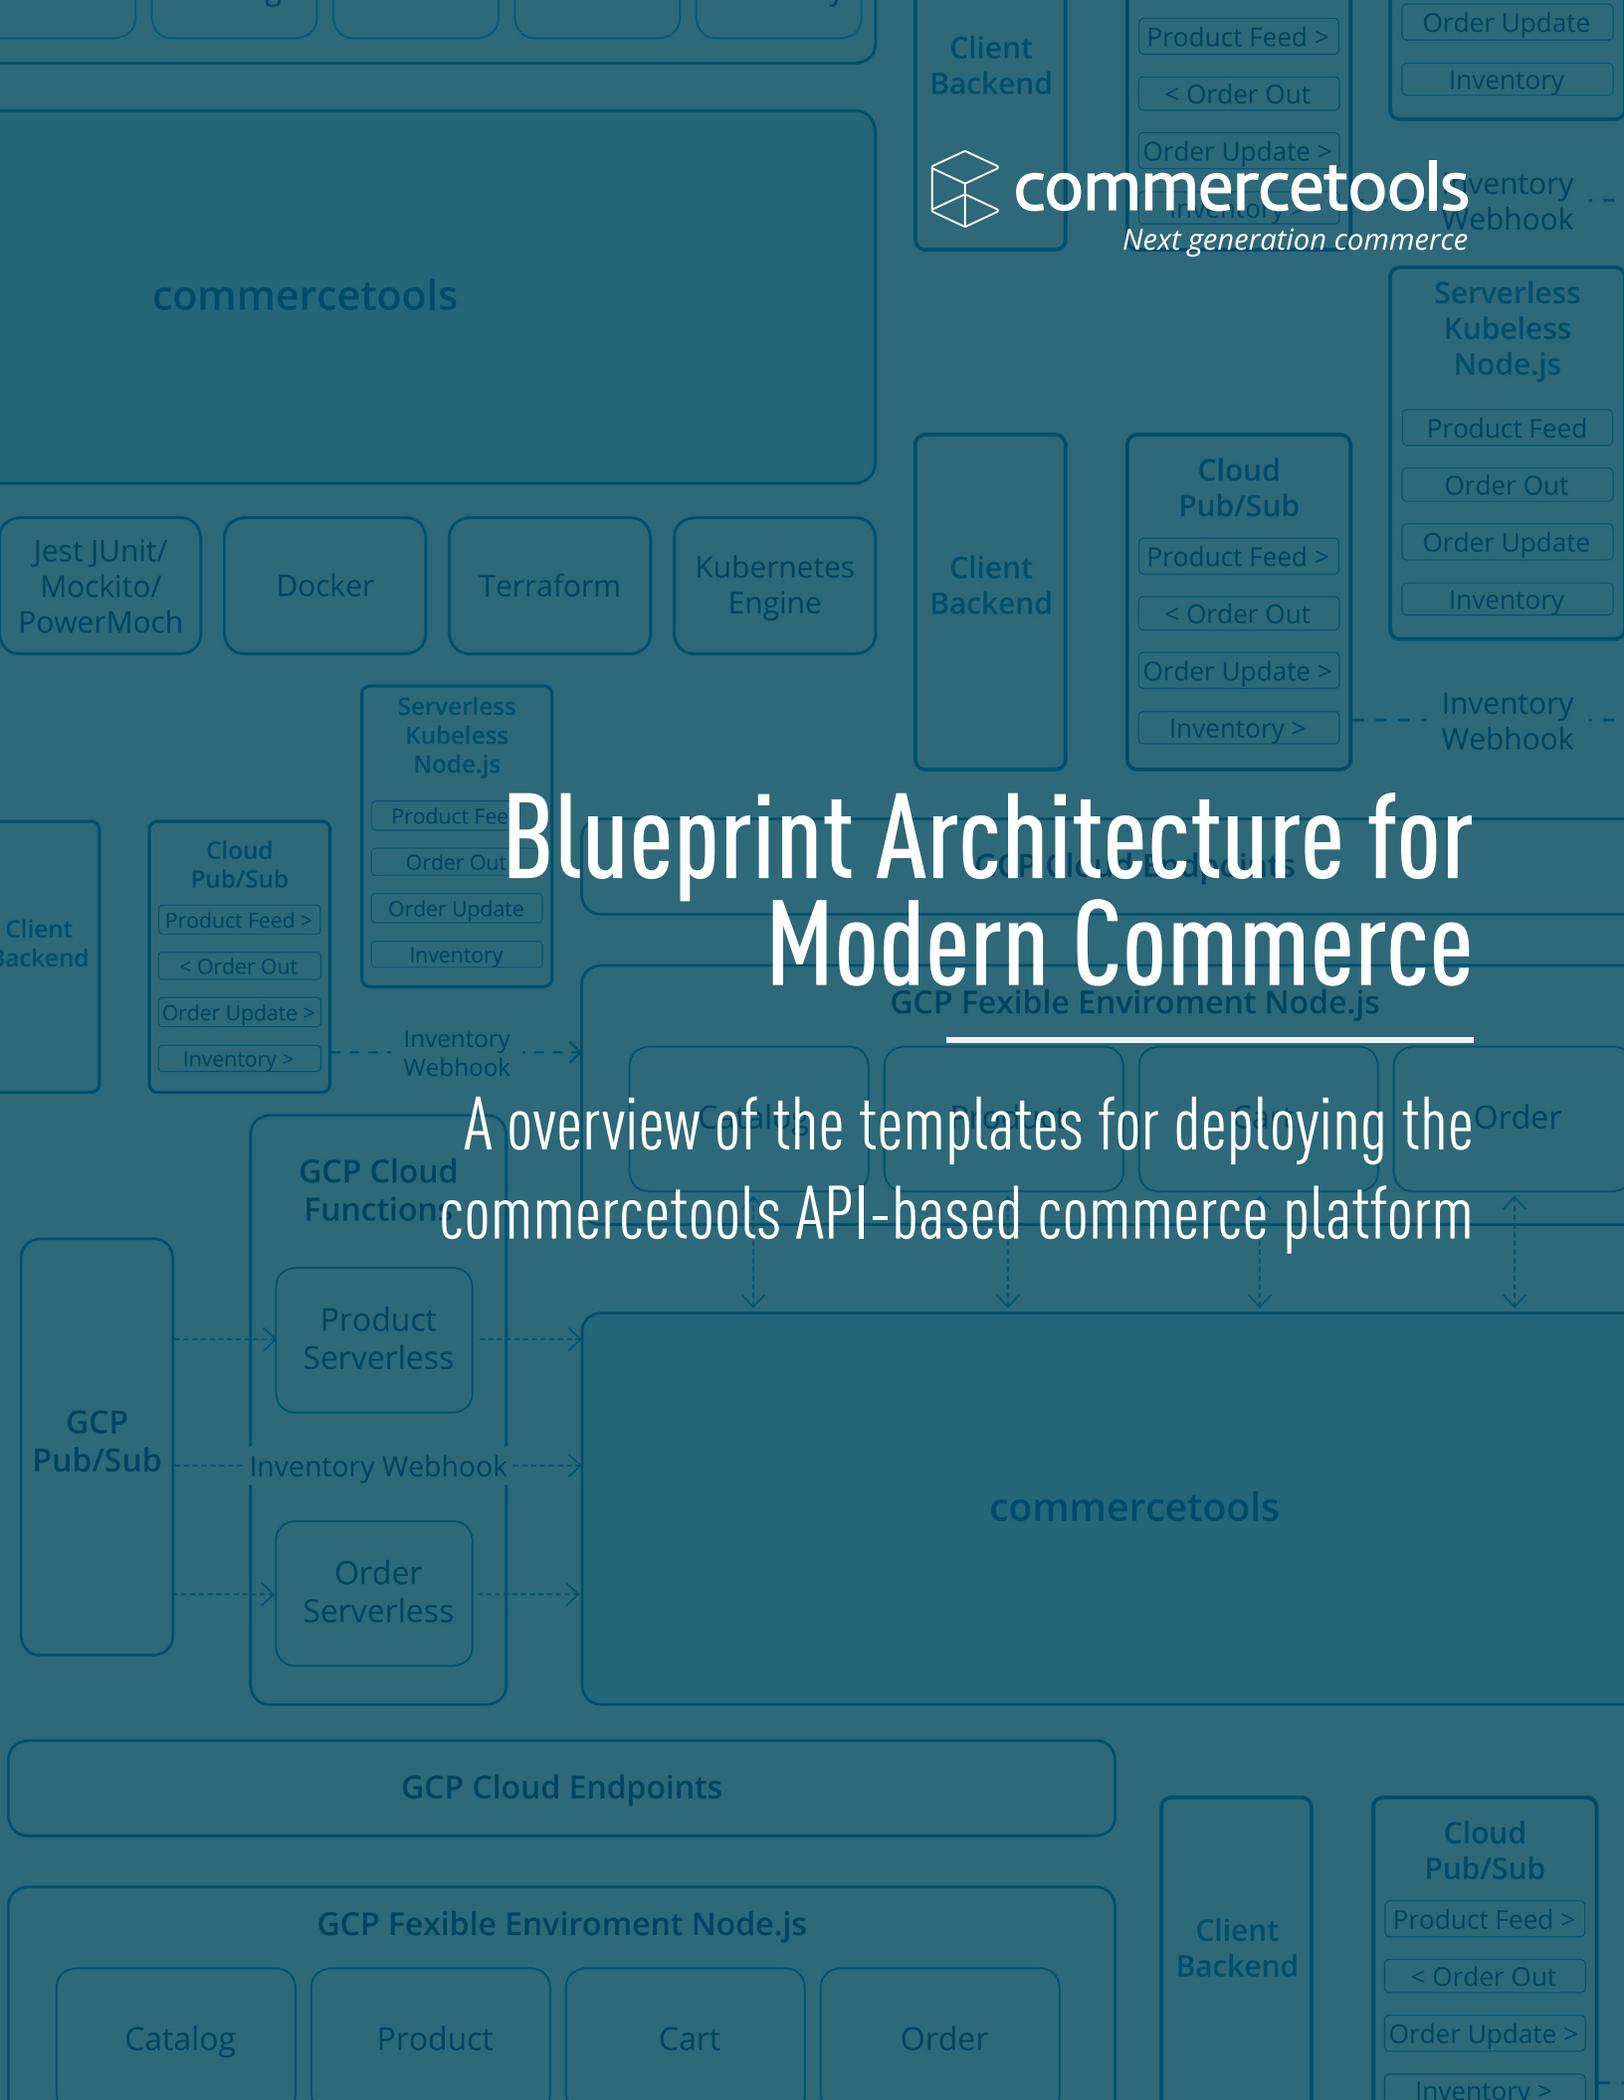


Table of Content

Executive Summary	3
Business Issues	5
Cloud Transformation	6
Continuous Integration / Continuous Deployment: End-to-End Ownership	7
Microservice/API/Event Based Architecture	8
Why Monoliths with APIs Added On Do Not Work	9
Omnichannel Consumer and Corporate Applications	9
Solution: Create Your Own Cloud Native Architecture	10
Proven Methodology to Rapidly Build, Test and Deploy	11
What is in the Blueprint?	12
Architectural Strategy of the Blueprint	13
Blueprint Offerings	15
Conclusion	19
About the Author	19



