

DEVELOPER SURVIVAL GUIDE

EXPRESS SERVERLESS PLATFORM FOR IBM vs.
IBM CLOUD FUNCTIONS + IBM API CONNECT



LunchBadger

Express Serverless Platform vs IBM Cloud Platform (IBM Cloud Functions + IBM API Connect)

ABSTRACT

We've created this comparison page to make it easy to understand the major differences (and similarities) between two popular platforms for the microservices and API development use case.

In this review we'll be comparing Express Serverless Platform for IBM and IBM Cloud's offerings.

It's important to note to achieve a comparable experience to Express Serverless Platform, one must utilize multiple offerings in IBM Cloud including:

- IBM Cloud Functions
- IBM API Connect
- IBM Kubernetes Service

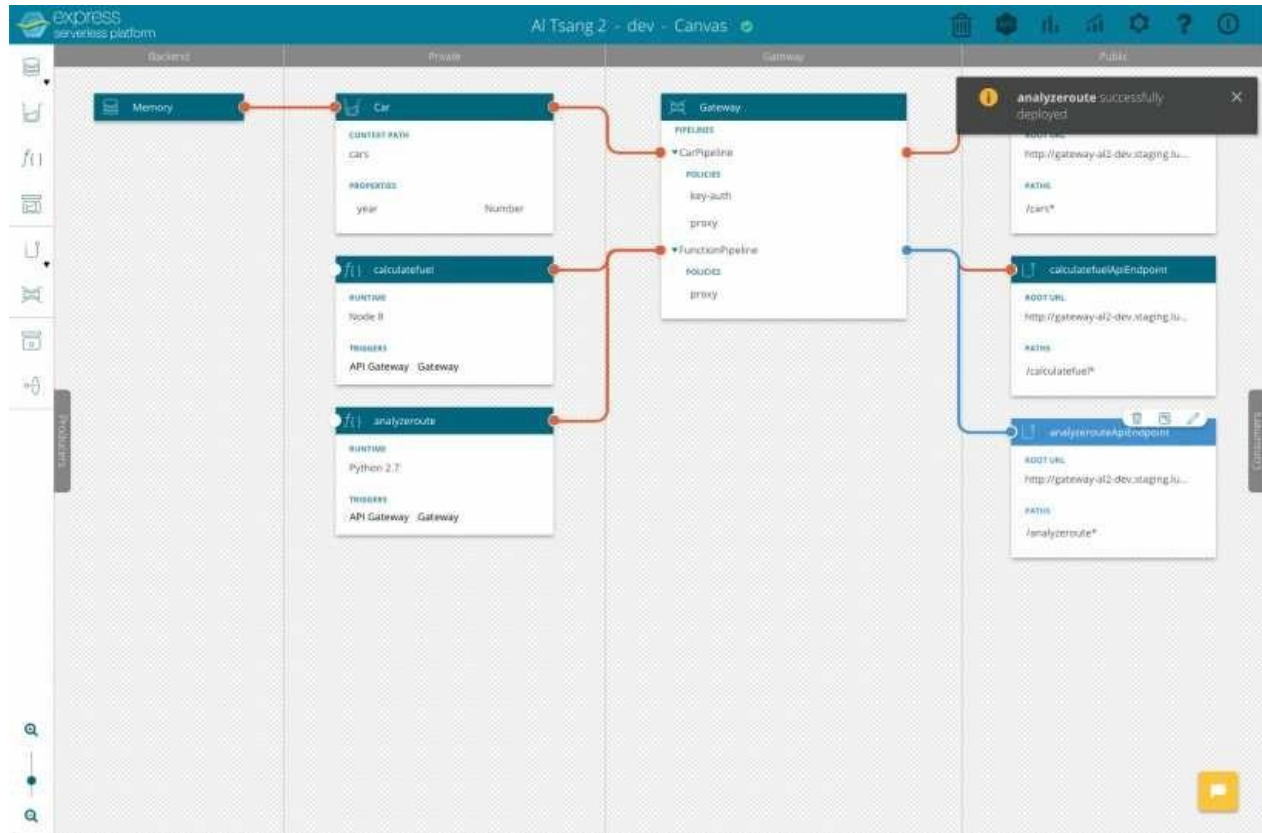
What is Express Serverless Platform?

