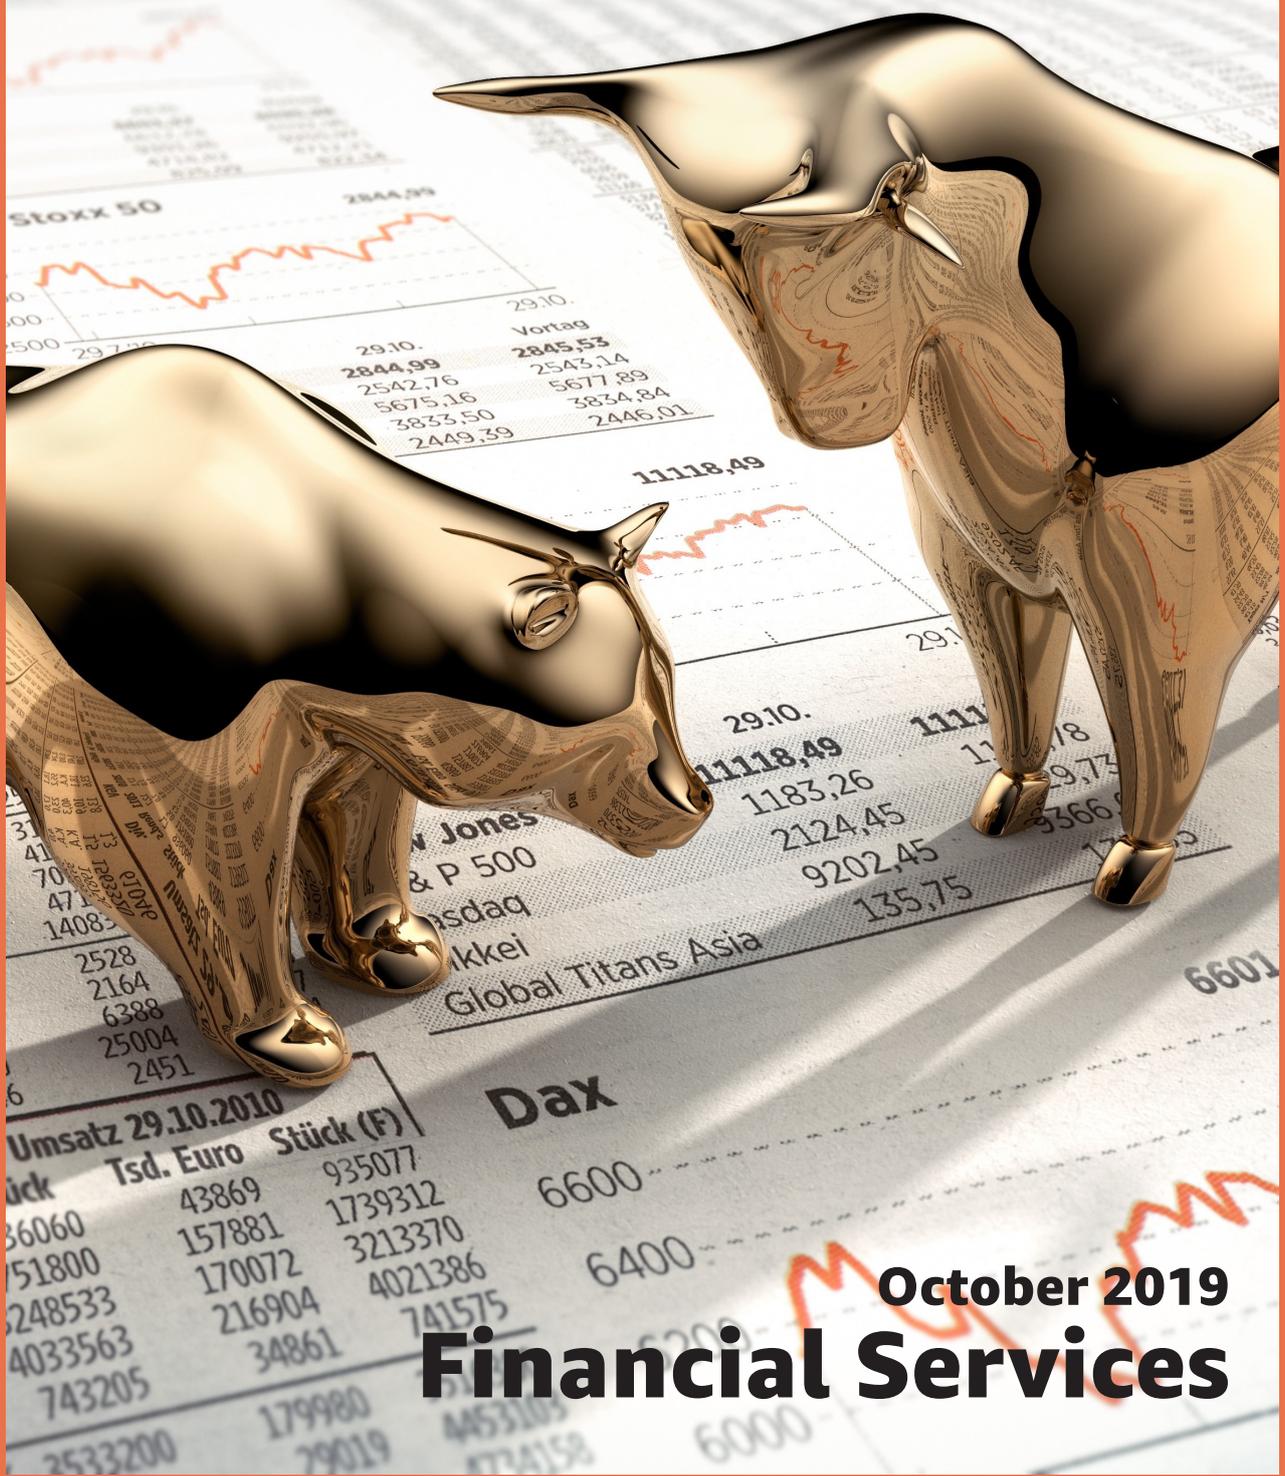


AWS Architecture Monthly



October 2019
Financial Services



Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided “as is” without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.

© 2019 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Editor's Note

This month's magazine delves into the high-stakes world of banking, insurance, and securities.

From capital markets and insurance, to global investment banks, payments, and emerging fintech startups, AWS helps customers innovate, modernize, and transform.

We're featuring two field experts in October's issue. First, we interviewed Ed Pozarycki, a Solutions Architect manager in the AWS Financial Services vertical, who spoke to us about patterns, trends, and the special challenges architects face when building systems for financial organizations. And this month we're rolling out a new feature: Ask an Expert, where we'll ask AWS professionals three questions about the current magazine's theme. In this issue, Lana Kalashnyk, Principal Blockchain Architect, told us three things to know about blockchain and cryptocurrencies.

We hope you'll find this edition of Architecture Monthly useful, and we'd like your feedback. Please give us a star rating and your comments on Amazon. You can also reach out to aws-architecture-monthly@amazon.com anytime. You can also view past issues at <https://aws.amazon.com/whitepapers/kindle/>.

For October's magazine, we've assembled architectural best practices about Financial Services from all over AWS, and we've made sure that a broad audience can appreciate it.

- **Interview:** Ed Pozarycki, Solutions Architecture Manager, Financial Services
- **Blog post:** Tips For Building a Cloud Security Operating Model in the Financial Services Industry
- **Case study:** Aon Securities, Inc.
- **Ask an Expert:** 3 Things to Know About Blockchain & Cryptocurrencies
- **On-demand webinar:** Simplify Machine Learning with AWS
- **Whitepaper:** Financial Services Grid Computing on AWS

Annik Stahl, Managing Editor

Interview:

Ed Pozarycki
Manager, Solutions Architecture



I recently sat down with Ed Pozarycki, a Solutions Architecture manager for AWS Financial Services, to talk about his experience and observations on architectural trends in the industry.

