

AWS Architecture Monthly



Biopharma
April 2021



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.

© 2021 Amazon Web Services, Inc. or its affiliates. All rights reserved.

Editor's note

From R&D through commercialization, biopharma organizations are looking to bring advanced and differentiated therapeutics to market faster. For example, the Moderna mRNA vaccine was developed and approved for emergency use in record time compared to past vaccines.

AWS is the secure, trusted technology partner to help the highly regulated biopharma industry confidently increase the pace of innovation, unlock the potential of scientific data, and develop a more streamlined approach to therapeutic development.

In this issue, we'll look at some of the architectural patterns, processes, and services available for scientific research, drug manufacturing, compliance, and scaling.

We hope you'll find this edition of Architecture Monthly useful. We'd like to thank Kelli Jonakin, PhD, Global Head of Life Sciences & Genomics Marketing, AWS, as well as our Expert, Patrick Combes, Tech Leader, Healthcare and Life Sciences for their contributions. Please give us your feedback! Include your comments on the [Amazon Kindle](#) page. You can [view past issues](#) and reach out to aws-architecture-monthly@amazon.com anytime with your questions and comments.

In the Biopharma issue:

- **Ask an Expert:** Patrick Combes, Tech Director, Healthcare and Life Sciences, AWS
- **Case Study:** Inspire Uses ML to Connect Millions of Patients and Caregivers on AWS
- **Quick Start:** AWS Biotech Blueprint - Core
- **Blog:** AWS and Novartis: Re-inventing pharma manufacturing
- **Executive Summary:** Why AWS for GxP regulated workloads?
- **Reference Architecture:** GxP Compliance Automation
- **Solution:** Service Workbench on AWS
- **Case Study:** Bristol Myers Squibb Builds Enterprise Data Lakes on AWS
- **Training:** Intro to Healthcare and Life Sciences (HCLS) Compliance on AWS
- **Reference Architecture:** Patient Engagement with Amazon Pinpoint and Amazon Personalize
- **Blog:** Executive Conversations: Accelerating COVID-19 vaccine development with Marcello Damiani, Chief Digital and Operational Excellence Officer at Moderna
- **Executive Brief:** Biopharma on AWS
- **Reference Architecture:** Lab instrument log acquisition and analytics
- **Related Videos:**
 - Streamline Innovation in Biopharma with AWS
 - Evidation Health Enables Next Generation Clinical Studies on AWS

- PwC: Building an Efficient, Intelligent, Serverless Pipeline for the Pharma Industry
- Bristol-Myers Squibb: Enabling DevOps with AWS Service Catalog
- **Podcast:** AWS Innovation with EB Research

Jane Scolieri, Managing Editor

Ask an Expert

Patrick Combes, Technical Leader for Health and Life Sciences at Amazon Web Services



What are the general architecture pattern trends for biopharma in the cloud?

Biopharma has long been on the cutting edge of cloud adoption—some of the earliest customers of AWS were in this space. Thirteen years ago, I was in biopharma myself. Personally, I used [Amazon Elastic Compute Cloud \(Amazon EC2\)](#) and [Amazon Simple Storage Service \(Amazon S3\)](#) to obtain resources that were difficult to get. Storage had to be forecast at least a year in advance—so if unplanned projects came up, we would scramble to buy desktop hard drives. Compute resources took at least 3 months to be ordered, delivered, and installed, which limited the ability of researchers to experiment with new ideas. Today, we're seeing cloud adoption across almost all efforts in biopharma. For example, development of drugs and therapeutics is data intensive—especially with the emerging use of artificial intelligence (AI)/machine learning (ML). While customers are now able to access all of the compute and storage resources they need, now they're looking for expanded access to their critical data. Patient data, trial data, data from earlier research, "real-world" data, and more are required for development, manufacturing and distribution, commercialization, and monitoring effectiveness and safety. As a result, biopharma has been adopting a more decentralized approach to data sharing and consumption. This includes data mesh architectures and methods for fully federated analytics and compute. These architectural patterns give biopharma organizations access to necessary data without having to exhaustively centralize it first. This helps them to more easily overcome internal and external organizational boundaries.

When putting together an AWS architecture to solve business problems specifically for biopharma customers, what are some of the considerations?

Helping our customers deliver safe and effective products while maintaining compliance with government regulations and industry standards is critical. They need to closely adhere to regulations and standards from the moment a researcher records their idea in a lab notebook to establishing the specifics of manufacturing and distribution. For obvious reasons, governments around the world regulate the safety and effectiveness of everything that can be prescribed to patients. This means our biopharma customers need to submit products for approval, retain complete records of the product's development, and enforce rigid controls to ensure consistency and safety. It's a balance to ensure that these

considerations do not become constraints. They should be implemented as a complementary, functional layer within every biopharma customer's architecture.

The next thing to consider is the diversity of workloads. There is no common architectural approach for the process of developing and delivering a drug or therapeutic in biopharma (also known as the “value chain”). The value chain spans fundamental scientific research to post-market surveys of therapeutics for safety and effectiveness with things like manufacturing along the way. The architectural requirements for each step are very different. In early research, tightly coupled high performance computing (HPC) is active for workloads like computational chemistry or molecular modeling. This has a different implementation from the more familiar analytics we support on [Amazon EMR](#) used during the commercialization of drugs and therapeutics.

Whatever architecture we help customers implement, we need to create something adaptable and extensible and not just another “silo.” The architecture should address their specific requirements but fit with all the other architectures being used to support workloads at other points of the value chain.

Do you see different trends in biopharma in the cloud versus on premises?

Not especially—as I noted earlier, biopharma has seen significant adoption of the cloud already. This experience has fed into overall design considerations. Now, most new on-premises deployments are capable of being hybridized or at the very least, extendable into the cloud. For the foreseeable future, most new workloads will continue to be deployed into the cloud. Those workloads with extremely low-latency requirements, like those seen in robotic laboratories or in manufacturing, will require on-premises infrastructure. We’ve already seen the trend for on-premises infrastructure to be cloud-connected, and I expect that to evolve into hybridized implementations like [AWS Outposts](#).

What’s your outlook for biopharma, and what role will the cloud play in future development efforts?

My outlook for biopharma is overwhelmingly positive—I've always felt it is one of the most scientifically exciting domains I could work in (the other being space 😊). Even with the current wide adoption of cloud, we’ve only just scratched the surface of what can be accomplished within biopharma. With emerging technologies and developments for AI and ML, the cloud is becoming more essential due to the requirements for extreme scale, flexibility, and data. These technologies are already beginning to be applied to difficult problems in biopharma like [computational chemistry](#), computational biology, [toxicology](#) and so many more. Beyond that, the cloud can accelerate every point along the value chain. It can deliver therapeutics more quickly and precisely and ultimately help close the loop between biopharma and patients to accelerate early research and development.

Anything else you'd like to add?

COVID-19 has been a devastating worldwide pandemic. But the rapid development, approval, and ongoing distribution of the vaccine is one of humanity's greatest scientific and industrial achievements. One of reasons it was so quickly developed is the cloud. Early in the pandemic, all observational and laboratory data regarding COVID-19 was exchanged and analyzed using the cloud. As the vaccines were being developed, all of the modeling, toxicology, and construction of clinical trials was performed using the cloud. Even prior to obtaining approval, the supply chain for the vaccine was established, the manufacturing capability was expanded, and the distribution network was built out using the cloud. Going forward, I know we will continue to apply what we have learned from this pandemic. And we will apply the core mission of our biopharma customers: to improve and help save the lives of all patients, everywhere.

About the expert



Patrick Combes is the Worldwide Tech Leader for Healthcare and Life Sciences at Amazon Web Services. He began his career in scientific computing at Fermi National Accelerator Laboratory (FNAL) where he developed embedded systems for gathering and processing instrumentation data in real-time. In 2005, he joined Rosetta Biosoftware (Merck) to build and integrate products for genomics and proteomics into HPC environments. He was part of the Amalga Life Science Platform within the Health Solutions Group (HSG) at Microsoft and then led the industry solution architecture team for EMC as the Principal SA for Life Sciences and HPC. Patrick has a B.S. in Computer & Electrical Engineering from the University of Illinois at Champaign-Urbana with additional graduate work focused on programming language design.