Express Serverless Platform provides a uniform and seamless development experience for microservices and APIs as simple functions.

Express Serverless Platform can be deployed to any public or private cloud.

Enterprises can have multiple deployments of Express Serverless Platform at the same time to achieve a multi cloud and or hybrid cloud strategy.

Each deployed instance of Express Serverless Platform allows you to take advantage of a cloud's native infrastructure without coupling your code to its proprietary interfaces.



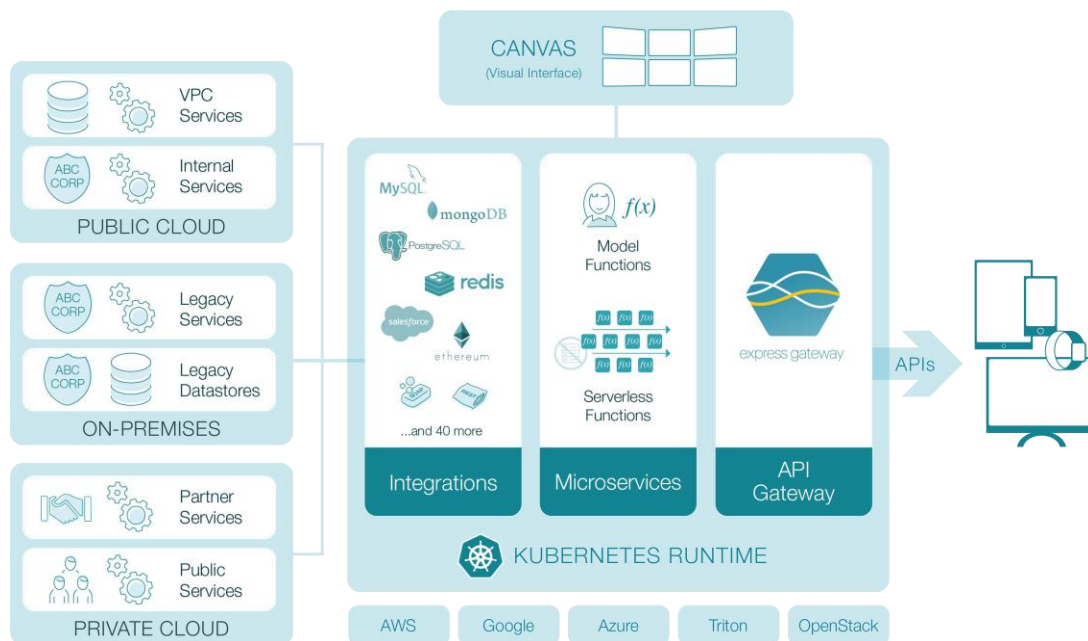
Express Serverless Platform’s GUI is called the Canvas. Actions done inside the Canvas are orchestrated and automated within Kubernetes as pods running container based microservices in real time for development and modeling purposes.

Other actions in the Canvas can also automate and deploy native cloud infrastructure such as serverless offerings (e.g. IBM Cloud Functions) running alongside the cloud’s Kubernetes offering.

Express Serverless Platform automates and manages both container and serverless microservices seamlessly giving you a unified view of what your microservices application looks like and how they’re orchestrated across different pieces of infrastructure.

Features & Architecture

Express Serverless Platform Features & Architecture



Express Serverless Platform is fully self contained. Express Serverless Platform can be installed on bare metal or VMs with all the necessary infrastructure components included: an API Gateway, Kubernetes Runtime, Serverless Engine and visual Canvas that serves as a "single pane of glass" to visualize your microservices and API development.

