

AWS Whitepaper

Migration and Modernization Strategy for Integrated Eligibility Systems



Migration and Modernization Strategy for Integrated Eligibility Systems: AWS Whitepaper

Copyright © 2024 Amazon Web Services, Inc. and/or its affiliates. All rights reserved.

Amazon's trademarks and trade dress may not be used in connection with any product or service that is not Amazon's, in any manner that is likely to cause confusion among customers, or in any manner that disparages or discredits Amazon. All other trademarks not owned by Amazon are the property of their respective owners, who may or may not be affiliated with, connected to, or sponsored by Amazon.

Table of Contents

Abstract and introduction	i
Introduction	1
Motivation for cloud migration and modernization	3
Improve customer and beneficiary experience	3
Enhance workforce productivity	3
Improve insights into program operations	4
Handle data at scale and improve security	4
Handle large call volumes	4
Reduce costs and minimize TCO	4
Modernization focus areas for IES benefits administrators	5
IES – High-level cloud migration strategy	6
Assess	6
Mobilize	7
Migrate and modernize	7
IES – Legacy patterns and migration approaches	9
Migration options	9
Legacy system migration	9
Oracle or SQL Server database migration	10
Middleware	10
Application migration	10
VMWare based workloads	10
Red Hat OpenShift based workloads	10
Migration implementation approach	11
Data migration	11
Database migration	12
Application and server migration	12
License migration	12
IES modernization – User perspective	13
Beneficiary and end user experience	13
Deploy scalable omni-channel contact centers	13
Provide artificial intelligence (AI) powered speech analytics	13
Develop self-service capabilities	14
Provide language translation capabilities	14
Build effective campaign management strategies	14

Case worker and manager productivity	14
Build review and approval work-flow automation	15
Use technology to minimize errors	15
Build ML models to identify anomalies and fraud	15
Program leadership and benefits administrators	16
IES modernization – System and application perspective	17
Application modernization	17
Microservices architecture	17
Serverless and fully managed services	18
Containers	18
Database modernization	19
Data and storage modernization	20
Operations and DevOps modernization	21
Automating deployment with continuous integration and continuous delivery	21
Managing infrastructure as code	21
Monitoring and logging	22
IES migration and modernization summary	23
Example – Windows application modernization	23
IES reference architecture summary	25
Conclusion	26
Contributors	27
Further reading	28
Document history	29
Notices	30
AWS Glossary	31

Migration and Modernization Strategy for Integrated Eligibility Systems

Transforming benefits and eligibility systems using cloud computing

Publication date: **November 10, 2021** ([Document history](#))

Introduction

In the fiscal year 2019, the US federal, state, and local government agencies spent approximately \$2.5 trillion on various social and safety net programs, which included Social Security, Medicare, Medicaid, Supplemental Nutrition Assistance Program (SNAP), and others, that assist low-income families. (Source: <https://www.cbpp.org/research/federal-budget/where-do-our-federal-tax-dollars-go>) Tens of millions of customers apply for these benefits every year. (Source: <https://www.medicaid.gov/medicaid/program-information/medicaid-and-chip-enrollment-data/report-highlights/index.html>) The enrollment can include complex eligibility determination and adjudication processes. In most cases, beneficiaries have to wait several weeks before their cases are adjudicated due to the high volume of applications. Additionally, it takes a large work force to review and process these applications which are submitted multiple ways (for example, web, mail-in, or contact center communications).

Integrated eligibility systems (IES) provide the enabling technology for some of these health and human services programs. They provide multiple channels for application and recertification intake—complex rules engines that use data to calculate benefit amounts. Additionally, IES includes reporting, appeals, and investigations modules that monitor program integrity and federal compliance. After the passage of the Affordable Care Act (ACA) in 2010, state agencies took advantage of enhanced matching funds to modernize their IES.

IES have traditionally been monolithic and tightly coupled systems that are difficult to customize for specific state needs. As a result, these systems often lack the flexibility to adapt to program and policy changes. Additionally, benefits administrators face a number of challenges with IES. These include:

- Increasing application backlogs and delayed benefits to citizens
- Complex case review and adjudication processes
- The inability to handle increasing call volumes for follow-up activities, including interviews, application status, and appeals

- The inability to provide deep insights into program operations, such as incidents of fraud, waste, and abuse, and the means to drive improvements

The benefits of IES cloud migration include improving business and mission agility, lowering the total cost of ownership (TCO), and improving system resiliency. IES modernization results in better beneficiary experience, improved case worker productivity, high visibility into program operations, and reduced fraud, waste, and abuse.

Motivation for cloud migration and modernization

Many of the legacy IES make it challenging for the agencies to be agile and flexible in dealing with eligibility and enrollment processes. There are three main personas that are relevant for IES:

- The end-user or beneficiary
- Agency workforce and case workers
- Program leadership and benefits administrators

The following sections focus on these personas and outline the challenges and motivation for administrators to migrate and modernize IES.