Case Study

Inspire Uses ML to Connect Millions of Patients and Caregivers on AWS

[Inspire](#), the vital online health community and important partner for life science companies, has a two-part mission. First, Inspire connects patients suffering from thousands of conditions and their caregivers with online tools, resources, and one another in condition-specific support groups. Second, Inspire connects pharmaceutical companies and other medical institutions conducting clinical-trials research—real-world evidence studies—to health-outcome studies on patients suffering from those diseases. “We seek to accelerate life-changing discoveries through our vital community of connected patients and caregivers,” says Richard Tsai, senior vice president of marketing at Inspire. More than 50 million people in 150 countries have used Inspire since 2015, and more than 2 million registered people with over 5 million reported medical conditions use the health community as of February 2021. Thousands more register every week, making it the largest and fastest-growing virtual support community for patients living with cancer, rare diseases, and chronic conditions, enabling them to actively share experiences and learn about diagnoses and treatments.

As Inspire looked to build on its success and keep up with its own growth, it needed to overcome the scaling challenges posed by its legacy on-premises infrastructure. Using managed solutions from Amazon Web Services (AWS), the company discovered numerous advantages in the cloud, including faster iteration, greater flexibility, and multiregion availability. Inspire found particular success using [Amazon SageMaker](#), a fully managed service that provides every developer and data scientist with the ability to build, train, and deploy machine learning (ML) models quickly. Using a solution driven by Amazon SageMaker, the company saw a substantial increase in user engagement across all channels. Inspire’s use of AWS also helped streamline the process by which pharmaceutical companies conducting clinical trials or medical research could connect with relevant patient data—an important step in the development of life-saving therapies.



We’ve migrated most of our development to AWS Lambda functions. Between that, caching, and Aurora, we honestly don’t pay attention to scaling anymore.”

Anthony Sheetz
Vice President of Engineering, Development
Infrastructure, and Data Science, Inspire

Scaling Automatically and Accelerating Innovation

Before using AWS, Inspire was operating around 20 physical servers in Ashburn, Virginia. The company ran into scaling problems, often having to wait 2 months to order and install a server and up to 3–6 months total to expand its capacity. In one instance, a database server upgrade resulted in a time-to-market delay of 3 months. Inspire explored the possibility of migrating to the cloud beginning in 2016, ultimately migrating its database to [Amazon Aurora](#), a MySQL- and PostgreSQL-compatible relational database that is built for the cloud and combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open-source databases. “The most compelling thing to us was the database offering, Aurora,” says Brian Loew, founder and CEO of Inspire. “That blew everything else out of the water.”

Also involved in the new infrastructure was [AWS Lambda](#), which lets customers run code without provisioning or managing servers. “We’ve migrated most of our development to AWS Lambda functions,” says Anthony Sheetz, vice president of engineering, development infrastructure, and data science at Inspire. “Between that, caching, and Aurora, we honestly don’t pay attention to scaling anymore.” Able to scale automatically and iterate more quickly on AWS, Inspire would eventually increase the frequency of its releases from one every 2 weeks to several a day, thus accelerating innovation and broadening its operations.

“We essentially forklifted our entire infrastructure onto AWS, keeping the same architecture,” adds Sheetz. “And then, once we were on AWS, we started playing with the toys.”

Using Machine Learning to Boost Engagement

A big part of Inspire is its content recommendation engine, through which it directs users living with particular conditions to relevant posts or articles. Integral to this engine is Amazon SageMaker, which Inspire uses in its development process to build and modify custom deep-learning models in 1–2-week cycles. “We’re now able to match users to relevant content by analyzing behavioral patterns and deploy these models with ease—all using Amazon SageMaker,” says Teja Talluri, director of data science at Inspire. “Amazon SageMaker provides a more scalable way of recommending content that we couldn’t manually curate ourselves.”

The sophisticated ML solution improved the content recommendation engine’s ability to suggest relevant content to two million registered users, pulling from Inspire’s massive library of 1.5 billion words written about 3,600 conditions. Ultimately, this solution enabled Inspire to accurately connect patients and caregivers with more personalized content and resources—including rare-disease information and treatment pathways.

When Inspire ran tests comparing new and old versions of its content recommendation engine, the metrics clearly showed more-robust engagement because of the company's ML-driven personalization efforts. Personalized email subject lines led to a 281 percent email open rate increase. And once users opened these emails, the new recommendation engine boosted the click-through rate by 914 percent, contributing to a 119 percent increase in average page views on the site. Inspire also saw that its retention rate, the number of users who remained active after 4 weeks, had increased by 550 percent since adopting the new content-recommendation engine.

Although the numbers are impressive, the human impact they represent is what matters most to Inspire. "We've received a lot of testimonials in which patients have said, 'The content you've recommended to me is so relevant,'" says John Novack, head of patient engagement and senior director of communications for Inspire. "We've never had that in the past. Now we have people telling us we've changed their lives—or even saved their lives."

Changing the Way Pharmaceutical Companies Find Critical Research Data

Inspire's other critical mission is to connect pharmaceutical companies researching new therapies with patients who may benefit from these therapies or at least provide useful data to those patients. Central to this use case are [Amazon Redshift](#), a fast, simple, cost-effective data warehousing service, and [Amazon Comprehend Medical](#), a natural language processing service that makes it simple to use ML to extract relevant medical information from unstructured text.

When [Boston Children's Hospital](#) and the pharmaceutical company [Pfizer](#) sought specific insights to aid in the development of [new lung cancer treatments](#), they faced a considerable challenge in finding data from a narrow set of patients: those with some combination of lung cancer and an autoimmune disease. Traditionally, researchers would have to reach out to individual investigators and clinicians to find patients who could supply relevant data—an extremely time-consuming process that could take years and still only find a few similar cases. However, Inspire's AWS-backed natural language processing capability enabled Inspire to search the profiles of tens of thousands of users who had consented to show up in such searches, and the company ultimately found more than 100 qualifying participants within a few weeks. Stefan McDonough, who is no longer at Pfizer but was the executive director of genetics at the time of the project, described the Inspire community as "an extraordinary resource," citing its rich pool of patients who are willing and eager to share medical information to advance treatments.

Connecting People with Profoundly Relevant and Impactful Information

After migrating to AWS, Inspire has seen a significant change in the way it does business. “AWS has given the whole software side of the house the ability to do what the business side needs done quickly and simply,” says Sheetz. “Now we can spend far more time focusing on giving our members new things.” Inspire expects that AWS data science–focused tools will help with the next phase of its business and play a key role in increasing revenue by an order of magnitude.

And at the center of Inspire’s mission are its users: the patients and caregivers who look to the company for help finding everything from practical information about rare illnesses to communities of people going through similar experiences. “It’s very important to be able to unite people around the world with similar health experiences,” says Sheetz. “This ability—to bring patients who have rare diseases together in a single place where they can share their experiences, regardless of where they live or what language they speak—has a profound impact.”