Express Serverless Platform is fully modularized and integration ready. Any deployment of Express Serverless Platform can take full advantage of a particular public cloud's native proprietary offerings such as its own serverless infrastructure.

The Express Serverless Platform utilizes the best of breed open source technologies in its reference implementation that is ready to be used out of the box as is:

- [Express Gateway](#), an open source API Gateway written in Node.js and built on top of [Express.js](#)
- [Loopback.js](#), an enterprise Node.js framework for building microservices as model based functions and vast microservice integration library of connectors
- [Kubeless](#), a Kubernetes native serverless engine to run polyglot functions
- [Serverless](#) (the framework), a framework that provides abstraction of functions to multiple serverless implementations - AWS Lambda, Azure Functions, Google Cloud Functions, Openwhisk, Kubeless and others
- [Kubernetes](#), the leading container orchestrator to run microservices running in containers, supported by all public and private clouds

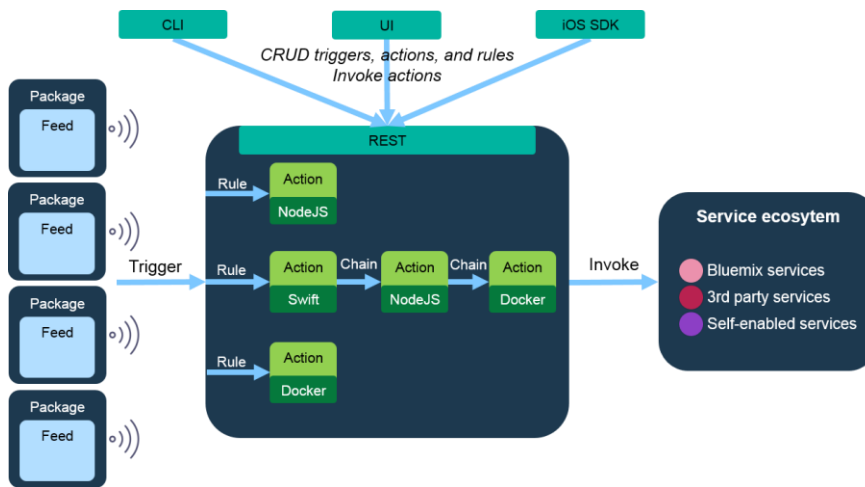
Express Serverless Platform Features & Architecture Cont.

Features Include

- ◆ Composition of Serverless Functions
- ◆ Composition of Model based Functions
- ◆ Supports Node.js, Python, Ruby, Go, PHP, .Net Core, and Java
- ◆ Completely extensible through open source modules
- ◆ Suite of Enterprise connectors to tap into legacy systems for Model based Functions
- ◆ Connects multi-cloud environments (private or public cloud)
- ◆ Supports any microservice use cases, patterns and designs
- ◆ Auto deployment to a Kubernetes Runtime
- ◆ Works with any Kubernetes cluster (EKS, GKE, etc)
- ◆ Auto exposes functions as APIs through the gateway
- ◆ Plugs directly into existing DevOps tooling and pipelines
- ◆ Can take advantage of public cloud's proprietary infrastructure services
- ◆ Reference implementation comes complete with all infrastructure components like API gateway, serverless engine, container orchestrator, and GUI

Further reading: [LunchBadger Documentation](#)

IBM Cloud Functions Features & Architecture



Like many other public cloud offerings, IBM Cloud has separate offerings for serverless, API management, and Kubernetes capabilities. These offerings are developed as separate products and the integration among them can be quite cumbersome.

IBM Cloud Functions are IBM Cloud's equivalent to AWS Lambda. Like Lambda, IBM Cloud Functions let you write custom functions and expose them via an HTTP interface through either IBM Cloud Function's minimal HTTP configuration or IBM API Connect for more sophisticated use cases.