What are the general architecture pattern trends for Financial Services business?

Financial services customers are coming to AWS for the same reasons that our other millions of active customers are, and the first and foremost thing is agility. They want the ability to innovate, spin up resources quickly, try different ideas, and develop new services.

If a company had to purchase, their own equipment and such, it could be five months before they try an idea out. And by the time they try it out, the opportunity may be lost. Agility spawns innovation, and I'll give you an example of that.

At the most recent Financial Services symposium that we had in London, Jo Hannaford, Head of EMEA Technology & Global Head of Quality Assurance Engineering at Goldman Sachs, said the company wanted to launch a new business in the U.K. called Marcus, an online-only savings product. They were able to launch this new business, end-to-end, in 11 months, from development to production. Jo said that timeframe would have been unthinkable if they would have had to provision their own hardware and not use a cloud provider like AWS.

The second reason is scalability and elasticity. Jo went on to say that with Marcus, they weren't able to predict how the public was going to react to this brand new business or if they'd have enough capacity. She said that having the ability to auto scale plus with the global footprint and elasticity of AWS, the company didn't have to guess, and within the first hour of launching Marcus, they had 600 times more customers than they'd anticipated. Jo said that if it had been built with anything other than a cloud provider like AWS, they wouldn't have been able to scale.

The third and final point would be around resiliency, which is extremely important for Financial Services firms that have a regulatory requirement where they need to provide business continuity plans for their applications, where they can prevent, detect, respond, and recover from disruptions that could affect the world's economy. The global

infrastructure that we provide—the points of presence, the geographic regions, the availability zones and the ability to seamlessly architect fail over of an application that takes advantage of that—is extremely important to Financial Services firms.

When putting together an AWS architecture to solve business problems specifically for Financial Services customers, do you have to think it all differently?

100% yes. And I believe that's why AWS created the Financial Services vertical three years ago. There was a customer need to have a team of folks that understood the compliance and regulatory obligations that Financial Services organizations are faced with and I alluded to that earlier, when talking about resiliency.

Consider [Depository Trust and Clearing Corporation](#) (DTCC), which clears 90% of the world's trades—it's basically the settlement house for every trade that goes on. If their systems or applications were to go down, that means trades can't be cleared, and there's a huge impact on Financial Services around the globe. So we have to think a little bit differently about our Financial Services customers.

The first way we go about it is the same way every other AWS Solutions Architect (SA) would: with the [AWS Well-Architected](#) framework, which is what we've learned from our million plus customers about best practices to help them build secure, high performing, resilient, efficient architectures. One of the pillars in that framework is around security. We dive very deep into the security with our customers and talk about best practices around confidentiality, integrity of data, who can do what around privilege management, how to protect systems from cyber threats, and finally controls that customers can put in place to understand if somebody violates their security posture.

The second thing is that we've created the "Financial Services Advisory Board." We're big at AWS about listening to the voice of our customer. As much as we like to think in our Financial Services vertical that we understand Financial Services very well, there are certainly some things that we don't know, so we put together an advisory board to further understand the specific requirements each company needs, and then we use that information and work very closely with the teams architecting these applications. There are cases where services don't necessarily meet the requirements of Financial Services Tier 1 applications. So what do we do? Well, at AWS, we work with our service teams. One of the jobs of an SA is to make services and features better and meet the needs of our customers. For example, in order to run an exchange on the cloud, there are very strict regulatory requirements around latency - so we would work with our customers very deeply to understand what those requirements are, and then help them run tests to see if we can meet their needs. If we can't, we would continue to iterate bring it to our service teams, who are more than happy to help us meet the needs of our customers.