[Read Case Study online](#)

Quick Start

AWS Biotech Blueprint - Core



Core template for deploying a preclinical, cloud-based research infrastructure and optional informatics software on AWS

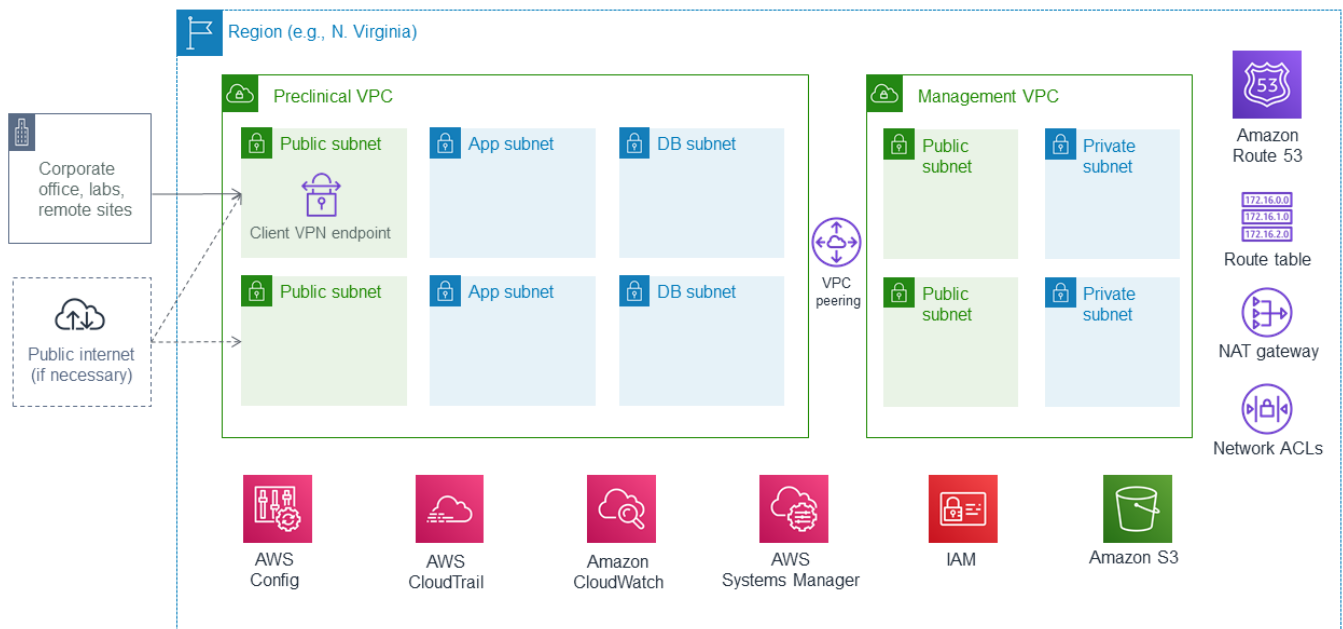
This Quick Start builds an informatics infrastructure for a biotech company on the Amazon Web Services (AWS) Cloud. It sets up a virtual data center by creating virtual private clouds (VPCs) and automatically configures this infrastructure for identity management, access control, encryption key management, network configuration, logging, alarms, partitioned environments (for example, to separate preclinical, clinical, and management processes), and built-in compliance auditing.

You can also use the Quick Start to automatically deploy the industry's leading scientific research applications into this informatics infrastructure. Currently, you can deploy an electronic lab notebook (ELN) and molecular R&D solution, a chemical compound registry, freezer management software, a biotech informatics and research infrastructure, and a genomics analysis environment. For more information, see the **Add-on products** tab.

This Quick Start sets up the following:

- A highly available architecture that spans two Availability Zones.
- A preclinical virtual private cloud (VPC) configured with public and private subnets according to AWS best practices, to provide you with your own virtual network on AWS. This is where informatics and research applications will run.
- A management VPC configured with public and private subnets, to support the future addition of IT-centric workloads such as Active Directory, security appliances, and virtual desktop interfaces.
- Redundant, managed NAT gateways to allow outbound internet access for resources in the private subnets.
- Certificate-based virtual private network (VPN) services through the use of AWS Client VPN endpoints.
- Private, split-horizon Domain Name System (DNS) with Amazon Route 53.
- Best-practice AWS Identity and Access Management (IAM) groups and policies based on separation of duties, designed to follow the U.S. National Institute of Standards and Technology (NIST) guidelines.

- A set of automated checks and alerts to notify you when AWS Config detects insecure configurations.
- Account-level logging, audit, and storage mechanisms designed to follow NIST guidelines.
- A secure way to remotely join the preclinical VPC network by using the AWS Client VPN endpoint.
- A prepopulated set of AWS Systems Manager Parameter Store key/value pairs for common resource IDs.
- (Optional) An AWS Service Catalog portfolio of common informatics software that can be easily deployed into your preclinical VPC.



[Read QuickStart online](#)

[View deployment guide for details](#)

Blog

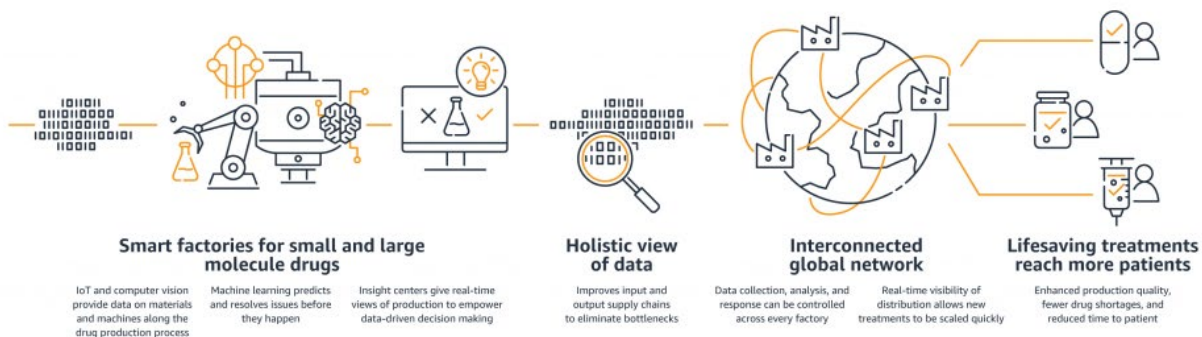
AWS and Novartis: Re-inventing pharma manufacturing

By Shez Partovi, MD and Ian Meyers

Today, AWS announced a new strategic collaboration with Novartis that will transform Novartis' core supply chain, manufacturing, and delivery operations through the use of AWS services. For several years, many AWS partners and life science customers have been using AWS services to modernize various aspects of manufacturing and supply chain. This strategic collaboration with Novartis is distinct by its broad strategic approach to transformation, and marks an acceleration of focused investment by AWS in this area. In this blog post, we will identify why manufacturing and supply chain transformations have become an urgent priority for pharmaceutical companies, and how AWS is uniquely positioned to assist customers along this journey. While the nature of the collaboration between Novartis and AWS spans supply chain, manufacturing, and distribution, the technical detail that follows will focus on the manufacturing portion. We will discuss how AWS is helping Novartis transform its manufacturing process by unifying access to all information and enabling Novartis to make quick and informed critical decisions.

Reinventing drug manufacturing and supply chains

Harnessing the power of AI, IoT, and ML to scale production globally



In recent years, targeted and more effective therapies and vaccines have been made possible by new biologic molecules. These therapies have had a positive impact on a broad range of diseases, including asthma, psoriasis, and various forms of cancer. These “large molecule”

therapies can be more complex to manufacture, and can create new logistical challenges since they include living organisms. At the same time, pharmaceutical companies need to reduce the cost of manufacturing their “small molecule” therapies as those drugs lose patent protection, while needing to scale up production to address growing demand from developing markets, such as China. Moreover, the pharmaceutical manufacturing process has historically been focused on producing very large quantities of a particular therapy. Newer, genetically engineered therapies are designed for smaller cohorts, and may even be personalized down to an individual level. The combination of these challenges requires manufacturers to make their upstream supply chains, manufacturing processes, and downstream distribution more visible, predictable, efficient, and adaptable. For many of our customers the ultimate goal is an automated system that responds in real time to all relevant information. This is true not just at the factory level, but also at the enterprise level — managers want to see throughput data for all factories globally in order to ramp certain production lines up or down, thereby minimizing drug shortages and optimizing capacity.