Improve customer and beneficiary experience

Timely and accurate application processing is critical in helping the customers that need immediate assistance with healthcare, cash, or nutrition benefits. Legacy systems often provide a poor consumer experience as they cannot scale to meet the surge during an open enrollment or an unexpected situation (such as COVID-19). These systems also lack the capabilities required to provide a digital experience to customers, such as self-service, mobile application submission, text communications on case status, and online scheduling for interviews.

Enhance workforce productivity

Agency workforces have to deal with large volumes of applications and manual processes that result in increasing application backlogs and delayed benefits to consumers. Benefits application documents can include federal tax forms, pay stubs, social security cards, and other documents. These documents are in multiple formats, including PDFs and images, and are submitted from various sources, including mobile devices, the web, mail-in, and contact centers.

The workforce spends a significant amount of time processing, and validating these documents. Often, this process is manual, which is error prone and requires specific domain expertise. Additionally, manual processes take time away from case workers which can otherwise be spent in high value activities such as ensuring that a constituent is getting the best available support.

Improve insights into program operations

Having access to data-driven insights allows agencies to build programs and advocate for innovative policy changes to better serve constituents. These insights include visibility into enrollment models, backlogs, budget forecasting, fraud, waste, and abuse. Often, benefits program leadership lacks the level of visibility needed to optimize the program operations and costs.

Handle data at scale and improve security

There are millions of customers applying for benefits every year, resulting in tens of millions of documents and hundreds of terabytes of data. For example, Healthcare.gov alone handled data for approximately 12 million consumers during the 2021 open enrollment period. (Source: <https://www.cms.gov/files/document/health-insurance-exchanges-2021-open-enrollment-report-final.pdf>)

This data not only comes in from the applicants, but is also exchanged with federal, state, and local agencies, as they communicate with each other to validate the application. Hosting a storage solution at scale is a major challenge for many agencies. Data security and privacy are also of primary concern, as agencies have a responsibility to protect sensitive customer data.

Handle large call volumes

Handling large call volumes is a major challenge for various agencies as the follow up for each application results in multiple calls to the agency on activities related to pending applications.

Better digital experiences and customer collaboration can eliminate the need for repeated calls about basic questions (the most common one being “where are my benefits?”) so that agents can spend more time on interviews, appeals, or more complex benefits related questions. Another challenge is to ramp up call center operations during disasters, such as COVID-19, which impacted the US workforce and resulted in unemployment insurance-related calls and applications.

Reduce costs and minimize TCO

Legacy systems and software for IES require a significant maintenance investment, with increasing annual costs. These systems require multi-year enterprise licensing agreements for proprietary hardware or software and the underlying infrastructure becomes harder and more expensive to maintain as it ages. If they are built on obsolete technologies, legacy system support requires a

specific (and sometimes hard to find) set of skills and expertise. As the developers who built the software retire or switch to other technologies, it becomes increasingly harder to find and retain the right talent. Cloud migration will help address some of these challenges and minimize total cost of ownership (TCO).

Modernization focus areas for IES benefits administrators

Based on the challenges described earlier, the following figure summarizes how you can utilize cloud capabilities to modernize an IES and outlines major focus areas for benefits administrators.

The slide content is as follows:

- Improve Beneficiary Experience**
 - Digital experience (Mobile app, submission / mobile first design)
 - Automated self service via chatbots; online interview scheduling
 - Automated identity and financials verification
 - Effective communication on case status, missing information (e.g. SMS updates)
- Modernize Contact Centers**
 - Provide Omni-channel contact centers at scale
 - Customer sentiment analysis, voice translation/transcriptions
 - Intelligent operations to improve agent productivity
- Enhance Case Manager Productivity**
 - Automated data extraction process
 - Augmented AI to help case review / adjudication
 - Machine learning models to flag data anomalies
- Provide high visibility into Program Operations**
 - Enable data driven decision making
 - AI/ML for forecasting budgets, enrollment models etc.
 - AI/ML models to detect & minimize fraud, waste and abuse
- Modernize Systems and Applications**
 - Legacy to Cloud migration / Cost optimization
 - Improve system security, scalability and reliability
 - Data lakes and large scale storage
 - Develop systems with high flexibility (e.g. microservices) to allow policy and program changes

Focus areas for IES and benefits administrators – including a better beneficiary experience

IES – High-level cloud migration strategy

A well-planned strategy includes readiness, both in terms of technology and culture. The idea of migrating IES to the cloud might seem overwhelming due to the complexity of the systems and migration process. AWS provides a structured approach including architecture best practices, a partner community, professional services, and training and support to help you through the migration and modernization journey.

AWS has helped thousands of organizations migrate to the cloud and free up resources by lowering IT costs while improving productivity, operational resiliency, and business agility. Refer to our case studies on [Health and Human Services](#) (HHS) agencies including [Maryland DHS](#) and [Centers for Medicare & Medicaid Services](#) (or CMS, which runs the healthcare.gov web site) are using AWS to run mission-critical applications that administer healthcare and social benefits programs for millions of beneficiaries.