The other thing that's a little bit different about Financial Services firms is that they're required to [whitelist](#) services. For example, if XYZ Financial Services firm wants to use [Amazon Elastic Kubernetes Service \(EKS\)](#). They're not allowed to use EKS until they whitelist it, which means going through an internal certification process. So they dive deep into our managed services and they check to see things like, *Is this service network isolated? Is it physically isolated? How does it handle data at rest and data in transit in terms with encryption? Can I use my own keys on this?* So there are a series of things that they look at before we can even architect an application using that service. And what we've learned over time is that the process is similar across our large Financial Services firms, so we've actually written a bunch of papers called "Service Adoption Acceleration Documents" that answers those questions and can help a customer in terms of what they need for their cyber team to get a service whitelisted.

What are your observations about machine learning (ML) and artificial intelligence (AI) in the fraud prevention space?

I'd like to talk about a couple of customers that are using our AI/ML products specifically for fraud detection.

I'll start with [NuData Security](#), a company that was purchased by MasterCard. For every transaction that happens, MasterCard has to confirm that the person on the other end is authorized to do so, and NuData's business is essentially reducing and eliminating credit card fraud. AWS makes the approach that they've taken possible because of their need to store a lot of data. Basically, they're moving past the traditional *Hey, what's your password? Let me ask you some security questions. What's your date of birth?* etc. And they're moving to *Let's analyze the digital profile of this person on the other end*, and they do that in real time—things that can't be replicated by another party. Data they've collected are stored in an [Amazon S3 data lake](#), which holds petabytes of data, and they're using services like [Amazon Kinesis](#), [Amazon Redshift](#), [Amazon EMR](#), and [Amazon Athena](#), to do analysis in real time to detect fraudulent behavior.

There's also the [Financial Industry Regulatory Authority \(FINRA\)](#), an organization that was authorized by Congress to protect America's investors, making sure that the broker-dealer interactions are on the level and operating fairly and honestly. In the past, FINRA would get millions of documents with unstructured data, and had to manually go through it to determine *Hey, is something fraudulent happening here? Is this trade in compliance?* and so on. But now they're using [Amazon Comprehend](#). So, instead of humans having to go through these documents, they're pulling out and extracting what they need in terms of anomalies from the documents, and as a result, they've exponentially increased the scope of the investigations that they're able to handle and they're also shortening the amount of time to conclusion of fraud or not.

Do you see different trends in Financial Services in cloud versus on-premise?

Why are capital markets coming to AWS? What are the trends that are making them look at cloud? A lot of it goes back to what I just said about AI and ML: there are increased regulatory reporting requirements that these firms have to adhere to. They have to store a lot more data, they have to go back further in time to prove to regulatory authorities that what they're doing is on the up-and-up, that the market risk they're holding is reasonable, and they're not doing anything outside of regulatory bounds. So in order to do that, there's a whole lot more data that they need to store. They also have to do risk calculations and increased analytics on that data more frequently and go back a longer period of time to report back to these agencies.

If you break it down into its simplest form, you've got to store a lot more data and you've got to do a lot more calculations on that data more frequently, which means you need more compute. So do you purchase additional compute in your data center, and run these workloads (that get very bursty) four to eight hours a day, and then shut them down leaving the servers that they've purchased essentially idle the rest of the day? Or is a better use case to use something like AWS where you can store the data, spin up the instances that you need only when you need them, run the calculations, do the reporting, and shut them down? AWS gives firms the ability to run these bursty workloads, meet their regulatory requirements, and save some money.

Another reason is innovation. Financial services have to differentiate themselves from one another, and a good example of this would be Bloomberg, a global leader in business financials that provides real time and historic price data—like tick feeds—to Wall Street firms. Typically, to get access to these data, a customer needs to install Bloomberg circuits on their premises. In their data center they need to install servers and switches that could take months to get up and running, and of course there's the ongoing maintenance. So, Bloomberg just came out with a new innovation by using AWS called [B-PIPE](#), which essentially provides their services instantly to clients via secure access through AWS PrivateLink. So by using our global infrastructure (Infrastructure in more than 110 countries delivers 80 billion picks per day to more than 15,000 customer locations), they're able to get to more customers in more countries, sell this service to more folks, and make it a heck of a lot easier and quicker to consume.