Amazon has spent years developing extremely sophisticated supply chain and automation systems that enable millions of products to be delivered to hundreds of millions of individuals, within hours. These technologies and experiences uniquely position AWS to help companies reinvent their supply chain and production systems. For example, in March of this year AWS and Volkswagen [announced](#) a collaboration to transform automotive manufacturing. When Novartis inquired about a potential collaboration in manufacturing and supply chain, our AWS Life Science practice saw an opportunity to help a pharmaceutical industry leader positively impact the lives millions of patients.

Novartis is a large multinational medicines company, operating more than 60 manufacturing sites that produce therapeutics used by nearly 1 billion patients in 155 countries, annually. Like any modern manufacturer, site control systems provide operational metrics to site managers about the efficiency of individual machines, and facilitate day-to-day maintenance tasks. Shift handover is a mission critical point where information about operational performance is transferred from one crew to another. Engineers know production lines inside and out, but have lacked the operational metrics to support moving from anecdotal to data driven decision making. Across the Novartis Technical Operations (NTO) group, it has been difficult and costly to develop standard metrics for global site efficiency, see global operational performance in a “single pane of glass”, and build advanced machine learning models to predict site performance.


Prior to this collaboration, Novartis manufacturing site metrics had already been centralized into a traditional Hadoop-based big data platform, which has enabled the creation of extensive operational reporting. Data is periodically ingested from local historians into Hadoop Distributed File System (HDFS), where it is batch processed and prepared for reporting. While this has delivered value to the business, the operational reports are based upon fixed datasets which may be out of date in relation to operational decision-making. Due to the use of legacy third-party vendors, this batch mechanism has not scaled to meet

the needs of the business, including: more users, additional metrics, advanced analytics and Machine Learning, and real-time data stream processing.

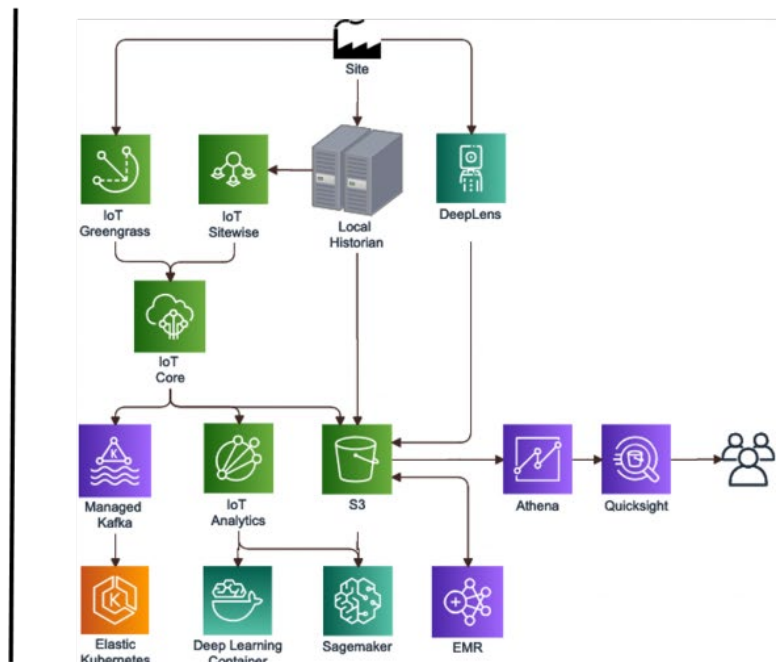
Through this collaboration, AWS and Novartis are jointly developing “Insight Centers” that provide real-time, interactive operational information to both site operators and corporate users around the globe. Insight Centers ingest site operational metrics through existing components such as local and global historians – platforms that capture industrial sensor data and add context and environmental meaning – as well as through new sensor integration. Existing historians will supply data from ‘brownfield’ sensors through nearly 20 SAP systems, while AWS IoT Greengrass edge devices will deliver richer data sources, such as images and video, and all data will be stored on Amazon S3. In addition to highly available and low cost object storage for this data, AWS IoT Core and IoT analytics provides data forwarding to an IoT optimized time series database where Novartis can perform real-time interpolation to facilitate predictive models.

Insight Centers will provide a cloud-native and ultra-scalable environment where existing big data processing technology can be run with minor changes due to compatibility with the most common frameworks and tools, but at lower cost, and enabling new uses cases that weren’t previously possible like computer vision powered line inspection. Daily production line & site operations will move to real-time data feeds, and add risk assessment models for proactive maintenance. Each site Insight Center will integrate with the others, providing a global view of the production capacity for each therapy. Finally, global Insight Center integration will deliver per product and process traceability, including manufacturing status, inventory levels, and cost.

Current Reporting Model



New Insight Centers



As the volume, variety, and veracity of real-time data flow into the Insight Centers increases, the AWS AI/ML services will enable Novartis to create sophisticated ML models to advance its operational forecasting. The proposed technical architecture will create dynamic and flexible views of the Novartis global manufacturing processes — views that were previously un-achievable. Novartis will soon be able to fully leverage AWS's machine learning and AI services to predict machine failure, develop digital twins, and build forecasting models of demand that will enable more efficient supply. These sophisticated use cases build upon the current initiatives already under way at Novartis. For example, Novartis is currently using Amazon SageMaker to build a computer vision-based model that will determine line clearance. Amazon SageMaker is a fully managed service that covers the entire machine learning workflow to label and prepare data, develop an algorithm, train the model, tune and optimize it for deployment, make predictions, and take action. In addition to Amazon SageMaker, AWS offers fully managed API-based AI services such as Amazon Textract and Amazon Comprehend. Novartis is also currently using these AI services to automate analysis of data extracted from printed manufacturing documents.

The Insight Centers will also be able to take advantage of AWS Identity & Access Management (IAM) to build a robust security posture and simplify entitlement management. IAM is a key co-traveler to data governance to ensure secure and appropriate sharing of manufacturing data among data scientists. Once approved, data scientists will have an open environment in which they can experiment with new supply chain and manufacturing optimization solutions that were previously constrained by legacy operating environment.

With these key features (and others to follow), Novartis Insight Centers will make quality manufacturing data available in real-time, and enable fast, informed decision-making. Improving the agility of the manufacturing process ultimately improves production quality while reducing cost, unnecessary inventory, and machine downtime. Transforming manufacturing will add value to Novartis operations, and will ultimately deliver the most important value to patients receiving Novartis therapies. This transformation roadmap will enable a future where small batch and even personalized treatments can be developed and manufactured faster, with higher quality, lower costs, and delivered on schedule.

[Read Blog post online](#)

Executive Summary

Why AWS for GxP regulated workloads?

Move, or build, GxP workloads in the cloud with help from the AWS Life Sciences Practice

The AWS life sciences practice and dedicated GxP experts help biopharma and medical device companies build new, or move, regulated workloads to the cloud. Organizations can achieve higher productivity, increased operational resilience, greater agility, improved transparency and traceability, and a [lower TCO](#)¹ and [carbon footprint on AWS](#)² compared to on-premises systems. AWS offers a secure cloud platform for automating compliance processes and providing enhanced management related to running your GxP workloads.

Some of the biggest enterprises and innovative startups in life sciences have moved, or built new GxP systems on the AWS Cloud. Decades of expertise in the life sciences space enable us to help you move away from on-premises regulatory paradigms as your trusted advisor to becoming regulated cloud native.

¹ IDC, [Fostering Business and Organizational Transformation to Generate Business Value with Amazon Web Services](#), 2018

² 451 Research, [The Carbon Reduction Opportunity of Moving to Amazon Web Services](#), 2019

Evidence of success: AWS Customers

[Moderna](#) power their R&D and manufacturing environments on AWS to accelerate the development of mRNA therapeutics and vaccines.

[Idorsia](#) house 90% of IT operations on AWS, including GxP compliant virtual data centers, giving them 75% cost savings compared to physical data centers

[Bristol Myers Squibb](#) use AWS to create consistent, scalable, and repeatable process to streamline compliance with GxP and other regulatory requirements.