Executive Summary

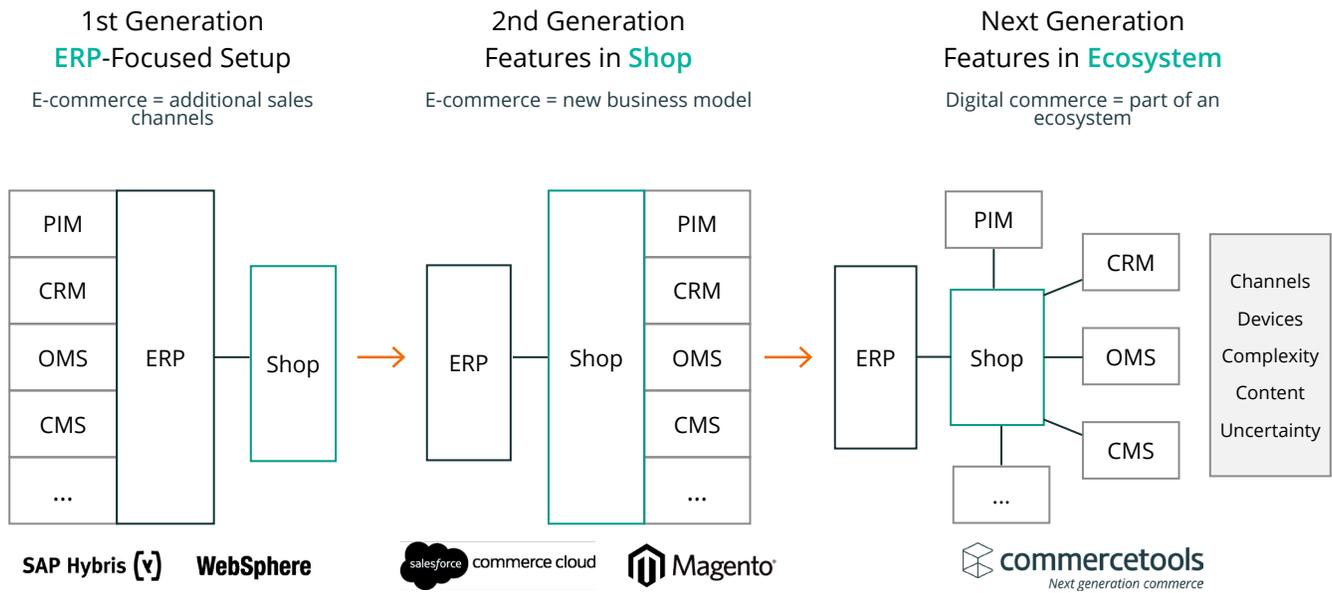
Customers now expect commerce to come to them across many different touchpoints. Combine that with increased pressures on operational costs while requiring constant innovation are forcing retailers to rapidly modernize their commerce architecture. Not since ecommerce first became an imperative during the 1990s have enterprises required such a monumental shift in the way they engage, deliver content across channels, sell and fulfill. Customer-facing sales channels have grown from the standard storefront and desktop webstore to contact points on mobile, television and streaming media, social media and IoT devices.

Modern commerce requires the capabilities of cloud and the flexibility of microservices to rapidly anticipate and respond to consumer behavior.

The large monolithic applications that were the commerce platform solution of choice in the 1990s have become costly, burdensome and ineffective at rapidly adapting to consumer needs. Modern commerce requires the capabilities of cloud and the flexibility of microservices to rapidly anticipate and respond to consumer behavior. *Only cloud-native, high-speed API, microservice and event-based architectures allow for rapid application development with unified real-time data as required by consumer centric applications.*

While enterprises realize the need, they often have limited budgets and limited knowledge of how to transform their architectures. commercetools is meeting this need with the only API-first, microservice-backed, cloud-native commerce platform on the market and a publicly available open-source blueprint architecture that allows enterprises to rapidly transform to a modern commerce architecture without starting from zero.





Evolution of commerce platforms beginning with onsite monolithic applications to the modern API-first, microservice-backed, cloud-native commerce platform architectures available today.

In this paper, we will cover:

- How monolithic applications impact the enterprise
- How conflicting objectives can be overcome with the blueprint architecture
- How to use the commercetools Blueprint to rapidly create your own microservices commerce architecture



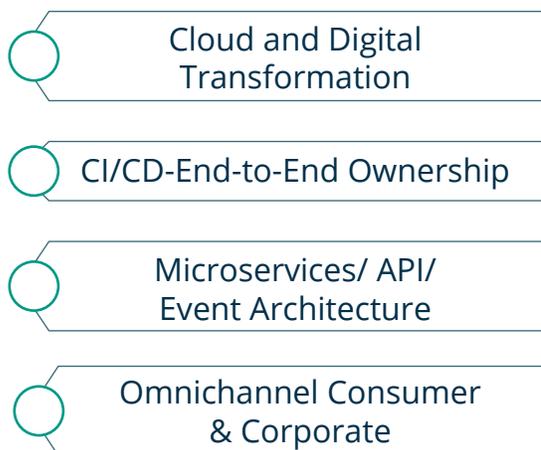
Business Issues

Once the organization has determined that an API-first, microservice-backed, cloud-native commerce architecture is the right solution, the inevitable next question is 'now what?'. The organization may have some existing applications with APIs built on top or perhaps small teams have even built a few one-off microservices exposed as APIs but there is not institutional knowledge around what or how to transform the organization.

The existing enterprise architecture is typically so complex and interdependent that the business suffers paralysis regarding where and how to start. Modern cloud-centric organizational transformation typically has four tracks (see figure 1) that are simultaneously vying for dollars and priority. All need to be in place to recognize full transformative value, but a big-bang approach is both reckless and costly.

Priority needs to be given to those capabilities that are the most consumer and cost focused. Thereby meeting both the revenue and cost reduction goals of the organization while presenting the highest ROI.

Transformation is typically centered around the following four capability tracks:



Four interdependent initiatives organizations face in modernizing their commerce architecture.



Cloud Transformation

There is more to 'going to the cloud' than simply copying your footprint into someone else's data center. Cloud transformation is a shift in technical operations, infrastructure costing and most importantly, the way applications actually work. The cloud, in turn, becomes the fabric on which the digital commerce platform is built, operated and grown. The simple 'lift and shift' of an application does little to gain the value of a cloud strategy.

When utilizing a cloud-based architecture, organizations are looking for the following capabilities from applications:

1. **DevOps:** The IT organization moves to a culture and practice of unifying software development (Dev) and software operations (Ops). The goal being shorter development cycles, more frequent deployments, and more dependable releases, in close alignment with business objectives. This is done with a focus on automation and monitoring throughout the development, testing, deployment and operations stages.
2. **Fully Elastic Consumption:** Scaling up and down, in an automated fashion, based on machine-based monitoring to maintain optimal performance at the lowest possible cost is typically the biggest benefit of a cloud transformation. The infrastructure scales up and down based on usage - even to the point of only existing when required. The need to estimate server/network usage a year in advance so hardware can be budgeted, ordered and deployed becomes a thing of the past. A virtually unlimited set of hardware is available the instant it is needed - and discarded the moment it is under-utilized.
3. **Geo-redundancy:** Duplicating data centers (DCs) and applications based on concerns of a statistically unlikely occurrence is an expensive form of insurance. Even when we optimize the cost spent by running live out of both data centers and sizing so that normal traffic could be entirely run from one DC in the event of disaster, it means our insurance policy only covers us when it matters least. It is not peak, Black Friday or Cyber Monday. When 70% of sales happen in two months, true geo-redundancy in the event of a disaster needs to be in place. In an active/passive model, cloud based geo-redundancy allows your off-DC to be turned off with near-zero cost while your live-DCs are operating and scaling as designed. Cloud native models typically run active/active as the need for a passive environment is typically redundant or unnecessary due to auto-scaling capabilities.



