

Contents

Overview	1
Modules	1
Foundational AWS Acceleration: Organizational Layout	1
Foundational AWS Acceleration: CI/CD	1
Serverless Acceleration: Intro to Serverless	2
Serverless Acceleration: Project Layout.....	2
Serverless Acceleration: Testing and Debugging	2
Serverless Acceleration: Migrating Legacy Architectures	2
Serverless Acceleration: Security	2
Serverless Acceleration: Release Engineering	3
Serverless Acceleration: Monitoring and Observability.....	3
Serverless Acceleration: API Design.....	3
Serverless Acceleration: Serverless Data Stores	3
Serverless Acceleration: Gotchas and Pain Points	3
Serverless Acceleration: Architecture Patterns	3

Overview

Customers can mix and match from 2 Foundational AWS Acceleration modules and 11 Serverless Acceleration modules. Each module delivery contains:

- Approx. 1 hour of produced content (slide deck and potential supporting material such as code or demo)
- Approx. 1 hour of collaborative guidance, whiteboarding/applying best practices to customer situation

Modules

Foundational AWS Acceleration: Organizational Layout

If you're new to AWS, this "landing zone" experience is a great introductory module. We'll help you build out an AWS Organization, review core IAM and network security concepts, and set you up for future success in the cloud.

Foundational AWS Acceleration: CI/CD

Continuous integration and delivery are the backbone of a cloud native development shop. In the AWS world, that means getting savvy with CodePipeline, CodeBuild, and CodeCommit. We'll show you some awesome best practices, including our legendary dynamic feature branch pipelines, and

reinforce the importance of CloudFormation as the Infrastructure-as-Code tool underlying all the serverless goodness to come.

Serverless Acceleration: Intro to Serverless

We know, serverless still has servers. What it doesn't have is a clear definition everyone can agree on. In this overview session, we'll get your team up to speed on the development and operational benefits of a serverless-first approach, and the rationale for an architecture built on "cloud-native" managed services. We'll even mention the "c" word -- containers. As with all our sessions, there will be plenty of time for Q&A with our opinionated experts. (Just don't throw anything at us.)

Serverless Acceleration: Project Layout

Don't just lift and shift your app to create a "server in a function"; getting the most out of serverless requires a whole new approach to code layout and application architecture. In this session, Trek10 will help you make sense of modern tools like the Serverless Framework and AWS's Serverless Application Model, address the thorny monorepo vs multiple repo question, demonstrate Lambda code and dependency best practices (routing, logging, environment variables, what code goes in a function handler), and much more. You'll get all the code examples when we're done!

Serverless Acceleration: Testing and Debugging

How do you test an application made up of cloud services? There's no easy answer, but Trek10 will walk you through tried and true tools and best practices including local function execution, IDE plugins for serverless, and even cloud-native coding. Let us show you how to reduce latency and frustration in your development cycles while ensuring that what gets tested is actually what you want to deploy. You'll get all the code examples when we're done!

Serverless Acceleration: Migrating Legacy Architectures

How do you turn a monolithic, on-premise application into a flexible architecture that takes full advantage of the cloud? In this session, we'll cover some best practices for migrating data and converting legacy architectures to serverless, but the real value comes when we unpack your specific migration use cases and identify the cloud services and architecture patterns that will simplify your move to the cloud. Bring a whiteboard and prepare to have your assumptions challenged.

Serverless Acceleration: Security

With great serverless power comes great responsibility, and it's up to you to make sure your serverless apps follow security best practices. The Trek10 team will dive into least-privilege access design, help you right-size your Lambda functions, review best practices for dependency and secrets management, and introduce you to some remarkable third-party tools that can take your security posture to the next level. You'll get all the code examples when we're done!

Serverless Acceleration: Release Engineering

In this module, Trek10 will show you how to deploy and update serverless applications in production with zero downtime. We'll look at safe Lambda deployments and rollbacks with SAM and CodeDeploy, DNS-based failover with multiple CloudFormation stacks, versioning shared libraries with Lambda Layers, and managing shared application components with the Serverless App Repository. You'll get all the code examples when we're done!

Serverless Acceleration: Monitoring and Observability

There's no "SSH" in serverless, but you can still use best practices to keep tabs on how your applications are running in production. (Hint: we don't just mean squinting at CloudWatch Logs.) We'll look at monitoring and observability best practices involving AWS CloudWatch and X-Ray, introduce you to third-party tools such as IOPipe, Epsagon, and Datadog, and help figure out what to do when your tools aren't helping. You'll get all the code examples when we're done!

Serverless Acceleration: API Design

AWS offers several intriguing ways to build a fast, cost-effective API without any servers at all. But the "classic" API Gateway -> Lambda -> DynamoDB design is only the beginning. We'll walk you through advanced REST API designs (did you know you can put a Lambda function behind an application load balancer?), show you the power of AWS AppSync for managed GraphQL back ends, and even dig into AWS's new websocket capabilities for CloudFront and API Gateway.

Serverless Acceleration: Serverless Data Stores

Data is the beating heart of your serverless application, and we're here to make sure you get it right the first time. DynamoDB, the "granddaddy" of fully managed data stores, will get a workout as we demonstrate patterns and antipatterns for serverless architectures. If you prefer SQL in your life, we'll dive deep on the performance and security considerations of Aurora Serverless and its new data API. We'll also pull some surprising power from the depths of good old Amazon S3.

Serverless Acceleration: Gotchas and Pain Points

We love serverless as much as anyone, but we'll be the first to agree there are plenty of challenges when it comes time to apply it in the real world: cold starts, stateful connection issues at scale, microservices sprawl, cost considerations, time and space constraints, and probably two more things you just ran into. We'll lay out plenty of mitigations and workarounds to keep you moving forward -- and more importantly, we'll let you know when to panic. (Don't panic. You can do this!)

Serverless Acceleration: Architecture Patterns

Trek10's experts have built leading-edge serverless architectures for some of the world's largest cloud-native companies, giving us unparalleled insight into serverless architecture patterns that work and perform at scale. In this advanced session, we'll dive into function fan-out and fan-in, workflows and long-running tasks, asynchronous and synchronous APIs, real-time streaming, cross-account service discovery, and more.