Evidence of success: AWS Partners

[Aizon](#) use AWS to enable real-time visibility and predictive insights that address GxP compliance.

[Core Informatics](#) developed a platform on AWS that enables GxP-validated applications in the cloud for its customers.

[Tracelink](#) choose AWS to support its life science cloud solution to meet customer compliance needs.

Take advantage of AWS cloud security and security management best practices

With AWS, you benefit from the significant resources applied to ensuring security of the cloud. **The AWS cloud is designed and managed in alignment with IT security standards such as SOC 1, 2, 3; ISO 9001/ISO 27001/ISO 27017/ISO 27018; HITRUST; FedRAMP; and CSA Security, Trust and Assurance Registry (STAR).** Plus, the underlying cloud infrastructure for your GxP workloads are monitored and protected 24 x 7 by redundant and layered controls and continuous verification and testing, and automation.

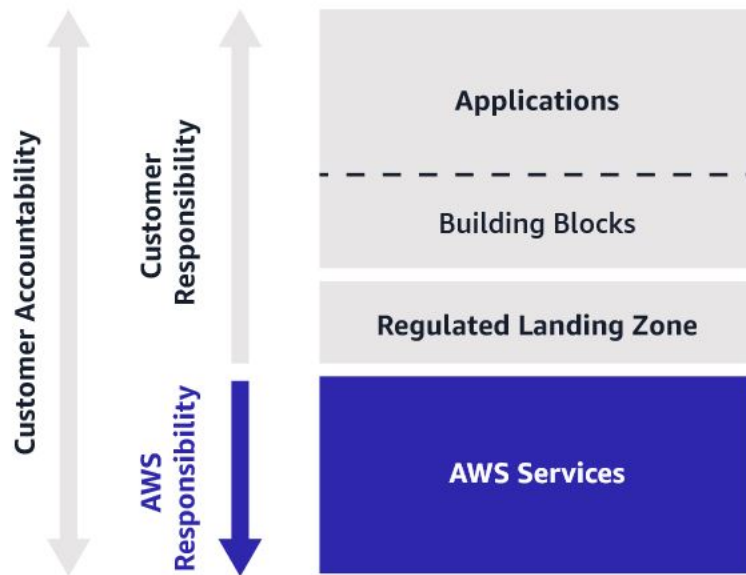
In addition, security best practices from AWS help you define and build a set of security policies and processes for your organization that protects your data and assets. You get an overview of how to protect your assets on AWS and manage access to AWS resources, along with suggestions for how you can secure your data, operating systems, applications, and overall infrastructure in the cloud.

Leverage the AWS Quality Management System for your GxP-regulated products

Using an ISO 9001:2015 certified provider like AWS can make your own certification process [easier](#). Formal, documented policies and procedures from AWS life science experts guide you through AWS operations and information security and address purpose, scope, roles, responsibilities, and management.

Build on a cloud backed by the AWS shared responsibility model

In the shared responsibilities model for compliance, AWS is responsible for the compliance **of the cloud**, and customers are responsible for compliance **in the cloud**.



AWS' responsibility for compliance "of" the cloud includes physical hardware, global infrastructure, and how regions, availability zones, and edge locations function and operate. It also encompasses the virtual infrastructure for compute, storage, database, and network services. There is also a set of AWS services available for setup and governance, security and user management, operations, DevOps, and DevSecOps. The Regulated Landing Zone is your GxP compliant production environment.

Leverage the most comprehensive set of cloud services

Setup and governance services

- [AWS Control Tower](#) Set up and govern a new, secure, multi-account AWS environment based on best practices.
- [AWS Config](#) Assess, audit, and evaluate the configurations of your AWS resources.
- [Private Amazon Machine Images \(AMIs\)](#) Create Amazon Machine Images that only you can use to call an instance.
- [AWS Audit Manager](#) Automatically collect and organize evidence as defined by each control requirement.

Security and user management services

- [AWS CloudTrail](#) Leverage governance, compliance, operational auditing, and risk auditing of your AWS account.
- [Amazon CloudWatch](#) Log, continuously monitor, and retain account activity related to actions across your AWS infrastructure.

- [Amazon Inspector](#) Automatically assess applications for exposure, vulnerabilities, and deviations from best practices.
- [Amazon Cognito](#) Enable the delivery of scoped temporary credentials to mobile devices and other untrusted environments.
- [AWS Identity and Access Management \(IAM\)](#) Manage access to AWS services and resources securely.
- **End-to-end encryption** Get the configuration needed for data security.

Operational services

- [AWS Lambda](#) Run code without provisioning or managing servers with a serverless compute service that creates workload-aware cluster scaling logic.
- [AWS Security Hub](#) Get a comprehensive view of your security alerts and security posture across your AWS accounts.

DevOps and DevSecOps services

- [AWS CloudFormation](#) Leverage a service that treats infrastructure as code.
- [AWS Service Catalog](#) Create and manage catalogs of IT services that are approved for use on AWS.
- [AWS Code Pipeline](#) Get help with automating your release pipelines for fast and reliable application and infrastructure updates

Simplify change management with AWS Config and other services

Change management doesn't have to be complex. When your GxP regulated workloads move to AWS, these processes are automated and simplified. Continuous deployment methodology ensures changes are automatically built, tested, and deployed, with the goal of eliminating as many manual steps as possible. For example, AWS Config controls how changes to configuration items are made, including an assessment of the potential impact on the GxP applications that it supports.

Use AWS services that make configuration easy

AWS services help simplify your cloud configuration process by applying tags, labeling, collecting, and organizing resources and components within services. AWS services continuously monitor and record your AWS resource configurations and lets you automate the evaluation of recorded configurations against configuration goals.

Track and measure your resources and infrastructure with AWS monitoring

You can monitor resources and the applications on AWS, collect and track metrics, monitor log files, set alarms, and automatically react to changes in AWS resources. You can also track and log the behavior of the customer application landscape, triggering events based on that behavior.

Ensure business continuity

You can't predict a disaster, but you can be prepared for it. For disaster recovery, AWS Regions provide multiple physically separated and isolated Availability Zones connected with low-latency, high-throughput, and highly redundant networking. [Availability Zones](#) are more highly available, fault tolerant, and scalable than traditional single or multiple data center infrastructures and are well-suited for keeping your GxP workloads online. A multi-Availability Zone architecture can withstand a complete failure of one or more zones.

The ability to back up and restore is required for all validated applications. It is therefore a common capability that can be centralized as part of the regulated landing zone. [AWS Backup](#) is a fully managed backup service that makes it easy to centralize and automate the backup of data across AWS services.

AWS is right for GxP

AWS helps your life sciences organization innovate while providing the infrastructure, services, and industry experts that best support your GxP workloads. Access the new [GxP on AWS whitepaper](#) to dive deeper on best practices for architecting for regulated workloads in the cloud.