You can take advantage of the [AWS Migration Acceleration Program](#) (MAP) to help structure the lifecycle of your cloud migration. The [AWS Cloud Adoption Framework](#) can help you focus on business and technical perspectives to achieve a successful migration outcome. The migration strategy includes three main phases: assess, mobilize, and migrate and modernize.

Assess

In this phase, you review your organization's current readiness for operating in the cloud, identifying the business outcomes and developing the business case for migration. You can start by using the [Migration Evaluator](#), which provides a TCO projection for AWS, based on your actual utilization of resources together with optimizing the compute, storage, database, networking, and software licenses on AWS.

This can help your agency move from assumptions regarding the cost savings of cloud to an actual forecast specific to your infrastructure needs. You can develop the business case using the [AWS Cloud Economics Center](#), a proven methodology based on hundreds of customer engagements that delivers a compelling, board-level business case and justification prior to embarking on a large-scale migration.

After you have your business case, review the migration and modernization strategies, guides, and patterns provided through [AWS Prescriptive Guidance](#). This provides guidelines and best practices to navigate the cloud landscape, plan your migration journey, and accelerate your transformation.

Mobilize

In this phase, you create a migration plan and refine your business case. You can address gaps in your organization's readiness that were uncovered in the assess phase, with a focus on building your baseline environment (the landing zone), driving operational readiness, and developing cloud skills. One critical aspect of developing your migration strategy is to collect application portfolio data and rationalize applications using the seven common migration strategies: relocate, re-host, re-platform, refactor, repurchase, retire, or retain. You can use the [AWS Application Discovery Service](#) to automatically collect and present detailed information about application dependencies and utilization so that you make more informed decisions on your migration. The [AWS Control Tower](#) automates the setup of a baseline environment for running secure and scalable workloads in AWS based on prescriptive best practices that enable guardrails for security, operations, and compliance. If you need to create a more customized baseline, you can use [AWS Landing Zone](#).


Migrate and modernize

The final phase involves application design, migration, and validation. The [AWS Migration Hub](#) allows you to get progress updates across all of your migrations, easily identify and troubleshoot any issues, and reduce the overall time and effort spent on your migration projects.

Due to the nature of legacy architectures for IES applications, the approach we recommend is to rapidly move to cloud and then re-architect the systems in AWS. You can use [AWS Application Migration Service](#) to quickly lift and shift (re-host) a large number of servers from physical, virtual, or cloud infrastructure to AWS. Application Migration Service automatically converts your source servers to run natively on AWS, and simplifies your migration by allowing you to use the same automated process for a wide range of applications. Refer to the [How to migrate on-premises workloads with AWS Application Migration Service](#) blog post for a walkthrough of the migration.

In cases where you can't install an agent-based migration service on your server, [AWS Server Migration Service](#) is an agentless service that makes it easier and faster for you to migrate thousands of on-premises workloads to AWS from a snapshot of the existing server. If you have VMware Cloud Foundation-based environments, [VMware Cloud on AWS](#) allows you to quickly relocate hundreds of applications virtualized on vSphere to the AWS Cloud in just days, while maintaining consistent operations with your on-premises environments.

Refer to the [cloud migration and modernization blog](#) for further tips and best practices. The migration and modernization strategy is summarized in the following figure.



High-level framework for IE cloud migration and modernization

IES – Legacy patterns and migration approaches

There are a number of common legacy architecture and design patterns for current IES that are in on-premises data centers. Some of these are outlined as follows.

Note

This is not an exhaustive list; it is only a representation of some common patterns.

- Application architectures: Three-Tier / N-Tier web applications, mainframes / IBM Curam-based architectures; Java or .NET-based applications
- Operating systems: Linux (RedHat or others), UNIX, Windows, VMware-based workloads
- Middleware: Enterprise Service Bus (ESB), IBM MQ
- Databases: Oracle, Microsoft SQL Server, IBM DB2, Open Source database (such as MySQL or PostgreSQL)
- A combination of the previous technologies, tools, and operating systems, and databases

Migration options

There are seven standard patterns for cloud migration. These include: relocate, re-host, re-platform, refactor, repurchase, retire, and retain. Depending on the state of your workload and application rationalization, you can choose any of these options for your IES workload. This section focuses on a few applicable patterns for IES based on the application architecture and operating systems.

Legacy system migration

Some agencies possess legacy non-x86 systems in their data centers: mainframe or legacy UNIX systems. Short-term migration approaches mainly use hardware emulation, middleware emulation, automated refactoring, or middleware re-platforming. Refer to the [Demystifying Legacy Migration Options to the AWS Cloud](#) blog post (outlines the approaches for migrating legacy systems to the cloud) and [mainframe migration](#) page, for additional resources and guidance.