4. **Rapid, low-cost development:** Spinning up new or modified applications requires minimal investment as the underlying infrastructure is on-demand. Additionally, public cloud offered SAAS and PAAS solutions allow for focus on coding applications or configuring for a business need instead of focusing on deployment, scaling, hardware and data storage. If the application is not successful, it can be quickly removed with minimal sunk-costs.

Continuous Integration / Continuous Deployment: End-to-End Ownership

Minimizing the time from ideation to launch is an ongoing quest in technology. With the advent of automation technologies around version control, testing and deployment, the speed of change can be nearly instantaneous with multiple deploys per day. The only limitation is how fast change can be coded. Fixes and updates go from version control to production in a matter of minutes – vs. the weeks, months or years we face with monolithic applications.

The way code is deployed also changes in the CI/CD world. Three examples are:

1. **Blue Green Deploys** - two identical environments where green hosts the present production version while blue hosts the upgraded environment. You run smoke tests or any other tests in the blue environment while you prime the operating system, cache, etc. When things are fully running and ready, you point the the load balancer / router to the blue environment, which now becomes production environment.
2. **A/B Testing** - allows the enterprise to test features in the application for things like usability, popularity, noticeability, etc. If a feature presents value, it is rolled out to the remainder of the platform.
3. **Canary Releases** - send out a version of the application to see how it will perform, integrate with other applications, utilize memory, etc. The idea here being that no matter how much we test in lower environments, issues always arise in production. This way the deployment issues are isolated and can be rolled back before broad release.



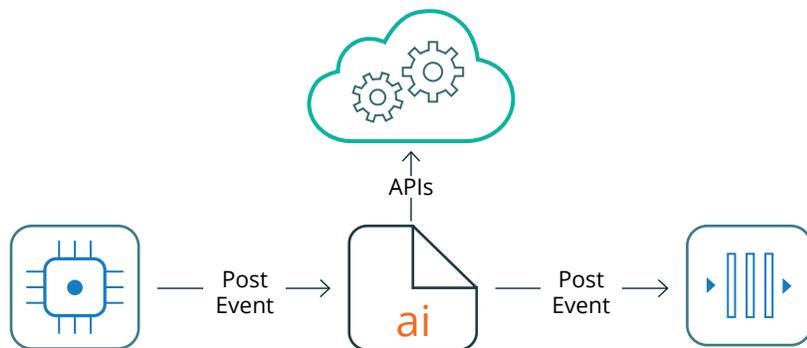
While the tooling around building development pipelines is part of the process, it is only a part of what is required. Using a scaled agile development approach to replace the waterfall methodology required by monolithic applications is essential as the process matches the organizational structure.

Instead of organizing teams around skill sets; such as DBAs, middle layer coders, front end developers, operations, etc. Development is organized around a functionality set and teams become vertical in nature. All skills required around a given functionality are available within the team. For more information on this topic, I recommend reading here on how Spotify has implemented a [scaled agile development process](#).

Microservice/API/Event Based Architecture

In an ideal world, all data and all processes would be available real-time to any valid application. This is the idea behind the API/Event economy. Simply put, the idea is that any data or processing is made available the instant it is required. An API allows you to request or post data to a process. Events, in turn, notify consuming entities that changes have occurred within the system.

An order is a simple example of how APIs and events work together. An order may be submitted via an API call while changes to the order state, such as received or shipped would be posted as events. Events become particularly interesting as IoT connected devices and artificial intelligence-driven actions increase in usage.



IoT event based AI mechanism. The AI mechanism reacts to the incoming event by posting its own event and/or exposing as an API.

Why Monoliths with APIs Added On Do Not Work

In the days of the monolith, all data was stored in one massive database with a middle tier and view tier built on top. This pattern was duplicated for different business purposes over and over within the enterprise. Any application that was going to rely on the monolith's data would either have to be a part of the monolith or would need to store a duplicate version of the data. Batch jobs and ETL (Extraction, Transformation and Loading) became the wiring that held together corporate architectures.

As needs changed, software vendors added layers on top of the monolith to expose APIs for use by external applications. This was meant to modify ETL processes to be less batch in nature but was never intended to fully support the high-speed scale required for omnichannel commerce. Because the API was an afterthought of the monolith, it was never meant to be a truly high-speed entity. Typically, any large load on an API causes the internal processing of the monolith to slow or fail due to the unexpected processing burden. This meant duplicating data in caching and content delivery layers in an attempt to cover up the weaknesses of the system.

The modern commerce platform has highly available, highly scalable APIs exposed to multiple consumer end points.

Omnichannel Consumer and Corporate Applications

In the original retail model, there were two primary touchpoints: storefront and catalog. Ecommerce came along in the 1990s and the catalog was largely replaced with the desktop webstore. Thus, organizations typically had one set of operations for stores and one for ecommerce. The software that supported this dual front was logically monolithic in nature. Ecommerce was treated as just another store in many instances and it was fed the data it needed.

Customer touchpoints are everywhere now. While companies are trying to catch up by building mobile apps, innovative companies like Disney® are unifying customer experiences online and in parks with their Magic Band while Amazon uses 'Dash Buttons' to allow customers to order with the push of a



button. Touchpoints like these are created by building small apps on top of the API economy. The objective for every retailer should be to allow consumers to make a purchase everywhere they interact with the brand.

The monolithic nature of the old ecommerce platforms just can not be modified to handle dozens of simultaneous applications (both internal and external) all vying for the same instantaneous real-time updated data. They were built in an era of web only purchasing, not the omnichannel experiences of the modern shopper.