Express Serverless Platform offers a similar stack that lets you write Node.js APIs and expose those APIs through Express Gateway without vendor lock in. IBM Cloud is the new name for what used to be called Bluemix, so don't be surprised by the Bluemix.net URL. OpenWhisk is an Apache open source project that provides the backbone for IBM Cloud Functions.

Features Include:

- Composition and Deployment of Serverless Functions
- Create APIs from integration flows with databases and back-end systems
- Testing and integration of cloud services and capabilities from IBM, open source communities, and third-party developers
- Connect legacy systems and apps from a cloud platform through private network and API capabilities
- Real time resource management
- Publish APIs as private services
- Access to the IBM Cloud services available including IBM Watson, Internet of Things, Analytics, Mobile, and IBM DevOps services



Further reading: [IBM Cloud Function Features](#)



Features Comparison



The following table has IBM Cloud Functions and IBM API Connect (two separate products used together) compared with the Express Serverless Platform.

The API Gateway that come with IBM API Connect is a microservices gateway. IBM DataPower Gateway is a separate paid product.

- = Supported
- ◐ = Supported with caveats
- ◑ = Not supported

Feature		
General		
On Premise	◐ ¹	●
Runs on any Public or Private Cloud	◑	●
Kubernetes Support	n/a	●
Configuration & Administration	GUI, CLI, API	GUI, YAML, CLI and API
Auto Scaling	●	●
Visual Designer	◑	●
Visual Orchestration	◑	●
Git Access	◑	●

Serverless		
HTTP Functions	●	●
Event Functions	●	●
Model Functions	●	●
Node.js Functions	●	●
Python Functions	●	●
Swift Functions	●	●
Go Functions	●	●
.NET Core	●	●
Ruby Functions	●	●
PHP Functions	●	●
Java Functions	●	●
Docker Image	●	●
Auto REST Scaffolding	●	● ²
Events and Triggers	●	● ³
Pre-Built Connectors	IBM Cloud only	8/20+ ⁴

API Management		
HTTPS	●	●
CORS	●	●
Basic Auth	●	●
OAuth2	●	●
Key Authentication	●	●
JWT	● ⁴	●
LTPA Token	●	●
Finegrain Access Control	●	●
Rate Limiting	●	●
Quotas	●	●
Request Transformation	●	●
Response Transformation	●	●
Pipeline Driven Conditional Actions	●	●
Pipeline Driven Expressions	●	●
Consumer Management	●	●
API Portal	●	● ⁵
Plugins Framework	●	●
Open Source Ecosystem	●	●
Configuration Database	IBM DB2 and IBM Informix, MySQL, Cloudant and CouchDB, MongoDB	In-Memory Redis

1. IBM Cloud Functions is based on OpenWhisk, unclear on portability, IBM API Connect can be installed via VM
2. Auto REST scaffolding for Models only
3. HTTP triggers only at this time (general pub/sub in roadmap)
4. JWT is only available if you use the DataPower Gateway, another IBM paid offering
5. API REST Testing Interface can be integrated into an existing portal

>> Getting Started with the Express Serverless Platform

Getting started with Express Serverless Platform is dead simple.

LunchBadger offers a free 14-day trial of Express Serverless Platform with no credit card required.

The trial runs in LunchBadger's cloud so you don't need to have a public cloud account nor worry about installing anything on premises - simply sign up.

The trial provides a complete walkthrough building a microservices application and its API through a series of guided and interactive steps - all in less than five minutes.

Once built, you'll trace your API workflow from public API Endpoints to your in-memory data source and be able to experiment with a live and running microservices application to see the value of the time savings and functionality it provides.

> QuickStart

The following is an overview of the getting started experience provided by the self-guided walkthrough.