[Read Whitepaper online](#)

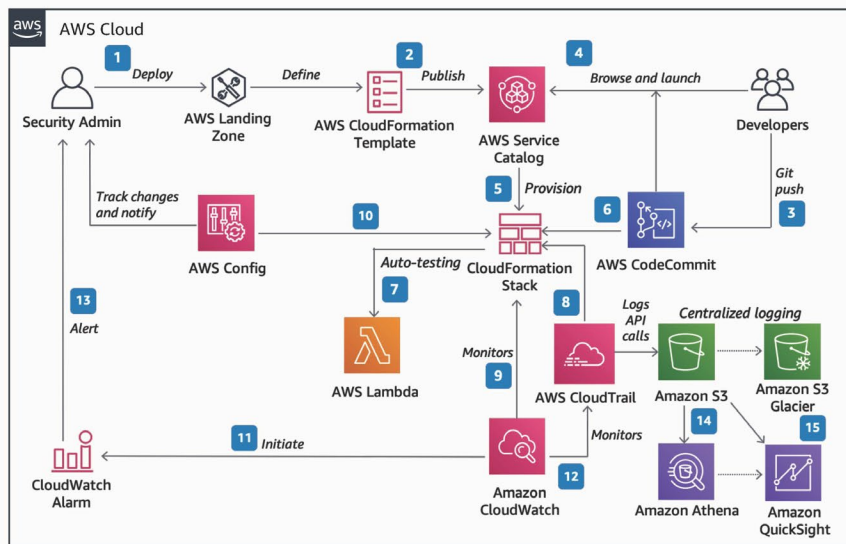
Reference Architecture

GxP Compliance Automation

Building a Secure and Compliant GxP Workload on AWS

GxP Compliance Automation

Building a Secure and Compliant GxP Workload on AWS



- 1 **AWS Landing Zone** allows the security administrator to automate the set-up of an environment for running secure and scalable workloads. Security admin defines an **AWS Service Catalog** product (for example, a GxP application) using **AWS CloudFormation** templates.
- 2 Security admin publishes the template for developers in the **AWS Service Catalog**. Developers use this framework to further enhance the template based upon the application requirements.
- 3 Developers take the framework and modify applications to further enhance it under **Git** source control and use **AWS CodeCommit** to fully manage the private code repository.
- 4 Developer deploys the modified code from **CodeCommit** to their GxP infrastructure, using **AWS Service Catalog** to launch the product they need as an **AWS CloudFormation stack**.
- 5 The stack automatically provisions the necessary AWS resources based on what has been committed to the code repository as specified by the developer.
- 6 **AWS Service Catalog** is at the center of this architecture, so developers can release their source code without needing to access to any underlying resources or go through security administrators.
- 7 Automate the testing/Installation qualification process using **AWS Lambda** or Python program and create a test summary/qualification report automatically in an **Amazon S3** bucket.
- 8 All individual **CloudTrail** logs, **VPC flow logs**, and **AWS Config** changes are aggregated into a centralized **S3** bucket in a separate **AWS** account.
- 9 The security administrator configures, monitors, and sets up automated alerts on changes and on the health of the stack via **Amazon CloudWatch**.
- 10 When the stack is changed, change events are recorded and tracked through **AWS Config**. Out of compliance events are displayed in dashboard.
- 11 To indicate that something may be out of compliance, **CloudWatch** can initiate alarms based on rules that you design.
- 12/13 **CloudTrail** monitors API calls made against the AWS environment. The administrator is notified/alerted by **CloudWatch Events** when something changes that could cause the system to be non-compliant.
- 14/15 Log Data is queried and converted into a human readable format like CSV using **Amazon Athena**, for any audit purpose. Visualize **CloudTrail Logs** using **Amazon QuickSight**.



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

AWS Reference Architecture

[View Reference Architecture online](#)

Solution

Service Workbench on AWS




What does this AWS Solutions Implementation do?

Service Workbench on AWS is a cloud solution that enables IT teams to provide secure, repeatable, and federated control of access to data, tooling, and compute power that researchers need. With Service Workbench, researchers no longer have to worry about navigating cloud infrastructure. They can focus on achieving research missions and completing essential work in minutes, not months, in configured research environments.

With Service Workbench on AWS, researchers can quickly and securely stand up research environments and conduct experiments with peers from other institutions. By automating the creation of baseline research setups, simplifying data access, and providing price transparency, researchers and IT departments save time, which they can reinvest in following cloud best practices and achieving research reproducibility.

AWS Solutions Implementation overview

The diagram below presents the architecture you can automatically deploy using the solution's implementation guide and accompanying AWS CloudFormation template.



Service Workbench on AWS Solutions Implementation architecture

The [AWS CloudFormation template](#) deploys a test environment of Service Workbench with default settings, which provides the solution's website hosted on [Amazon Simple Storage Service](#) (Amazon S3), made available through [Amazon CloudFront](#). The web frontend authenticates using [Amazon Cognito](#) and leverages [Amazon API Gateway](#) to invoke the solution's microservices ([AWS Lambda](#) functions and [AWS Step Functions](#)).

These microservices interact with [Amazon DynamoDB](#) to manage the content, users and AWS accounts in [AWS Organizations](#) to access data in S3 and instantiate out-of-the-box compute instances for [Amazon EMR](#) with Hail, [Amazon SageMaker](#) with Jupyter notebooks and [Amazon EC2](#) with Windows and Linux operating systems. The CloudFormation templates for these compute instances are hosted in [AWS Service Catalog](#) for flexibility and to allow the simple addition of custom templates.

After you install the solution, you can log into the web frontend using the root password and URL provided by the install process, create administrator- and end-user logins, associate an AWS account with Service Workbench on AWS and start testing the solution and its capabilities.

Features

Compute provisioning leveraging Service Catalog and CloudFormation

Service Workbench on AWS provides an easy-to-use, self-service portal for researchers to provision compute environments. Administrators can use AWS CloudFormation and AWS Service Catalog to create environment templates that are secure and compliant with organizational policies.

Simplified account vending through AWS Organizations

Service Workbench on AWS provides a simple mechanism for creating new AWS accounts via AWS Organizations. This enables fine-grained control associating research projects with budgets and billing.

Data and compute sharing to speed collaboration

Service Workbench on AWS connects to multiple identity providers simultaneously enabling different organizations to collaborate. Users can share data and compute templates across accounts subject to administrator review and approval.

[View Implementation Guide online](#)

[Find more Biopharma Solutions here](#)

Case Study

Bristol Myers Squibb Builds Enterprise Data Lakes on AWS



[Bristol Myers Squibb](#) is a global biopharmaceutical company committed to discovering, developing, and delivering innovative medicines that help patients prevail over serious diseases. Before moving to the Amazon Web Services (AWS) Cloud, Bristol Myers Squibb was building data lakes using various tools and methods, which created inconsistencies, technical debt, and a higher cost to run. Brian Zellner, associate director of R&D data lakes at Bristol Myers Squibb, says, “When we moved to AWS, we started with a multi-tenant concept that allows us to use services in a consistent way.”

In addition, Bristol Myers Squibb is leveraging [AWS Glue](#) to build out a scalable infrastructure with a pay-as-you-go model, which “gave us great advantage over what we were doing before, so we’re able to ingest data faster and transform it and get faster insights,” says Zellner.

“Speed is a core tenant of what we’re trying to do at Bristol Myers Squibb. The faster that we can develop drugs, the sooner they can get to the patients who need them,” explains Zellner.

[Watch Video online](#)

To learn more, visit aws.amazon.com/health

Training

Intro to Healthcare and Life Sciences (HCLS) Compliance on AWS



Description

This self-paced course is designed to introduce you to fundamentals of cloud computing under healthcare and life sciences compliance regimes such as HIPAA, GxP, and international data privacy. This course allows you to test new skills and apply knowledge to your working environment through a variety of practical exercises.