Solution: Create Your Own Cloud Native Architecture

Innovative organizations want to combine the building of a sustainable, modern architecture with the highest impact customer facing capabilities. The commercetools Blueprint Architecture for Microservices Commerce gives you a foundation on which you can build and grow your digital platform.

The Blueprint is designed for focuses on a cloud-native, microservices based architecture with API and event-based transaction capabilities. The architecture is useful for interacting with both commercetools and other API based providers as well as connecting the internal ERP and other legacy systems.

At its core, the Blueprint is a cloud native foundation on which to build the omnichannel commerce capability that enterprises seek. The Blueprint is meant to be deployed into the public cloud. The enterprise is the owner of the deployment, all the code and all the networking. It is a fully stand-alone set of infrastructure and tools separate from commercetools and other vendors.

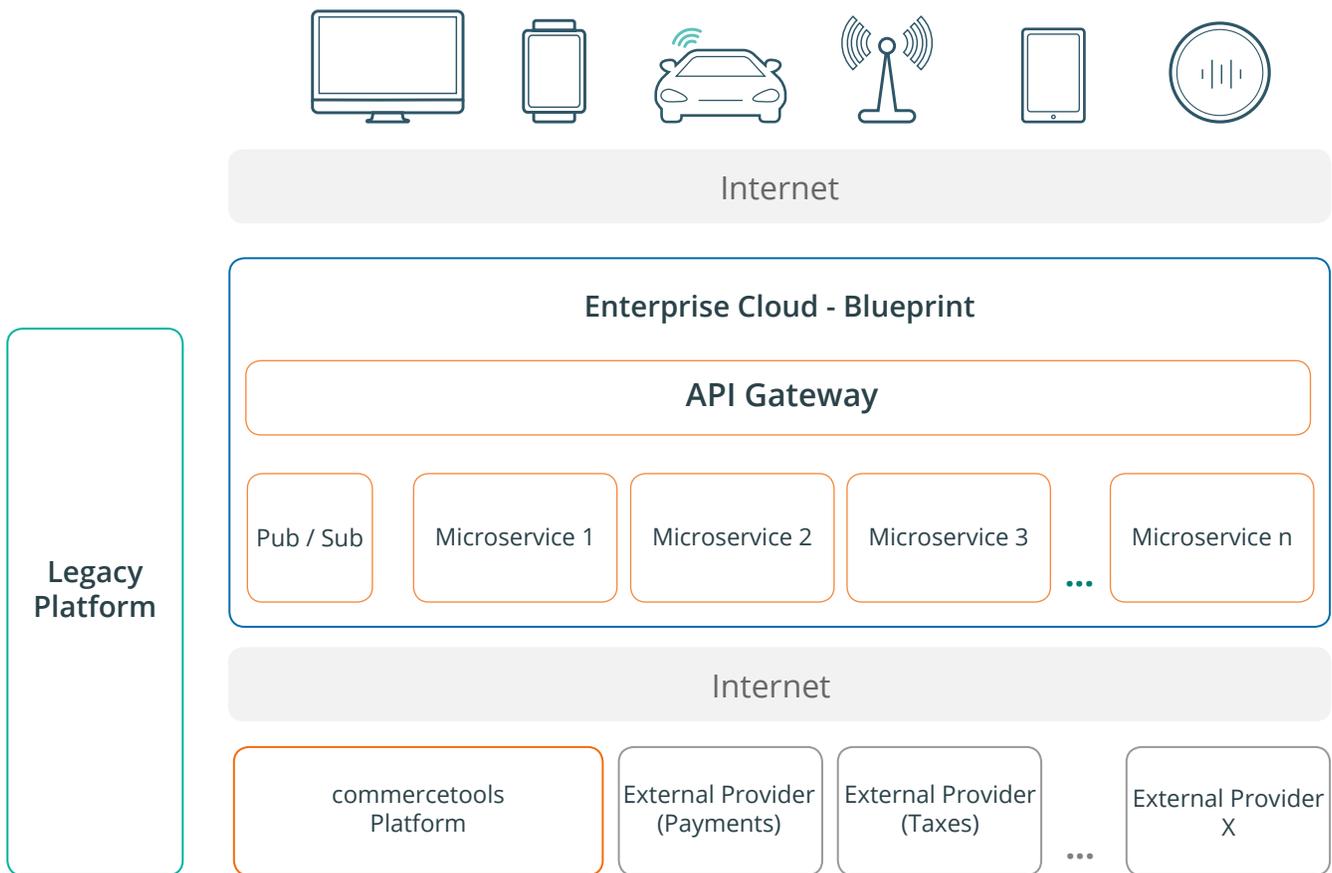
The Blueprint creates a level of abstraction between the corporate environment and providers. The enterprise then can define their own API contracts with multiple consuming applications. If the supporting component behind the API changes, consuming applications are unaffected and will not require code updates or testing. This means that providers can be replaced based on the value of their APIs vs competitors in the market. By introducing the blueprint abstraction, the enterprise does not tie their consuming applications directly to any licensed provider. The enterprise retains control of their public interface.



Proven Methodology to Rapidly Build, Test and Deploy

The commercetools Blueprint Architecture for Microservices Commerce allows companies implementing the commercetools platform to rapidly build, test and deploy a cloud-native, microservice based architecture just as commercetools uses, in a rapid, reproducible fashion into their own cloud environment.

The Blueprint creates a functional abstraction API, event and service layer with full CI/CD and auto scaling capabilities that is fully controlled by the enterprise. The commercetools APIs and events are abstracted away from direct consuming applications behind a corporate infrastructure.



Vendor software is obfuscated behind enterprise owned microservices and APIs. Full control of the API interface rests with the enterprise. All consuming applications engage with enterprise APIs, thereby allowing downstream vendors to be replaced or upgraded without any consuming application impact.



Unlike typical blueprints or open source samples, the commercetools Blueprints are full architectures that can be workable in a very short time period. If this were a lawn, the commercetools Blueprint would be the sod – not just seed. While there is effort to deploy and connect, there is a full set of workable services, development pipelines and connections ready to go.

What is in the Blueprint?