Oracle or SQL Server database migration

Many IES use commercial databases including Oracle or SQL Server. There are a few options for migration including re-hosting Oracle on Amazon Elastic Compute Cloud (Amazon EC2), re-platform with Oracle on Amazon Relational Database Service (Amazon RDS), or re-architect to use [Amazon Aurora](#). These migration patterns apply to SQL Server as well. These options also include bringing your own licenses to the cloud. Refer to the [Oracle migration whitepaper](#) and [SQL Server whitepaper](#) that outline the migration options and best practices for hosting these databases on AWS.

Middleware

You can implement enterprise application integration patterns using the [AWS messaging](#) services. Refer to the [blog post](#) for these patterns and options. If you need to implement the traditional ESB in the cloud, [AWS Marketplace](#) has multiple options.

Application migration

You can migrate Java or .NET based applications using the lift and shift approach or containerize existing applications running on-premises without code changes. Refer to the *Application Modernization* section for application modernization approaches.

VMWare based workloads

[VMware Cloud on AWS](#) provides an elastic and scalable solution to deploy VMware's Software-Defined Data Center (SDDC) and consume vSphere workloads on the AWS global infrastructure.

Red Hat OpenShift based workloads

[Red Hat OpenShift Service on AWS](#) (ROSA) provides an integrated experience to use OpenShift. ROSA delivers the production-ready OpenShift that many enterprises already use on-premises, simplifying the ability to shift workloads to the AWS Cloud as business needs change. You can accelerate your application development process by using familiar OpenShift APIs and tools for deployments on AWS. With ROSA, you can use the wide range of AWS compute, database, analytics, machine learning (ML), networking, mobile, and other services to build secure and scalable applications in the cloud.

Migration implementation approach

Implementing migration at scale requires expertise, tools, and alignment of business and IT strategy. Many organizations can accelerate their migration and time to results through partnership. You can take advantage of the [AWS Partner network solutions](#) for cloud migration. Important aspects of cloud migration implementation are outlined in the following sections.

Data migration

Developing a comprehensive data strategy is critical for successful application deployments, analytics workflows, and ML innovations. IES have eligibility, enrollment and other user data. AWS offers a portfolio of data transfer services to provide the right solution for data migration depending on the use cases, types of data you are moving, and the network resources available, among other considerations. The options include hybrid cloud storage, online data transfer, off-line data transfer, or data transfer using migration tools.

If you need to take advantage of the benefits of cloud storage but have applications running on-premises that require low-latency access to their data, or need rapid data transfer to the cloud, you can use hybrid cloud storage architectures to connect your on-premises applications and systems to cloud storage. The [AWS Storage Gateway](#) lets you seamlessly connect and extend your on-premises applications to AWS storage.

For online data transfer, you can utilize the [AWS DataSync](#) or [AWS Transfer Family](#). DataSync is a data transfer service that makes it simple for you to automate moving data between on-premises storage and Amazon S3, Amazon Elastic File System (Amazon EFS), or Amazon FSx for Windows File Server. [Transfer Family](#) provides fully managed support for file transfers directly into and out of Amazon S3 using File Transfer Protocol (FTP), Secure FTP, and FTP over SSL (FTPS).

For offline data transfer, there are a number of options including: [AWS Snowcone](#), [AWS Snowball](#), and [AWS Snowmobile](#). Snowcone is the smallest member of the AWS Snow Family of edge computing and data transfer devices; it's portable, rugged, and secure. You can use Snowcone to collect, process, and move data to AWS offline by shipping the device. Snowball is a petabyte-scale data transport and edge computing device that comes with on-board storage and compute capabilities. Snowmobile is an exabyte-scale data transport solution that uses a secure shipping container to transfer large amounts of data into and out of AWS.

Database migration

You can use the [AWS Database Migration Service](#) (AWS DMS) to migrate databases quickly and securely. The source database remains fully operational during the migration, minimizing downtime to applications that rely on the database. AWS DMS supports homogeneous migrations (such as Oracle to Oracle) as well as heterogeneous migrations between different database platforms (such as Oracle or Microsoft SQL Server to Aurora).

Additionally, AWS DMS helps consolidate multiple source databases into a single target database. This can be done for homogeneous and heterogeneous migrations, and you can use this feature with all supported database engines.

Application and server migration

The Application Migration Service is the primary migration service recommended for lift-and-shift migrations to AWS. This service helps you migrate your applications from physical infrastructure, VMware vSphere, Microsoft Hyper-V, Amazon EC2, and Amazon Virtual Private Cloud (Amazon VPC).

If you need an agentless service, you can use the [AWS Server Migration Service](#) for server migration. The AWS Migration Hub provides a single location to track the progress of application migrations across multiple AWS and partner solutions.

License migration

[AWS License Manager](#) provides a mechanism to allocate licenses to both your on-premises and cloud-hosted resources and allows you to automate the tracking of how these licenses are consumed. Refer to the [Manage license compliance during migration of workloads to AWS](#) blog post for more information.