Course Objectives

This course teaches you how to:

- Define the fundamentals of HCLS legal and compliance frameworks
- Understand the fundamentals, identify common scenarios and discuss customer considerations for HIPAA and HITRUST, GxP, and International Data Privacy in Healthcare and Life Sciences

Intended Audience

This course is intended for:

- IT business-level users and professionals interested in cloud compliance practices
- Cloud solution architects who are new to healthcare and life sciences compliance
- HCLS compliance professionals who are new to cloud computing or AWS

Course Outline

This course covers the following concepts:

- Introduction to healthcare and life sciences compliance frameworks
- HIPAA & HITRUST in the United States
- GxP Worldwide
- International Data Privacy

[Start course online](#)

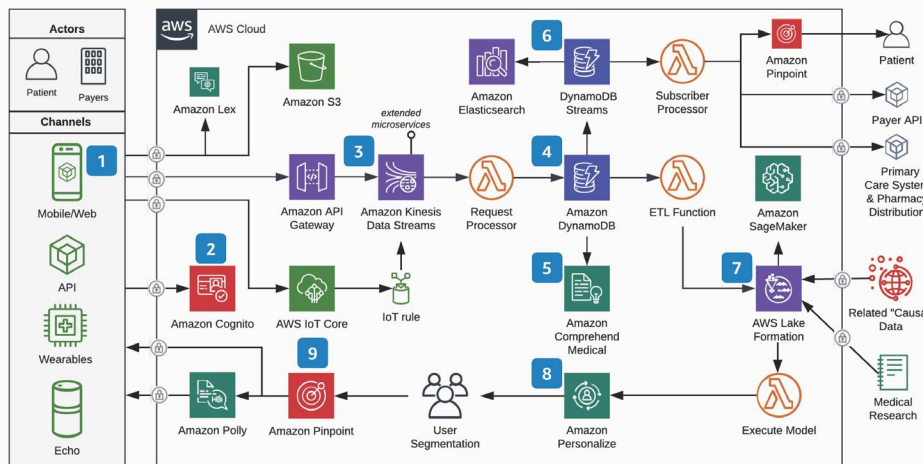
Reference Architecture

Patient Engagement with Amazon Pinpoint and Amazon Personalize

Targeted and meaningful patient engagement on AWS

Patient Engagement with Amazon Pinpoint and Amazon Personalize

Targeted and meaningful patient engagement on AWS



- 1 For ease of interaction, patients and payers can use a number of optimized channels in this solution, with this example describing how mobile apps, digital therapeutics devices, chatbots, or voice-enabled devices can be used for key interactions such as **patient registrations, appointments, payer status alerts, prescriptions, medicine distribution wellness advice, etc.**
- 2 **Amazon Cognito** uses Web-Identity Federation authentication to grant access to the system using familiar methods like a patient's Facebook, Amazon, or Google credentials.
- 3 **Amazon Kinesis Data Streams**, a scalable and durable real-time streaming service, decouples the system, provides event-driven messaging, and supports an extended microservices architecture.
- 4 Patient data, attributes, transactions and session state are encrypted at-rest and securely managed in the high-performance and scalable **Amazon DynamoDB** NoSQL datastore.
- 5 Patient symptoms and medical notes are analyzed and tagged for automated decision making by **Amazon Comprehend Medical** supporting next best actions.
- 6 **Amazon DynamoDB Streams** triggers an **AWS Lambda** function that sends notifications to subscribers via an API or via **Amazon Pinpoint**. **Amazon ElasticSearch** enables the analysis and visualization of data across the solution.
- 7 Raw and curated data generated through every channel is ingested into the **AWS Lake Formation** data lake for processing and analytics, allowing for combined external datasets and finding causation and correlation in data using **Amazon SageMaker** ML models.
- 8 **Amazon Personalize** trained model is executed against any new patient behaviors or insights in real-time to place patients in a more suitable segmentation in real-time.
- 9 After a patient is segmented into one or more groups, **Amazon Pinpoint** engages patients with relevant content on their preferred channel.

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

AWS Reference Architecture

[View Reference Architecture online](#)

Blog

Executive Conversations: Accelerating COVID-19 vaccine development with Marcello Damiani, Chief Digital and Operational Excellence Officer at Moderna



By Todd Weatherby and Kelli Jonakin, Ph.D

Marcello Damiani, Chief Digital and Operational Excellence Officer at [Moderna](#), joins Todd Weatherby, Vice President of [AWS Professional Services Worldwide](#), for a discussion on developing Moderna's COVID-19 vaccine, scaling systems to enable global distribution, and leveraging cloud technologies to accelerate processes. Moderna is pioneering a new class of medicines made of messenger RNA, or mRNA, with significant potential to improve the lives of patients.

This Executive Conversation is one of a series of discussions held with those progressing their industries where we seek to learn more about their discovery, ingenuity, and contributions to healthcare and life sciences.

Todd Weatherby: *How does the Moderna vaccine for COVID-19 work?*

Marcello Damiani: Our focus at Moderna is on messenger RNA, which is an information-based molecule that goes to and from DNA retrieving genetic code. Every day, your body naturally produces billions of messenger RNAs. At Moderna, we send messenger RNA with instructions for cells to produce specific proteins that can cure or prevent diseases. So when we received the genetic sequence of the COVID-19 virus from China, we identified proteins on the surface of the virus that would serve as the foundation for all our COVID-19 vaccine activities. Through our vaccine, we send instructions via the messenger RNA to produce these proteins that have exactly the same signature as the COVID-19 virus, so the immune system can mount a response.

TW: *How has cloud technology enabled vaccine development?*

MD: Six years ago, we started building databases and information-based activities to support all our programs. Today, we're fully cloud based, and our scientists don't go to the lab to pipette their messenger RNA and proteins—they go to our web portal, the Drug Design Studio that sits in the AWS Cloud. Through the portal they can access public and private libraries that contain all the messenger RNA that exist and the thousands of proteins they can produce. After that, they just press a button and the sequence goes to a fully automated, central lab where we collect data at every step. Over the years, data from the portal and lab has helped us improve our sequence design and production processes and

improve the way our scientists gather feedback, our data scientists built algorithms to accelerate the design of sequences for messenger RNA. In terms of research, all our algorithms rely on computational power from AWS to further our science. On the manufacturing side, everything is fully digitized, paperless, and sits on AWS—including our manufacturing execution system. In fact, we were among the first in the industry to build our manufacturing execution systems on AWS. Then for clinical trials, most of our clinical trial data resides on AWS and we use [Amazon Redshift](#) for analytics, including those that supported our COVID-19 vaccine development.

TW: *How did Moderna scale operations to reach billions of people?*

MD: As we were developing the COVID-19 vaccine, in parallel, we were also working on scaling our manufacturing and supply chain capabilities to be ready for massive distribution. We focused on scaling up our manufacturing execution system and enterprise resource planning (ERP) platform, working closely with AWS Professional Services and Answerthink. That work involved upgrading our SAP ERP so it could leverage high availability and infrastructure as code to meet distribution demand. And by consolidating workstreams, we were able to shorten our timeline for upgrading the system from five months to just two. Now we're able to deliver vaccines in more than 30 countries to people who desperately need it.

TW: *What's next for Moderna?*

MD: Now that we have the mechanisms in place that helped us launch our first vaccine, we're focusing on how to collect data across the globe to improve our clinical trials and build the next generation of drugs that can help the world. Scaling up with AWS is going to be key for accomplishing our vision—across manufacturing, science, clinical trials, and commercialization. Having the flexibility to scale computational power up or down for sophisticated algorithms will be crucial. Throughout the company we'll be relying on AWS infrastructure to support us.

[Read Blog post online](#)

[Find more Life Sciences blog posts](#)

Executive Brief

Biopharma on AWS

Bring advanced therapeutics to market faster with AWS

The pressure is on to develop advanced therapeutics quickly and cost-effectively. Accelerate innovation and streamline operations with a technology partner that helps you keep pace in a changing industry.

Accelerate innovation with a trusted partner

Join fellow industry leaders on the most secure and reliable cloud platform, with an unmatched breadth and depth of services, and a deep understanding of life sciences industry regulations and compliance.

Actualize the potential of your data

Unlock insights through data liquidity that unites R&D discoveries, clinical trial insights, supply chain information, sales results, and real-world data for a holistic view of the therapeutic lifecycle.

Achieve operational excellence and cost savings

Streamline operations and boost business agility by accessing the broadest choice of cloud computing, analytics, and machine learning (ML) services.

Personalize customer engagements

Build your own personalized customer experiences at every major touchpoint of care and sure-fit drug launches using the same customer-centric services that power Amazon.com.