There are multiple blueprints available. Each has a mono-repository and instructions on how to deploy a fully functional blueprint. In addition, individual pieces of the Blueprint can be utilized if the full architecture is not required. Examples of this would be the CI/CD pipelines, sample microservices (both serverless and standing service in nature).

The Blueprint addresses the transformative requirements enterprises seek in the following ways:

Cloud Transformation: The Blueprint is a cloud-native, auto scaling architecture that is API and Event based. While the core Blueprint microservices integration is with commercetools, the microservices can be cloned or passed through to connect easily to other API vendors. Ultimately, the architecture provides the connection functionality to commercetools along with the pattern of integration for other API based vendors. There are two architectures available:

- **Blueprint GCP PaaS** – Organizations that want to push as much of the operational activities as possible while still owning the code can utilize a blueprint built entirely on the Google Cloud PaaS offerings. This means that all control of server level functionality and autoscaling gets pushed to become Google’s responsibility. The enterprise need only provide the coding, testing and deploy of functionality. Google handles the rest.
- **Blueprint K8** – A cloud independent Blueprint with fully containerized code built within a Kubernetes deployment. This means that AWS, GCP or Azure can all be used. Each component is fully independent and multiple clouds can be used simultaneously. Ultimately, this means more control rests in the hands of the enterprise.



CI/CD – End-to-End Ownership: A full continuous integration/continuous delivery pipeline including Github-based source repository, automated testing scripts and deployment built on Kubernetes.

Microservices/API/Event Architecture: The blueprint is a microservice-based architecture with exposure via APIs and Events. Connections to the commercetools set of APIs and events are built into each of the representative services.

Omnichannel Consumer & Corporate: While no consuming applications are built as part of the architecture, a full high speed corporate API/eventing layer is available to access via the include API gateway. This allows for rapid development of both corporate facing and consumer facing applications on mobile, IoT, desktop, etc.

Architectural Strategy of the Blueprint

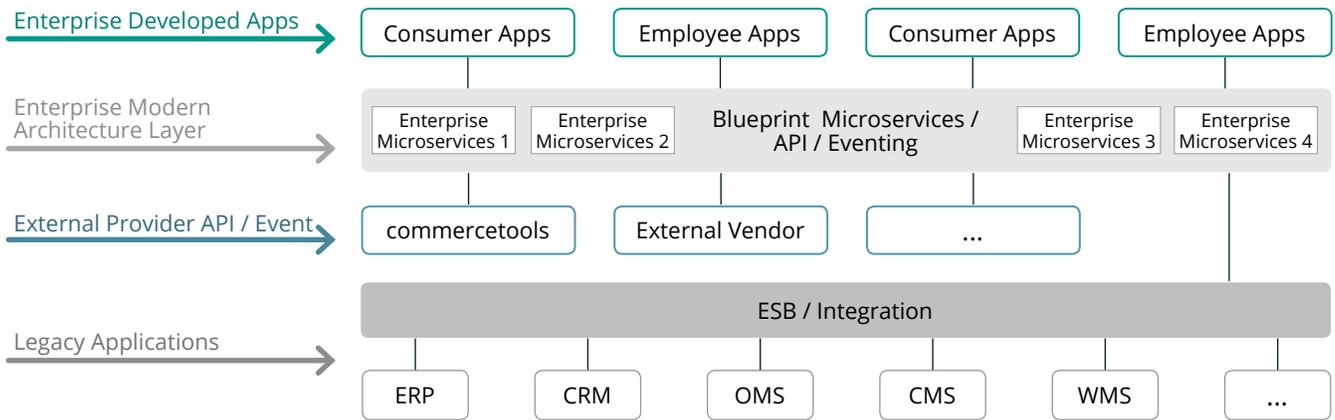
In an ideal world, the commercetools Blueprint Architecture would be an end-to-end enterprise architecture that could do everything any enterprise needed in a high-speed, cloud-native, microservice based API fashion. Replacing ERP, warehousing, and innumerable other large monoliths with a one time project is an unsound practice. Thus, the blueprint focuses on the consumer facing, commerce layer of the architecture and not on the operational side.

Inherent in deploying the Blueprint is that the underlying legacy architecture core is not modified. This strategy is often referred to as *Dual-Speed* or *Hybrid* architecture. The idea is that the consumer facing technologies and data are built on very high speed – scalable, microservices architecture while the core, infrastructure of the organization is not modified (at least in the initial phases).

This allows the organization to compete and lead in consumer facing applications while the core of the architecture can be modified over time in a reasonable and methodical manner. Other components of the enterprise can be removed or reduced through a strangle pattern while the outward face of the organization is cutting edge.

The
commercetools
Blueprint
focuses on
the consumer-
facing
commerce
layer.





Utilize the blueprint to create a high speed commerce face while removing any tight vendor coupling and legacy system dependency.

The blueprint layer is owned and managed by the enterprise. The endpoints it is exposing to consuming apps are then always controlled by the enterprise. If, for example, one of the commercetools microservices is replaced by a microservice built in-house, the consuming app exposed endpoint would not need to change.

Additionally the blueprint layer offers another place where customization can occur. As the enterprise defines its own API interface, the enterprise microservice layer defined by the blueprint architecture will contain enterprise-built custom microservices, transforms, customization on vendor APIs, GraphQL capabilities and multi-API consolidation.

As with any architectural structure, the enterprise should determine if the blueprint architecture has more benefits versus directly integrating with API vendors such as commercetools. Both have an element of dual-speed in the architecture, but the blueprint centralizes control of the enterprise API set into something wholly owned and controlled by the enterprise. See the comparison chart below.