To sum it up, they are able to get to more customers in more geographic locations, provide a better service, make it more resilient, and make the up time better. It's all about innovation and they're using on AWS to be able to do that.

Is there anything you'd like to add about AWS in the Financial Services space?

When Financial Services firms are looking at AWS, it's really for all the same reasons: they're looking to gain a competitive advantage, move faster, and the agility that we provide gives

them that. They're looking to have a global reach. But the thing that makes them a little bit different is this hyper focus, sensitivity on security, risk, compliance, and resiliency, which is why we have the Financial Services vertical. We have experts in this area and we're continuing to listen to the voice of our customers. We're also going back to our service teams to ensure that what we provide the services and features that meet the needs of Financial Services firms.

Blog:

Tips for Building a Cloud Security Model in the Financial Services Industry

Available online at: <https://amzn.to/AWS-finsrv-blog>

By Stephen Quigg

My team helps Financial Services customers understand how AWS services operate so that you can incorporate AWS into your existing processes and security operations centers (SOCs). As soon as you create your first AWS account for your organization, you're live in the cloud. So, from day one, you should be equipped with certain information: you should understand some basics about how our products and services work, you should know how to spot when something bad could happen, and you should understand how to recover from that situation. Below is some of the advice I frequently offer to Financial Services customers who are just getting started.

How to think about cloud security

Security is security—the principles don't change. Many of the on-premises security processes that you have now can extend directly to an AWS deployment. For example, your processes for vulnerability management, security monitoring, and security logging can all be transitioned over.

That said, AWS is more than just infrastructure. I sometimes talk to customers who are only thinking about the security of their AWS Virtual Private Clouds (VPCs), and about the Amazon Elastic Compute Cloud (EC2) instances running in those VPCs. And that's good; its traditional network security that remains quite standard. But I also ask my customers questions that focus on other services they may be using. For example:

- How are you thinking about who has Database Administrator (DBA) rights for Amazon Aurora Serverless? Aurora Serverless is a managed database service that lets AWS do the heavy lifting for many DBA tasks.
- Do you understand how to configure (and monitor the configuration of) your Amazon Athena service? Athena lets you query large amounts of information that you've stored in Amazon Simple Storage Service (S3).
- How will you secure and monitor your AWS Lambda deployments? Lambda is a serverless platform that has no infrastructure for you to manage.

Understanding AWS security services

As a customer, it's important to understand the information that's available to you about the state of your cloud infrastructure. Typically, AWS delivers much of that information via the Amazon CloudWatch service. So, I encourage my customers to get comfortable with CloudWatch, alongside our AWS security services. The key services that any security team needs to understand include:

- Amazon GuardDuty, which is a threat detection system for the cloud.
- AWS CloudTrail, which is the log of AWS API services.
- VPC Flow Logs, which enables you to capture information about the IP traffic going to and from network interfaces in your VPC.
- AWS Config, which records all the configuration changes that your teams have made to AWS resources, allowing you to assess those changes.
- AWS Security Hub, which offers a "single pane of glass" that helps you assess AWS resources and collect information from across your security services. It gives you a unified view of resources per Region, so that you can more easily manage your security and compliance workflow.

These tools make it much quicker for you to get up to speed on your cloud security status and establish a position of safety.

Read the full blog at: <https://amzn.to/AWS-finsrv-blog>

Real-World Example

Discover and Slalom: Automate and Deliver AWS Multi-Account Creation with Scaffolding

Discover needed to standardize and automate AWS account creation in order to quickly deliver consistent AWS accounts for their application teams. They worked with Slalom to create a system that delivers a basic standard account scaffold for every account, including identity management, logging, shared services, and connectivity back to on-premises. The system uses Yeoman, Terraform, and Jenkins along with AWS services such as AWS CloudTrail, AWS Identity and Access Management, AWS Lambda, and Amazon Virtual Private Cloud (Amazon VPC).

<https://amzn.to/AWS-finsrv-TMA-slalom>

Case Study:

Aon Securities, Inc.