IES modernization – User perspective

The key personas for IES include:

- The end user or beneficiary
- Agency workforce and case worker
- Program leadership and benefits administrators

This section outlines the modernization strategy with a focus on these personas.

Beneficiary and end user experience

Agencies can enhance the beneficiary experience by providing digital capabilities for a complete application process including submission, document processing, status tracking and reporting, interviews, and approval. The modernization strategy includes the following activities:

Deploy scalable omni-channel contact centers

[Amazon Connect](#) is a fast to deploy, easy-to-use omni-channel cloud contact center that helps companies provide superior customer service at a lower cost. Amazon Connect provides a seamless experience across voice and chat for your customers and agents. This includes tools for skills-based routing, powerful real-time and historical analytics, and easy-to-use management tools. Amazon Connect simplifies contact center operations, improves agent efficiency, and lowers operational costs. Additionally, Amazon Connect provides integrations for leading customer relationship management offerings, such as Salesforce.

Provide artificial intelligence (AI) powered speech analytics

[Amazon Contact Lens](#) is a set of machine learning (ML) capabilities integrated into [Amazon Connect](#). With Contact Lens for Amazon Connect, contact center supervisors can better understand the sentiment, trends, and compliance risks of customer conversations to effectively train agents, replicate successful interactions, and identify crucial feedback on benefits and claimant services. Additionally, [Amazon Transcribe](#) and [Amazon Transcribe Medical](#) provide speech-to-text capabilities. Recorded speeches can be converted to text and analyzed for further insights and opportunities for automation.

Develop self-service capabilities

Self-service capabilities will help minimize the burden on the agencies to provide large call center staff by minimizing the number of calls that are serviced by the agents. [Amazon Lex](#) is a service for building conversational interfaces into any application using voice and text. Amazon Lex provides the advanced deep learning functionalities of automatic speech recognition for converting speech to text, and natural language understanding to recognize the intent of the text, to enable you to build applications with highly engaging user experiences and lifelike conversational interactions. With Amazon Lex, the same deep learning technologies that power Amazon Alexa are now available, enabling you to quickly and easily build sophisticated, natural language, conversational bots ([chatbots](#)). You can integrate these chatbots with Amazon Connect call flows and provide self-service capabilities to customers so that agents' time can be optimized.

You can provide intelligent search capabilities with [Amazon Kendra](#), powered by ML. Amazon Kendra reimagines enterprise search for your IES and applications so that customers can easily find the content they are looking for, even when it's scattered across multiple locations and content repositories within your agency.

Provide language translation capabilities

You can use [Amazon Translate](#) to convert text from one language to another (for example, Spanish to English). Together, Amazon Transcribe and Amazon Translate can first transcribe calls in one language and then be translated into a different language.

Build effective campaign management strategies

Campaign management and effective communication with the public is critical during an emergency or during an open enrollment for benefits. [Amazon Pinpoint](#) helps the agencies engage with the public by sending them personalized, timely, and relevant communications through email, SMS text, and voice. These strategies can also help beneficiaries to act on important notices and forms, which is especially critical if they have recently moved but not updated their mailing address.

Case worker and manager productivity

An ideal way to modernize the IES and improve workforce productivity is to introduce AI/ML into the application process, augmenting the human workforce with process automation and having humans intervene only as needed. Specifically, there are three main areas of focus.

Build review and approval work-flow automation

[Amazon Augmented AI](#) (Amazon A2I) makes it simple to build process automation and workflows required for human review of ML predictions. Amazon A2I brings human review to all developers, removing the undifferentiated heavy lifting associated with building human review systems or managing large numbers of human reviewers.

Many ML applications require human review of low-confidence predictions to ensure the results are correct. For example, extracting information from scanned application forms for healthcare, unemployment, SNAP, or other benefits programs can require human review due to issues such as low-quality scans or poor handwriting. Amazon A2I provides built-in human review workflows for common ML use cases, such as text extraction from documents and audit of predictions. Using this service, you can allow human reviewers to step in when a model is unable to make a high-confidence prediction or to audit its predictions on an ongoing basis.

For example, there could be an automated application review and adjudication process based on a certain confidence level. AWS customers are implementing A2I with Amazon Textract to improve the efficiency of their document processing by combining the speed, efficiency, and cost savings of ML with A2I to include human-in-the-loop validation of results.

Use technology to minimize errors

Many of the federal programs administered by IE agencies have requirements for quality assurance and error detection, with error rates published by their oversight agencies. AI/ML and other technologies can help managers and workers bring the error rate down by identifying common error patterns and remediation through training or data cleanup. The IES itself can use this functionality to alert workers before an error is made, which can improve federal metrics, reduce fraud, and improve outcomes for eligible families.

Build ML models to identify anomalies and fraud

Automated identification of standard compared to a high-risk application is key to improving efficiencies in the review and approval workflow. To achieve this automation, historical data might need to be used to build and train ML models.