Benefits of Enterprise-Owned Service Layer Model	Benefits of Directly Connecting Consuming Applications to API Vendors
Full control of the app facing layer	No additional hop - transactionally, this is going to be in the 10 to 50 millisecond range
Consumers applications do not change when functionality vendors are replaced	No additional monitoring and issue diagnosis layer
Additional point of customization	No additional cloud environment costs
Optimize APIs for consuming applications	No additional development cost
Organizational change starting point	
Rapid development environment	
Automation built-in	

Blueprint Offerings

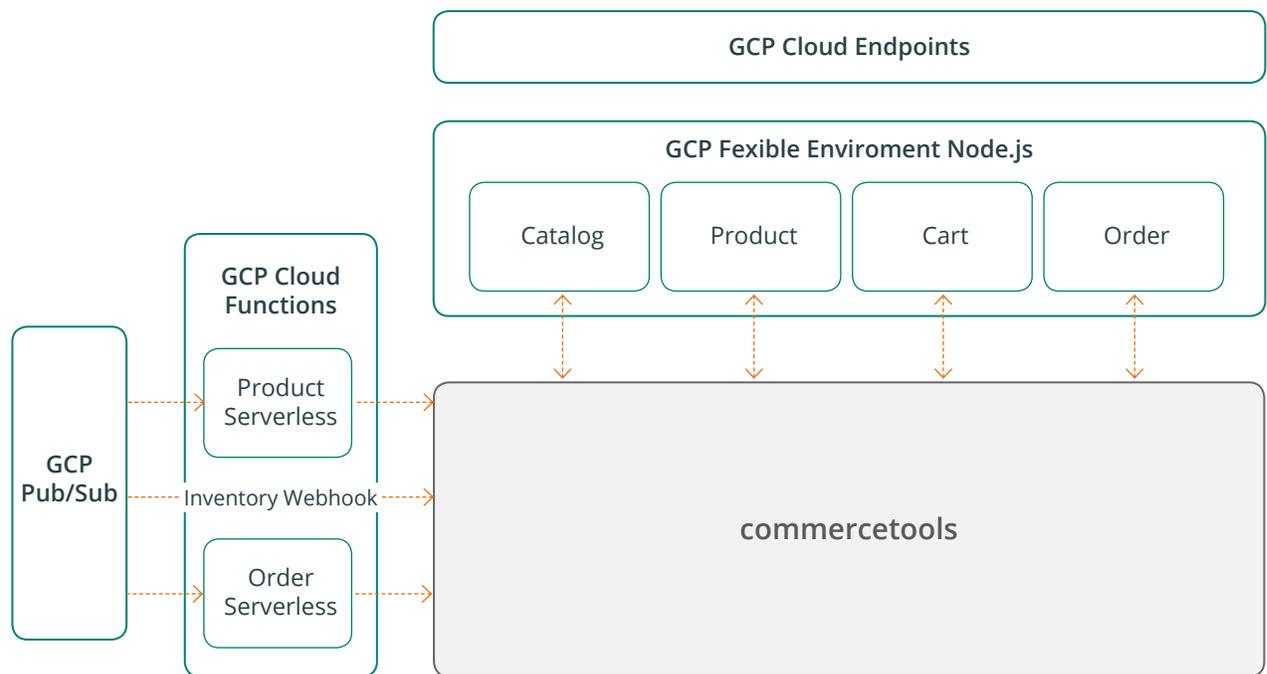
Two blueprint architectures are available. Blueprint GCP PaaS is meant for organizations who want to move most of the operations side of their cloud offerings to the cloud provider. Blueprint K8s has significantly more flexibility but requires the organization to take more control over what, where and how they deploy their architecture.

Blueprint GCP PaaS

Blueprint GCP PaaS is dependent on the use of Google Cloud Platform. It is an entirely Node.js based set of code. It utilizes the GCP's set of SaaS and PaaS offerings to let Google handle things like scaling and machine sizing. There is almost no required knowledge of hardware required in this blueprint. It is entirely a coding exercise and is meant to minimize the operational knowledge and resource requirements for deployment.



Components of Blueprint GCP PaaS



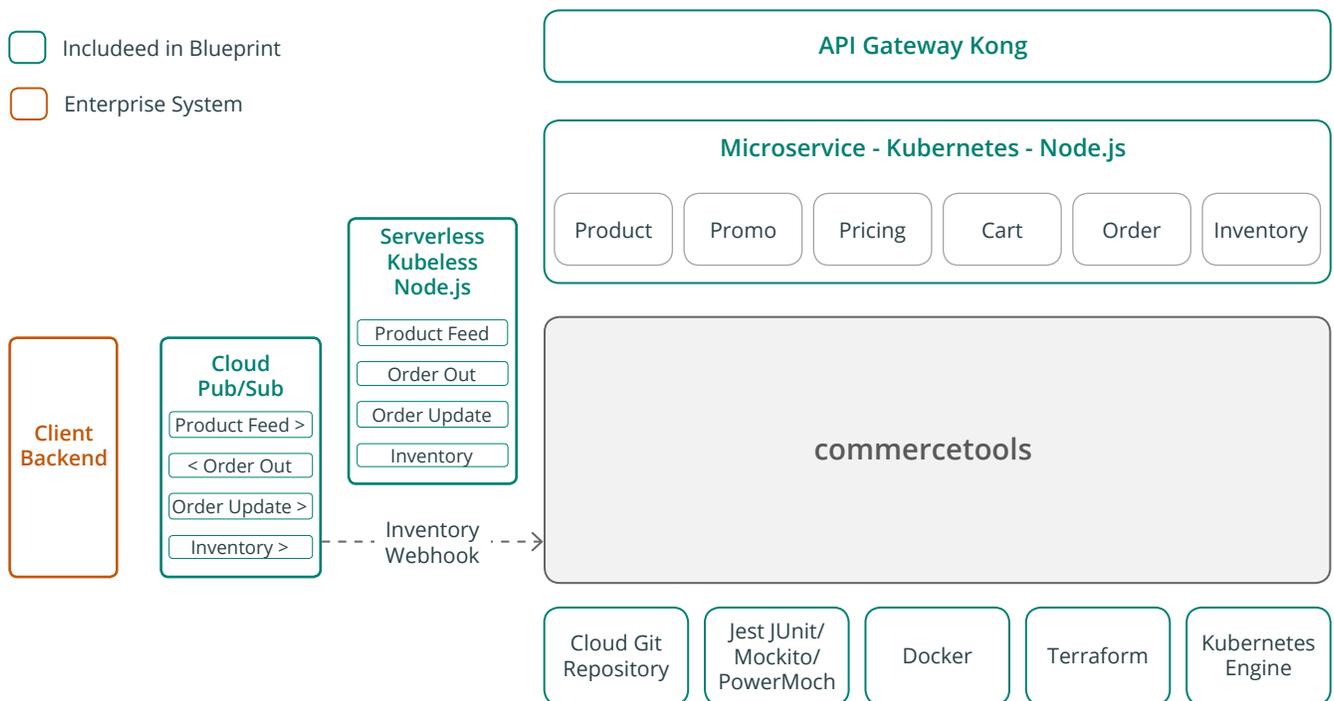
- **API Gateway:** Cloud Endpoints - Google Cloud Platforms API gateway offering.
- **Serverless Microservice:** Google Cloud Function - Node.js serverless function within the Google ecosystem.
- **Microservice:** Google Flexible Environment - Node.js implementation of 'always-on' microservices.
- **Cloud Pub/Sub:** Google Cloud Pub/Sub - Eventing environment
- **Monitoring:** Stackdriver
- **CI/CD Pipeline:** Pictured below, includes automated testing with Jest and deploy.