Available online at: <https://amzn.to/AWS-finsrv-case-study>

Aon Securities Inc. (ASI) is a registered broker-dealer focused on securities and other financial products of interest to insurance companies, including catastrophe bonds, contingent capital, sidecars, collateralized reinsurance, industry loss warranties, and derivative products. ASI offers services such as underwriting and placement of new debt and equity issues, financial and strategic advisory services, and a leading secondary trading desk. In addition, in conjunction with its affiliates, ASI provides distinctive analytics, modeling, rating agency advisory, and other consultative services.

The Challenge

Investing and risk go hand in hand. Many insurance retirement products contain financial guarantees, and these can span 30 years or more before paying off—so understanding whether an investment is a good risk over the coming decades is critical. A client company might run 5,000 different scenarios using a monthly time step for any given investment policy. Multiply that by five million policies, and the potential economic scenarios that might play out over the next 30-50 years, and the need to update risk analysis multiple times a year, and you have a massive computational challenge on your hands. It's work that requires high-performance, highly scalable computing and hardware that costs millions of dollars.

For a Financial Services provider like ASI, it is crucial to provide clients with effective business risk management solutions. ASI helps insurance companies price their investment products, analyze risk, and address regulatory requirements by using a financial modeling tool called PathWise, that simulates millions of potential economic scenarios and uses stochastic simulations to evaluate potential outcomes.

ASI customers typically maintain robust grids in their own data centers with as many as 2,000 CPU cores, but the company's intricate, complex growing financial modeling and reporting needs calls for a quantum leap in computing power: GPUs, or Graphical Processing Units. ASI's GPUs have tens of thousands of cores to process parallel workloads efficiently. Based on ASI's benchmarking, its clients stand to gain large increases in performance per dollar by switching to ASI's GPU-powered modeling tool. "Using GPUs is a massive leap forward for them—and that's without factoring in the huge spikes during quarterly reporting periods," says Peter Phillips, ASI managing director. "When you include that in the

analysis, using our business solution can be as much as 500 times more efficient in terms of performance per dollar for some clients.”

When ASI first launched PathWise, it used a broadband HPC processor in a collocated data center, but found that they needed a more scalable service than the collocation facility could provide. “Our challenge was being able to scale up as necessary without the need to build an expensive data center,” says Peter Phillips, president and CEO of the PathWise Solutions Group at ASI. “Specifically, we wanted GPUs — the rocket fuel of our industry. They’re also much more commoditized and less expensive than traditional data centers with CPUs.” The company needed a solution that would scale with it during peak periods, so it started to look at cloud computing.

Why Amazon Web Services

ASI found that by using Amazon Web Services (AWS) pay-as-you-go pricing, it could spin up large numbers of GPUs quickly and inexpensively, so it decided to move its infrastructure to AWS and deprecate its collocated data center. “We realized that by using AWS, we could have a whole turnkey environment up and running in no time,” says Phillips. “We didn’t have to spend time or effort building out a new infrastructure, either.”

ASI uses proprietary algorithms to generate economic simulations that span decades. The company uses the algorithms with PathWise Modeling Studio to run hedging simulations. Phillips says, “It’s called a Monte Carlo simulation, and our clients rerun those millions of times with different variables, all in parallel.”

ASI built a front end on AWS for its processing solution, using Pathwise Modeling Studio to compile a program using its algorithms and automatically running GPU instances on Amazon Elastic Compute Cloud (Amazon EC2) in an Amazon Virtual Private Cloud (Amazon VPC) for security. ASI uses Amazon Elastic Block Store (Amazon EBS) for persistent storage.

By processing on AWS, recalculating policies takes minutes rather than hours or days. “In regular quarterly financial reporting, it would take you two weeks and a small army of people to complete your regulatory reporting process,” Phillips says. “AWS gives us the computing power to shorten that time to hours and minutes.”