[Amazon SageMaker](#) is a fully managed service that provides developers and data scientists with the ability to build, train, and deploy ML models quickly. SageMaker removes the heavy lifting from each step of the ML process to make it easier to develop or apply [predeveloped high quality models available through AWS Marketplace](#).

SageMaker makes it simple to deploy your trained model into production with a single click so that you can start generating predictions on claims and application data. This is not only useful in training models to identify standard errors being made in applications, but also in training models to flag suspicious or fraudulent application patterns or anomalous activities identified.

Program leadership and benefits administrators

Program leadership often faces operational challenges including budget and enrollment model forecasting and better preventing and managing incidents of fraud, waste, and abuse. IES modernization helps to optimize costs and improve operational efficiencies. The modernization enables high visibility into program operations through real-time and batch analytics using services such as [Amazon Managed Service for Apache Flink](#) and [Amazon Athena](#). Program leadership can gain deep insights through operational dashboards that can be built using [Amazon QuickSight](#), a cloud powered business intelligence service. Additionally, you can use [Amazon Forecast](#) to forecast enrollment models and budgets. Using these capabilities, agencies can reduce forecasting from months to hours. Agencies can also proactively identify fraud, waste, and abuse within the benefits programs using services such as [Amazon Fraud Detector](#), or using a number of ML models available on the [AWS Marketplace](#).

IES modernization – System and application perspective

The IES application modernization strategy establishes a holistic approach for transforming the legacy systems by:

- Implementing decomposition patterns to break-up the monolith using microservices and choosing modern enterprise integration patterns.
- Re-platforming using managed container services, and cloud-native and NoSQL databases.
- Using serverless or fully managed services where applicable.
- Using modern approaches to deliver advanced capabilities, such as data analytics and AI/ML, as a means for improving program operations.

The main functional areas for system modernization include: application modernization; database modernization; data and storage modernization; and operations and DevOps modernization. The modernization approach for each of the functional areas is outlined in the following sections.

Application modernization

As the legacy IES applications expand due to policy or program changes, it becomes difficult to modify or add functionality to them, and to maintain what becomes an ever-expanding code base. As a result, even minor changes can require lengthy regression testing, and development of new features can be slow.

Microservices architecture

Agencies can become more agile and flexible by redesigning their IE applications into microservices architecture with loosely coupled components; with this architecture, you can implement new features or defect fixes at the level of a single service and released much more rapidly.

As an example, the eligibility module can be separated from enrollment module, or from financial and payment modules and further decomposed into their own set of microservices. Each team that builds microservices can also take clear ownership of their own design, development, deployment, and operations lifecycle activities.

To achieve the loose coupling, microservices in an application must communicate with each other. You can use published APIs or asynchronous message queues to communicate between separate services.

For example, separating an entire benefits eligibility determination process into different components that are connected by message queues creates clear transaction boundaries and enables services to operate more independently.

The microservices architecture, together with well-established APIs and asynchronous communication channels, enables each service to be operated and automated independently, improving reliability, scalability, and availability for the applications. As part of implementation, you need only orchestrate the overall workflow between services, configure various timeouts, and incorporate monitoring and auditing.

You can take advantage of AWS services to implement the application modernization. For workflow automation, you can use [AWS Step Functions](#), a low-code visual workflow service used to orchestrate AWS services, automate business processes, and build serverless applications. Workflows manage failures, retries, parallelization, service integrations, and observability so developers can focus on higher-value business logic.

For messaging, you can use [Amazon Simple Queue Service](#) (Amazon SQS), a fully managed message queuing service that enables you to decouple and scale microservices, distributed systems, and serverless applications. Amazon SQS eliminates the complexity and overhead associated with managing and operating message-oriented middleware, and empowers developers to focus on differentiating work.

For API management, you can use [Amazon API Gateway](#), a fully managed service that makes it easy for developers to create, publish, maintain, monitor, and secure APIs at any scale.

Serverless and fully managed services

Agencies can reduce operational complexity by utilizing serverless technologies and fully manages services. The system administrators do not have to provision and manage servers, spend time on OS patches, or maintain unused resources to address occasional peak usage. Reliability, fault-tolerance, and scalability are largely built-in by default, which eliminates much of the design and operations time required for these aspects of the system. AWS provides serverless computing services such as [AWS Lambda](#) and [AWS Fargate](#) to enable implementation of custom code without requiring provisioning or managing servers, creating workload-aware cluster scaling logic, maintaining event integrations, or managing runtimes.

Containers

Containers provide a portable, consistent, and lightweight software environment for applications to easily run and scale anywhere. Containers are also integral to breaking down traditional

monolithic IE application architectures, and they enable a transition to microservices for easier scale.

AWS offers the widest range of services for storing, managing, and running containers. [Amazon Elastic Container Service](#) (Amazon ECS) is a fully managed container orchestration service that provides the most secure, reliable and scalable way to run containerized applications that are deeply integrated with all of the other AWS services.