Blueprint K8s

Blueprint K8s has no cloud provider based dependencies. All functionality is containerized via docker as well as deployed and scaled via Kubernetes. This does mean that the cloud to which it is deployed must support Kubernetes; of which AWS, GCP and Azure all support. Where Blueprint GCP PaaS relied on Google Cloud autoscaling and networking, Blueprint K8s gives full control back to the enterprise.

The blueprint provides a pattern and deployment for the enterprise to move to a large scale cloud architecture in a matter of days.

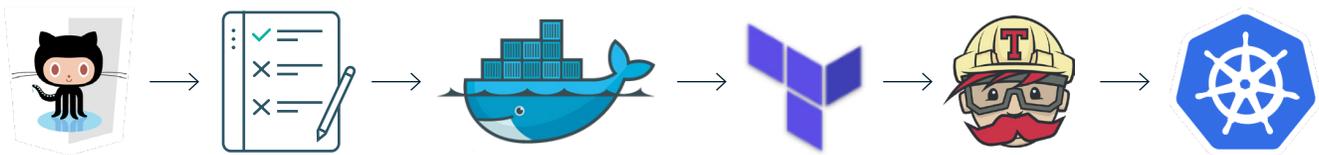


Components of Blueprint K8s Implementation

- **API Gateway:** Kong (<https://konghq.com/>)
- **Serverless Microservice:** Kubernetes native serverless offering called Kubeless. Written in Node.js for the included code, but does not have a language restriction like GCP or any of the cloud provider serverless offering out today.



- **Microservice:** Kubernetes based deploy of a Node.js and Java based 'always-on' microservices.
- **Cloud Pub/Sub:** Kubeless based offering written in Node.js
- **Logging and Monitoring** are left for the implementation. Every enterprise has tools in place and there is no need to force a second set as part of the implementation. Each cloud will have logging by default as part of the deploy which can be integrated or replaced.
- **CI/CD Pipeline:** Pictured below, includes automated testing in Jest or JUnit (Mockito/PowerMock libraries), Terraform and Kubernetes deploy.



Operational Components

Because Blueprint K8s can be deployed into any of the Kubernetes supporting clouds (AWS, GCP or Azure), the cloud dependent operational and management functions are left to the implementation.

Getting Started

The goal behind the Blueprints is to accelerate and define many of the early steps in an implementation. Organizations typically know the set of functionality they need to build in a project. The blueprint gives the development team a template from which to start. It defines everything from directory structure to deployment pipelines as well as containing coded solutions and template microservices to assist the team in jumping over the 'getting started hurdles'. The goal is to save a typical project six to eight weeks and allow the team to quickly deliver functionality.

Utilizing either of the blueprints as an accelerator for your platform can be done in a number of ways. The most obvious use is as part of a digital transformation project as you move to headless commerce built around commercetools. Enterprises who wish to slowly modify their platform because of business or technical reasoning would likely utilize the architecture to build out a greenfield site or replace portions of their digital experience piece by piece over time. If you have already started the journey, it is likely that you would only use components such as the CI/CD pipeline and a few of the sample services to accelerate a commercetools implementation.

The goal of the Blueprint is to accelerate an implementation by 6-8 weeks.

The Blueprint is available from commercetools and includes full descriptions of components as well as instructions on implementing. Much of the work on the part of the enterprise will be setting up the cloud environment and many of the administrative tasks such as billing, security, developer access and general cloud setups. At that point, the blueprint is simple to deploy and the enterprise is ready to expose its modern cloud-native API centric platform.

Conclusion

Enterprises recognize the need to upgrade their architecture. The monolithic structures relied on as the digital market grew have become burdens to progress. Many enterprises understand where they want to be but not how to get there.

Utilizing the commercetools Blueprint Architecture for Modern Commerce delivers immediate impact to the organization. By focusing on consumer centric portions of the organization, the enterprise becomes very nimble and is given the buffer room required to update the operational back end systems in a reasonable manner.

To learn more about the Blueprint or to sign up for a demo, go to: <https://ok.commercetools.com/blueprint-architecture>

About the Author

William Linebarger is a solution architect at commercetools, the world's leading API-based commerce platform. He has more than twenty years of ecommerce development experience ranging from developer to senior architect. William came to commercetools from Belk, where he led their transition to cloud as Director of Enterprise Architecture. Before Belk, he spent more than eight years leading the development of eBay Enterprise's first microservices-based commerce platform. He holds a bachelor's degree from the University of Rochester, an M.B.A. from the University of Kansas, and is a Google Professional Cloud Architect.





commercetools is a next-generation software technology company that offers a true cloud commerce platform, providing the building blocks for the new digital commerce age. Our leading-edge API approach helps retailers create brand value by empowering commerce teams to design unique and engaging digital commerce experiences everywhere – today and in the future. Our agile, componentized architecture improves profitability by significantly reducing development time and resources required to migrate to modern commerce technology and meet new customer demands. It is the perfect starting point for customized microservices.

Contact

Americas
American Tobacco Campus | Reed Building
318 Blackwell St. Suite 240
Durham, NC 27701, USA
Tel. +1 212-220-3809
mail@commercetools.com

Europe - HQ
Adams-Lehmann-Str. 44
80797 Munich, Germany
Tel. +49 (89) 99 82 996-0
info@commercetools.com

Join the Conversation

 @commercetools
 commercetools-gmbh
 commercetools