Being able to run these calculations at will means that ASI’s customers can more effectively judge the risk in their hedging programs for these insurance products, Phillips says. “When it comes time for our customers to make a trade, they’re going to take the current market information and re-calculate everything at once without any short cuts, and look at the intraday risk information from our system to help them make trading decisions as markets move. With our system, these calculations take minutes, not days or hours. Everything is synchronized with the latest market information, providing customers with situational

awareness as market conditions change, which is something legacy solutions cannot provide our clients with today.”

“Using AWS has enabled us to scale our work over hundreds of GPUs very efficiently, and deliver much more granular risk assessments to customers,” Phillips says.

The Benefits

By using AWS, ASI is able to deliver client solutions more quickly, with richer risk assessments and at a price that enables it to pass on savings to customers. “Being able to use GPUs so seamlessly means we can recalculate as often as we need to. We can run all 5 million policies in minutes, instead of the standard overnight run times,” Phillips says. “We can get a very accurate and unique picture of our customers’ market risk exposure—and there is no other solution that offers better performance at a lower cost for this business.”

Decreasing the amount of time it takes to complete calculations doesn’t just speed things up for ASI’s customers—it helps transform the way ASI thinks about those calculations. “Using AWS helps us reduce a 10-day process to 10 minutes. That’s transformative: it broadens our ability to discover,” Phillips says. “We are free to ask a lot more questions of our data now.”

“The beauty is that by using AWS, actuaries don’t need to know how to code. That’s not a productive use of their time,” Phillips says. “They also don’t have to schedule workloads or stand up an infrastructure for something they’ll only use a few times a year. Using AWS is a very efficient way to manage our business and help our clients manage their risk.”

Learn more about how AWS can help you with high-performance computing:

<https://aws.amazon.com/hpc/>.

Real-World Example

Vanguard: A Modern Network TAP Solution for the Cloud

Larry from Vanguard explains how his team built a lightweight, cloud-optimized, and scalable network TAP solution on AWS. You'll learn how they leverage CloudFront with Lambda@Edge to extract metadata from incoming traffic, how they designed an agent-based solution for capturing unencrypted data using the loopback interface on EC2, how their accounts are connected with PrivateLink, and how Kinesis and DynamoDB allow them to scale the solution on demand.

<https://amzn.to/AWS-finsrv-Tma-vanguard>

Ask an Expert:

3 Things to Know About Blockchain and Cryptocurrencies

By Lana Kalashnyk

Watch Lana's *This Is My Architecture* video, **Kaleido: Blockchain Cloud at AWS Speed:**

<https://amzn.to/AWS-FInServ-TMA-Lana>

1. Financial Institutions leverage cryptographic verification, which is at the heart of blockchain, to cut costs in regulated industries by providing a compute-inexpensive, fast way of verifying the integrity of data and applications (smart contracts) on audits, database restores, and systems of record reconciliation for mission critical workloads.
2. Pick the right tool for the job. Let the trust requirements for your use case drive the technology selection. Whether you have a central trusted authority administering the consortium, in which case a cryptographically verifiable ledger, like [Amazon QLDB](#) will excel as the immutable datastore, or if you want to maintain autonomy and sovereignty of the ledger data along with agreed upon immutable smart contracts, blockchain frameworks like [Hyperledger Fabric](#), [Ethereum](#) supported by [Amazon Managed Blockchain](#), or partner solutions in DLT (distributed ledger technologies) would be a great choice.
3. When deciding to build a digital asset or a cryptocurrency business, security is job zero. Make sure to leverage principles such as least privileges for access to sensitive information. Analyze the attack vector on your platform and leverage services such as [Amazon GuardDuty](#) and [AWS Shield](#) for perimeter protection, along with services like [AWS PrivateLink](#), [AWS Key Management Service](#) (KMS) and [AWS CloudHSM](#) for secure encryption key management and data transmission.



Lana is a Principal Blockchain Architect and a Global Partner Technology Lead at AWS Partner network. While at AWS, Lana led the technical incubation of the Blockchain Partners ranging from Intel, ConsenSys, R3, IBM, to Digital Asset, Accenture, Deloitte, Luxoft, and beyond. Prior to AWS Lana has over 12 years of lead engineering positions launching innovative solutions for burgeoning startups and Fortune 500 companies in multiple industries.