- Deploy and use an Memory Connector to connect to an in-memory database
- Create and deploy a "Car" that will be a Model based microservice
- Connect the Car Model to the Memory Connector to read and write Car data
- Deploy and configure an API Gateway – an instance of Express Gateway
- Connect the Car Model to the API Gateway
- Expose the Car Model microservice as a Car API Endpoint that we can call through an API Request using cURL
- Deploy a Function called MyFunction that will be a Serverless Function based microservice
- Connect the MyFunction Function to the API Gateway
- Expose the MyFunction microservice as a MyFunction API Endpoint that we can call through an API Request using cURL

> Running Express Serverless Platform

Express Serverless Platform can run in any Kubernetes cluster on any cloud - private or public. It's a true multicloud solution.

Express Serverless Platform for IBM is Express Serverless Platform installed and ran on IBM Kubernetes Service within IBM Cloud.

Out of the box, Express Serverless Platform runs serverless functions in the Kubeless serverless engine in your Kubernetes cluster via the Serverless Framework. By auto deploying serverless functions via the Serverless Framework, you can avoid vendor lock in and use any cloud provider or on premise solution.

Express Serverless Platform using the Serverless Framework will also allow you to run serverless functions in a public cloud provider's proprietary offering if you sign up for an enterprise level plan. In this comparison, the cloud provider's offering for serverless would be IBM Cloud Functions built on top of OpenWhisk.

>> Getting Started with IBM Cloud Functions

Getting Started with IBM Cloud Functions is straightforward. After you log in, you can create actions, set up Cloud Functions CLI plugins and then view the platform architecture. We'll quickly cover the basic quickstart you can use and templates.

QuickStart:

- Log in and begin by using the HelloWorld quickstart template.
- From the IBM Cloud Catalog External link icon Functions category, click Functions.
- Click Start Creating > Quickstart templates and select the Hello World template.
- Check out the code for the action, and then create the action by clicking Deploy. This will create an action called "hello"
- Click Invoke to run the action. Invoking an action manually executes the app logic that the action defines. In the Activations panel, you can see the "Hello stranger!" greeting produced by the action.
- Expose the MyFunction microservice as a MyFunction API Endpoint that we can call through an API Request using cURL

Further reading: [Developer Comparison of the Getting Started Experience](#)

> Running IBM Cloud Functions

[OpenWhisk](#) is an Apache open source project that provides the backbone for IBM Cloud Functions. IBM Cloud Functions is a Function-as-a-Service (FaaS) platform based on Open Whisk. So as you begin working with IBM Cloud Functions, keep in mind that it's a managed service on top of an open source project.

Supported environments include:

- Docker
- Swift
- Java
- Go
- Python
- PHP
- Node.js

The easiest way to get started with IBM Cloud Functions is to go to IBM's Cloud Functions landing page, log in and use one of the Quickstart Templates to begin building your Functions.

[IBM Cloud Functions Getting Started Documentation](#)

IBM Cloud Functions store its configuration in IBM DB2 and IBM Informix, MySQL, Cloudant and CouchDB, MongoDB. You also have the option to use a Docker container for an action.

Further reading: [IBM Cloud Functions System Details and Limits](#)



Pricing

Express Serverless Platform

Express Serverless Platform is offered as a free 14-day trial.

License plans come packaged with a number of users, microservice functions, and API gateways at a low cost that predictably scales as you realize the value of what the platform brings.

There are no hidden or extra compute fees for operating in IBM's Cloud Platform.

Operating costs are transparently passed onto you by running in your account.

IBM Cloud Functions

IBM Cloud Functions offer a free API Gateway Rate with no limits. However; Basic Cloud Functions Rate are billed at \$0.000017 per second of execution, per GB of memory allocated. Which is a mouthful. The first 500 mc, 128MB of Action Memory and 5M Executions are free every month.

Further Reading: [IBM Cloud Functions Pricing](#)

Quick Reference Links

Express Serverless Platform



- [Documentation](#)
- [Installation and Getting Started](#)
- [Free Trial](#)

IBM CloudFunctions

- [Documentation](#)
- [Installation and Getting Started](#)
- [Platform Architecture](#)

For a more in-depth comparison, you can check out the [just getting started experience comparison between IBM Cloud Functions and Express Serverless Platform.](#)