[Amazon Elastic Kubernetes Service](#) (Amazon EKS) is a fully managed Kubernetes service that provides the most secure, reliable, and scalable way to run containerized applications using Kubernetes. AWS Fargate is a serverless compute engine for containers that works with both Amazon ECS and Amazon EKS. Fargate removes the need to provision and manage servers, lets you specify and pay for resources per application, and improves security through application isolation by design.

Database modernization

There are a number of challenges with legacy IES databases, including lack of scalability to meet the growing needs of benefits systems, high costs, complexity of database administration, one-size-fits-all databases, and inefficient and complex hardware provisioning. AWS provides a variety of purpose-built, fit-for-use databases that are highly scalable, reliable, fully managed, and serverless.

For example, [Amazon DynamoDB](#) is a key-value and document database that delivers single-digit millisecond performance at any scale. It is a fully managed database with built-in security, backup and restore, and in-memory caching for internet-scale applications. DynamoDB can handle trillions of requests per day and easily support over double-digit millions of requests per second across its entire backplane.

[Amazon Aurora](#) is a MySQL and PostgreSQL-compatible relational database built for the cloud that combines the performance and availability of traditional enterprise databases with the simplicity and cost-effectiveness of open-source databases. Aurora is up to five times faster than standard MySQL databases and three times faster than standard PostgreSQL databases. It provides the security, availability, and reliability of commercial databases at 1/10th the cost. For migration and modernization from Oracle to Aurora, refer to the [Migration Playbook](#).

[Amazon RDS](#) makes it simple to set up, operate, and scale a relational database in the cloud. It provides cost-efficient and resizable capacity while automating time-consuming administration tasks such as hardware provisioning, database setup, patching, and backups. It frees you to focus

on your applications so you can give them the fast performance, high availability, security and compatibility they need. Amazon RDS is available on several [database instance types](#)—optimized for memory, performance, or I/O—and provides you with six familiar database engines to choose from, including: [PostgreSQL](#), [MySQL](#), [MariaDB](#), [Oracle Database](#), [Aurora](#), and [SQL Server](#). You can use the [AWS Database Migration Service](#) to easily migrate or replicate your existing databases to Amazon RDS.

[Amazon Redshift](#) allows you to query and combine exabytes of structured and semi-structured data across your data warehouse, operational database, and data lake using standard SQL.

Redshift lets you save the results of your queries back to your Amazon S3 data lake using open formats, like Apache Parquet, so that you can do additional analytics from other analytics services like Amazon EMR, Athena, and SageMaker.

For a detailed approach on database modernization, please refer to the whitepaper [Modernizing the Amazon Database Infrastructure](#).

Data and storage modernization

Legacy IE storage solutions can be expensive, might not scale well to meet the growing needs, and might have durability and availability issues. AWS offers a complete range of services for you to store, access, govern, and analyze your data to reduce costs, increase agility, and accelerate innovation.

You can select from object storage, file storage, and block storage services. You can elect to backup and utilize the concept of data lakes to build the foundation of a modern IES. A data lake allows you to store all your structured and unstructured data at any scale in a centralized repository. You can store your data as-is, without having to first structure the data.

From there, you can use a variety of tools and processes to transform the data, if required, and to run different types of analytics to guide better decision-making—from dashboards and visualizations, to big data processing, real-time analytics, and ML.

Tens of thousands of customers run their data lakes on AWS. Setting up and managing data lakes previously required a lot of manual and time-consuming tasks. [AWS Lake Formation](#) automates these tasks so that you can build and secure your data lake in days instead of months.

For your data lake storage, Amazon S3 is the best place to build a data lake because of its unmatched 11 nines of durability and 99.99% availability. You can also use [Amazon HealthLake](#), a

fully managed, HIPAA-eligible service that allows healthcare customers to aggregate their health information from different silos and formats into a structured, centralized AWS data lake, and extract insights from that data with analytics and ML.

When the data in your data lakes and purpose-built data stores continues to grow, you need to be able to easily move a portion of that data from one data store to another.

AWS allows you to combine, move, and replicate data across multiple data stores and your data lake. For example, [AWS Glue](#) provides comprehensive data integration capabilities that make it easy to discover, prepare, and combine data for analytics, ML, and application development, while Amazon Redshift can easily query data in your S3 data lake. For further details, refer to the blog post [Design patterns for an enterprise data lake using AWS Lake Formation cross-account access](#).

Operations and DevOps modernization

Legacy IES, in many cases, have development and operations practices that lead to longer resolution time for operational events and extended software release cycles. Development and operations agility can be improved with streamlined DevOps processes and by incorporating monitoring, logging, and auditing into the operations. These can be achieved through the processes outlined in the following sections.

Automating deployment with continuous integration and continuous delivery

Agencies providing benefits systems strive to deliver the most value they can to citizens as quickly as possible. To achieve this, modern applications use continuous integration and continuous delivery (CI/CD) to automate the entire release process: automated build processes, automated system and integration testing, promoting artifacts to staging, and the final deployment to production.