Why AWS for Biopharma

| | | | | | |
|--|--|---|--|---|--|
| 90% Of Top-10 global biopharma companies use AWS | 130+ HIPAA - eligible services | 500+ Features and services focused on security and compliance | 90% Potential savings on compute | 24 Regions and 77 availability zones to support global collaboration and data sovereignty | 200+ Fully featured services from datacenters globally |
|--|--|---|--|---|--|

Biopharma use cases

Explore six key areas where AWS helps biopharma organizations benefit from cloud technology

Research and discovery

Accelerate R&D discoveries and de-risk candidates faster with high performance computing, analytics, and ML. Leverage the most extensive cloud infrastructure to power secure, global collaboration across datasets

Clinical development

Power the entire clinical development process—from designing effective trials using analytics and ML capabilities, to unlocking virtual trial capabilities, to accelerating the regulatory submission process

Manufacturing and supply chain

Modernize manufacturing processes by breaking down data silos, deploying analytics and ML to generate insights, utilizing more cost-effective infrastructure, and enhancing regulatory compliance

Commercialization

Streamline commercial product launches with democratized data access, an optimized biopharma salesforce, and omni-channel communications to engage healthcare professionals and customers

Post-market surveillance and patient support

Accelerate real-world evidence collection, effectively scale to handle adverse event reports, and achieve greater visibility into post-market surveillance data for improved patient support

Security and industry compliance

Achieve near continuous compliance for GxP-regulated workloads and leverage services that help you comply with GxP, HIPAA, HITRUST, and GDPR standards throughout the entire biopharma value chain

[Download Executive Brief](#)

[Learn more about Biopharma on AWS](#)

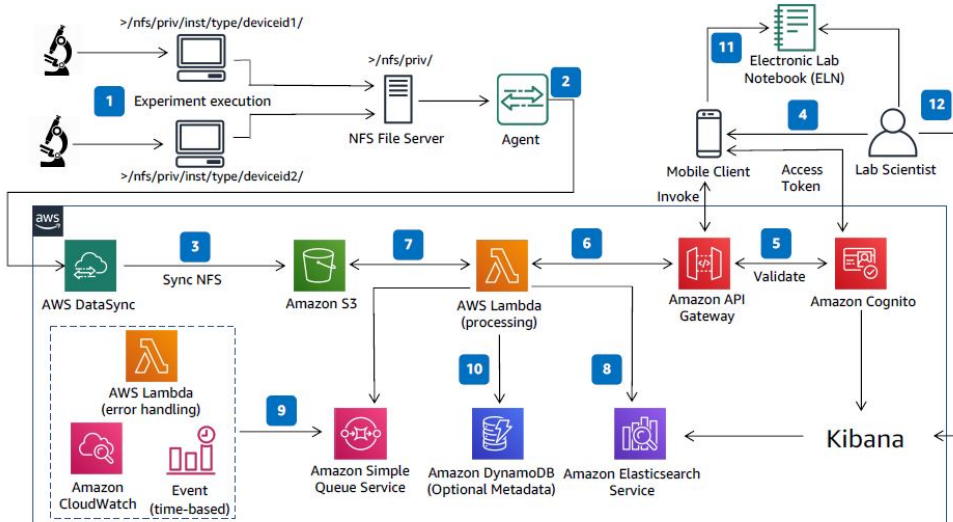
Reference Architecture

Lab instrument log acquisition and analytics

Build a data pipeline to automate lab instrument log ingestion, collate metadata, and transform and downstream to FDA compliant electronic laboratory notebooks. Eliminate multiple second-person reviews for a [drug discovery and development](#) process. Experience easy search capabilities and REST API support.

Lab instrument log acquisition and analytics

Build a data pipeline to automate lab instrument log ingestion, collate metadata, and transform and downstream to FDA compliant electronic laboratory notebooks. Eliminate multiple second-person reviews for a [drug discovery and development](#) process. Experience easy search capabilities and REST API support.



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

AWS Reference Architecture

- 1 All instruments create the output log file in their respective network file system (NFS) shares on the machine attached to the instrument.
- 2 An agent (virtual machine) is deployed in the user environment, and is used to read data from their NFS storage.
- 3 An AWS DataSync task created in the customer's AWS account syncs data from NFS to a customer-owned Amazon S3 bucket.
- 4 A lab scientist invokes a client application to look for an experiment output log file in Amazon S3 for an associated experiment ID.
- 5 The client application gets a token from Amazon Cognito and invokes the API with the experiment ID and the Amazon S3 file path (and possibly more IDs as needed from the client application).
- 6 Amazon API Gateway validates the token from Amazon Cognito, then calls the processing AWS Lambda function (one Lambda function for each type of log file), and returns response after the Lambda function runs.
- 7 The Lambda function reads the file from Amazon S3, and adds the experiment ID as a tag to the Amazon S3 log file.
- 8 The Lambda function transforms the S3 file, applies any business rules, converts the log to one or multiple JSON records, and pushes data to Amazon Elasticsearch Service.
- 9 If the processing Lambda function fails, it writes the S3 file info to an error monitoring queue. A Lambda function based on a pre-defined timed event can read from the queue to reprocess or to inform support staff.
- 10 Optionally, a processing Lambda function maintains the S3 file info with metadata (experiment IDs) in an Amazon DynamoDB table which can be consumed by other applications.
- 11 The processing Lambda function sends the JSON response to Amazon API Gateway, which is posted to an electronic lab notebook (ELN) via Mobile Client.
- 12 The lab scientist can access data in ELN, and can also discover and query data via Kibana for other experiments.

[View Reference Architecture online](#)

Related Videos, Podcast



Streamline Innovation in Biopharma with AWS

AWS is the trusted technology partner to help modernize every stage of the biopharma value chain so you can bring differentiated therapeutics to market faster. Organizations rely on AWS to help enhance data liquidity, optimize for operational agility and excellence, and personalize customer engagement.

[Watch Video online](#)

Evidation Health Enables Next Generation Clinical Studies on AWS

Evidation Health sought to understand everyday behaviors that create better health outcomes by running digital studies using “invisible” data—data collected passively in everyday life from devices like smartphones and wearable sensors. On AWS, Evidation Health gained automated elastic scaling, complete encryption, and simpler, predictable deployments.

[Watch Video online](#)

PwC: Building an Efficient, Intelligent, Serverless Pipeline for the Pharma Industry

Pharma and life science companies have to comply with strict regulations when studying and reporting on adverse patient outcomes. This new data pipeline can process more data, more intelligently and at lower cost using AWS Lambda and Amazon Comprehend Medical. Host: Simon Elisha, Head of Solutions Architecture, WWPS, A/NZ; Customer: David McKeown, Solution Architect

[Watch Video online](#)

[Watch more Life Sciences Architecture videos](#)

Bristol-Myers Squibb: Enabling DevOps with AWS Service Catalog

Assembling new application stacks on AWS can take time. In some cases, it can lead to duplication of CloudFormation elements already used for other application stacks. In this video, Bristol-Meyers Squibb (BMS) talks about their need to optimize application build and standardize application building blocks approved for re-use. To do so, BMS uses pre-approved application stack elements in their AWS Service Catalog, enabling developers and business owners to deploy new application stacks quickly and safely with approved components in a self-service model.

[Watch Video online](#)

AWS Innovation with EB Research

Host Sara Armstrong, Sr. Manager of Worldwide Prototyping at AWS speaks with Michael Hund, CEO of EB Research about their mission to accelerate treatments and unlock a cure for Epidermolysis Bullosa (EB), a devastating and life-threatening genetic skin disorder that affects children from birth. Get a behind the scenes look at how their unique approach to funding, "Venture Philanthropy", is leveraging to power of cloud computing to provide a common research and collaboration platform for a consortium of medical professionals, research scientists, as well as patients and parents. What does it take to achieve a more promising future for children born with life-threatening genetic skin disorders? You'll learn how an organization founded with support of Jill and Eddie Vedder, the lead singer of Pearl Jam took on a unique journey of philanthropy and healthcare technology to achieve this vision. Listen to this episode of Innovation Ambassadors for an enlightening discussion between host Sara Armstrong, Michael Hund of EB Research, and Ryan Jancaitis, Global Product Management Leader for the Envision Engineering Prototyping team at AWS.

[Listen to Podcast online](#)