On-Demand Webinar:

Simplifying Machine Learning with AWS

Available online at: <https://amzn.to/AWS-FinServ-webinar-2>

Details

Length: 30 minutes

Speakers: John Kain, AWS, Capital Markets Segment Lead and Russell Lewis, AWS, FSI Compliance Specialist

Abstract

Despite rapid improvements in tools and technologies, the process of developing, training, and maintaining machine learning models can be cumbersome—particularly those working in regulated industries like Financial Services. In this webinar, we will share how Amazon SageMaker can accelerate machine learning innovation while helping to address model governance requirements specific to Financial Services.



Whitepaper:

Financial Services Grid Computing on AWS

Available online at: <https://amzn.to/AWS-finsrv-whitepaper>

Abstract

Financial services organizations rely on high performance computing (HPC) infrastructure grids to calculate risk, value portfolios, and provide reports to their internal control functions and external regulators. The scale, cost, and complexity of this infrastructure is an increasing challenge. Amazon Web Services (AWS) provides a number of services that enable these customers to surpass their current capabilities by delivering results quickly and at a lower cost than on-premises resources. The intended audience of this paper is grid computing managers, architects, and engineers within Financial Services organizations who want to improve their service. It describes the key AWS services to consider, some best practices, and includes relevant reference architecture diagrams.

Overview

High performance computing (HPC) in the Financial Services industry is an ongoing challenge because of the pressures from ever-increasing demand across retail, commercial, and investment groups, combined with growing cost and capital constraints. The approaches to solving these problems have evolved over generations from centralized, monolithic solutions, to business-aligned clusters of commodity hardware, to modern grid architectures with centralized schedulers that manage disparate compute capacity. Regulators and large financial institutions are increasingly accepting hyperscale cloud providers, which has resulted in significant interest in how to best leverage new capabilities while ensuring good governance and cost controls. Cloud concepts such as capacity on demand and pay as you go pricing models offer new opportunities to teams who run HPC platforms. Historically, the challenge has been to manage a fixed set of on-premises resources, while maximizing utilization and minimize queuing. In a model with capacity that is effectively unconstrained, the focus shifts away from managing and throttling demand towards optimizing supply. With this model, decisions become more granular and tailored to each customer, and focus on how fast and at what cost, with the ability to make adjustments as required by the business. With this basically limitless capacity, concepts such as queuing and prioritization become irrelevant as clients are able to submit calculation requests and have them serviced

immediately. This also results in upstream consumers increasingly expecting and demanding near instantaneous execution of their workloads at any scale. Initial cloud migrations of HPC platforms are often seen as extensions or evolutions of on-premises grid implementations. However, forward-looking institutions see much in common with the patterns of HPC and serverless execution models, such as AWS Lambda. Both solutions focus on executing code on demand, and customers want the lowest cost allocation of capacity with no provisioning or management of servers. As HPC environments move to the cloud, the applications that are associated with them start to migrate too. Risk management systems which drive compute grids quickly become a bottleneck when the downstream HPC platform is unconstrained. By migrating applications with the compute grid, they also benefit from the elasticity that the cloud provides. In turn, data sources such as market and static data are sourced natively from within the cloud, from the same providers that customers work with today.

Many of the building blocks required for fully serverless solutions for risk management and reporting already exist today within AWS services. As financial institutions become increasingly familiar and comfortable with these services, it's likely that serverless patterns will become the predominant HPC architectures of the future.

Read the full whitepaper here: <https://amzn.to/AWS-finsrv-whitepaper>

Real-World Example

Fidelity Investments: Using Containers at Scale with Amazon EKS

Amr from Fidelity explains how the company built a number of tools and services for Amazon EKS. Governance and security are top priorities at this diversified Financial Services firm that serves millions of customers daily. You'll learn how EKS Connect, EKS Manager, and other tools and processes allow it to operate efficiently and securely at scale.

<https://amzn.to/AWS-finsrv-TMA-fidelity>