CI/CD can also automate certain security controls, such as scanning for known vulnerabilities and performing static analysis. You can take advantage of fully managed CI/CD services such as [AWS CodeBuild](#), [AWS CodePipeline](#), and [AWS CodeDeploy](#), in addition to open-source options and third-party marketplace offerings

Managing infrastructure as code

By modelling applications and infrastructure as code (IaC), you can achieve the full benefits of CI/CD. You can incorporate this model into your application development lifecycle, implement

infrastructure changes in your CI/CD pipeline, and get additional benefits, such as reducing configuration errors and faster provisioning.

AWS provides a number of IaC services. [AWS CloudFormation](#), for example, allows you to model a collection of related AWS and third-party resources, provision them quickly and consistently, and manage them throughout their lifecycles, by treating IaC. The [AWS Serverless Application Model \(AWS SAM\)](#) is an open-source framework for building serverless applications.

It provides shorthand syntax to express functions, APIs, databases, and event source mappings. With just a few lines per resource, you can define the application you want and model it using YAML. The [AWS Cloud Development Kit \(AWS CDK\)](#) (AWS CDK) is an open source software development framework to define your cloud application resources using familiar programming languages.

Monitoring and logging

Improving beneficiary experience is a major goal for benefits administrators. In service of this goal, operational teams should monitor the behavior of their applications at runtime using monitoring and logging services, and use that data to improve the beneficiary experience.

This can include business and user level metrics (such as benefits status inquiry response time), system level metrics (such as latency), and service level metrics (such as availability). AWS offers a number of services for monitoring and logging.

[Amazon CloudWatch](#) is a monitoring and observability service built for DevOps engineers, developers, site reliability engineers, and IT managers. CloudWatch provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, optimize resource utilization, and get a unified view of operational health.


[CloudWatch Logs](#) enables you to centralize the logs from all of your systems, applications, and AWS services that you use, in a single, highly scalable service. You can then view them, search them for specific error codes or patterns, filter them based on specific fields, or archive them securely for future analysis. This capability can simplify and streamline the various audit processes required by federal oversight agencies, further improving program compliance. Refer to [Centralized Logging](#) for implementing a centralized logging solution on AWS.

[AWS X-Ray](#) helps developers analyze and debug production, distributed applications, such as those that are built using a microservices architecture. With X-Ray, you can understand how your application and its underlying services are performing to identify and troubleshoot the root cause of performance issues and errors.

IES migration and modernization summary

The following figure provides a roadmap for IES migration and modernization and summarizes the steps in the process.


This can be a multi-year undertaking depending on the agility of the agency. As mentioned earlier, the [Cloud Adoption Framework](#) can help you focus on business and technical perspectives to achieve successful migration and modernization outcomes.



IES cloud migration and modernization roadmap

Example – Windows application modernization

The following figure provides a reference migration and modernization approach for a Windows-based application.




Cloud migration and modernization approach for Windows-based application

IES reference architecture summary

The following figure provides a reference architecture for a modernized IES. This includes migration to the cloud and enabling better resiliency, scalability, availability, and security.

The modernization can help break the monolithic applications into a set of microservices and containers together with fully managed services. These capabilities help improve both user experience and case worker productivity by providing flexible web interfaces and API gateways, cloud-based contact centers, self-service chatbots, data lakes, and AI/ML capabilities.



IE cloud migration and modernization reference architecture

Conclusion

Health and human services agencies face a number of challenges with legacy integrated eligibility and enrollment systems. This whitepaper outlined some of the challenges, and introduced migration and modernization strategies that will help address some of them. This whitepaper also outlined the three personas—end user, case worker, and the program leadership—that would benefit from IES migration and modernization.

Adopting the AWS Cloud can provide you with sustainable advantages for IES. As a next step on your journey to the cloud, you can develop a business case for cloud migration and modernization. AWS works together with your team and your chosen member of the AWS Partner Network to implement your enterprise cloud computing initiatives. You can reach out to an AWS Partner through the [AWS Partner Network](#). You can get started by visiting the [AWS Health and Human Services](#) webpage.

Contributors

Contributors to this document include:

- Srinath Godavarthi, Principal Solutions Architect, Health and Human Services
- John Velasquez, Technical Program Manager, Health and Human Services
- Trey Long, Business Development Manager, Health and Human Services
- Justin Ewald, Sr. Solutions Architect, Health and Human Services

Further reading

For additional information, refer to:

- [AWS Health and Human Services](#)
- [AWS state and local](#)
- [Cloud migration with AWS](#)

Document history

To be notified about updates to this whitepaper, subscribe to the RSS feed.

Change	Description	Date
Initial publication	Whitepaper first published.	November 10, 2021

Note

To subscribe to RSS updates, you must have an RSS plug-in enabled for the browser that you are using.

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.

AWS Glossary

For the latest AWS terminology, see the [AWS glossary](#) in the *AWS Glossary Reference*.