



**Amazon's Corporate IT
Migrates Business Process Management
to the Amazon Web Services Cloud**
April 2011

This paper describes how and why Amazon's corporate IT organization implemented its migration of Appian's Business Process Management Suite¹ (an enterprise-class, mission-critical financial application) to the Amazon Web Services (AWS) cloud. Throughout the process, we² engaged AWS like we would any vendor. In turn, AWS treated us as it would any enterprise customer.

This paper shares our story as an AWS reference customer of this three-month cloud migration project. Our story is broken down into the following sections:

- Where We Were
- Where We Wanted to Go and Why
- Our Application Selection Criteria
- How We Migrated
- Lessons We Learned
- What Benefits Emerged
- Where We Go From Here

Where We Were

Our company³ uses Business Process Management (BPM) software to automate a number of our mission-critical financial systems processes. These involve disparate teams across the company and include handling confidential financial data for payment, payroll, and accounting processes. Managed correctly, the automation and execution of these financial process applications helps ensure that the company maximizes profit, reduces errors, and eases compliance burdens.

Our company uses the Appian BPM Suite to automate tasks for a number of financial processes. For example, Accounts Payable processes thousands of invoices a day through various buyer departments. Quickly and favorably resolving any exceptions in these invoices has a significant positive financial impact.

In applying BPM to this process, we sought a technology platform that would provide simple power by combining extreme ease of use with comprehensive functionality. Appian's BPM Suite provides an intuitive and seamless user experience across a tightly integrated suite of components, including process design, rules, forms, content management, reporting, real-time analytics, identity management, integration, and collaboration.

¹ See <http://www.appian.com>.

² Throughout this document, the term "we" refers to Amazon Corporate IT.

³ Throughout this document, the term "our company" refers to Amazon.com.

The Appian BPM Suite system requirements are similar to those of many mission-critical platforms:

- Separate stacks for development, quality assurance, production, and failover
- Enterprise-class hardware and at least 64 GB RAM and 200 GB of disk space per stack
- Daily backups for disaster recovery
- NFS mounted file servers to support the data and application libraries (shared between the front-end Appian BPM web server and back-end Appian BPM engine)
- A link to Microsoft SQL Server databases
- A link to our financial data warehouse

Figure 1 shows our BPM architecture before migrating the Appian BPM components to the cloud.

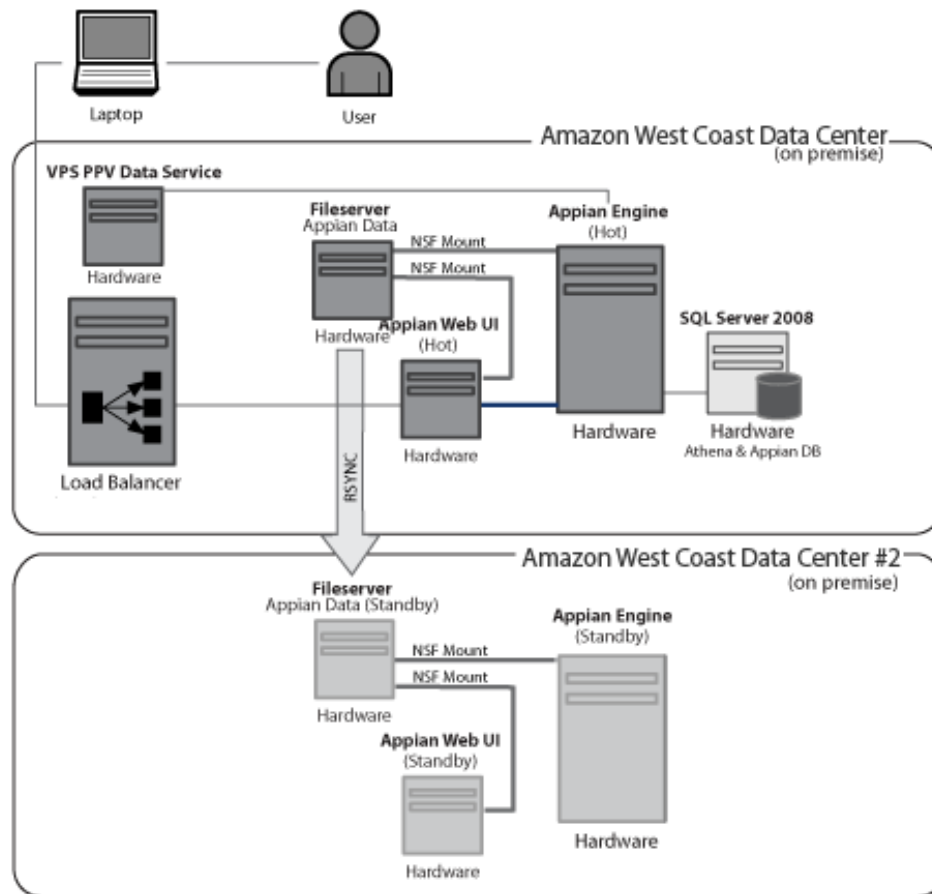


Figure 1. The Appian BPM Suite Component Breakdown before Migration

Higher-memory servers are best suited to meet these requirements. As our company's application needs have grown, so have the demands on the Appian BPM Suite infrastructure. Scaling to meet the company's increased business needs requires additional servers and more storage to accommodate the increasing amounts of process application data.

For efficiency, our company uses predefined server configurations, but because the Appian BPM engine has high memory requirements, we were overspending on disk capacity. The Appian BPM file server presented the opposite problem; it required more disk space, and we were therefore overspending on memory.

Typically, our new hardware procurement was taking four to six weeks. We had reached the point where we needed to be more agile with infrastructure scaling to keep up with the business demand. Corporate IT was looking to simplify management of the hardware infrastructure, including lease returns, provisioning new hardware, supporting data center moves, and expanding BPM to support new geocentric requirements.

Clearly, we needed to implement a better system-level solution for managing services so central to the company's business.

Where We Wanted to Go and Why

After reviewing the landscape of solution enhancements, we determined that an AWS-based infrastructure was the best-in-class option. Our reasoning included the following:

- Linux servers that typically take four to six weeks to procure, can be acquired and provisioned in minutes.
- Amazon EC2⁴ hosts provide better flexibility for scaling with our increasing BPM application system needs.
- The memory and compute requirements can be selected independently from the required Amazon Elastic Block Store (EBS)⁵ storage capacity, which can be scaled up or down as needed over time.
- Operational support could be simplified by consolidating the Appian BPM web server and engine to a single high-memory EC2 instance that would share the same file system via an attached Amazon EBS volume.
- Using a Virtual Private Cloud (VPC)⁶ keeps things inside the corporate perimeter.
- Our team could be more agile when prototyping application changes, testing patches and upgrades, and developing new BPM applications using the cloud infrastructure—all at a reduced cost and without any service disruptions to customers.
- Amazon EBS-based snapshots are more efficient, enable more frequent automated backups, and have a much quicker and automatable recovery process.
- We wanted to prove cloud resources would offer cost savings compared to our corresponding hardware total cost of ownership (TCO).
- We could shrink the on-premise footprint, reduce hardware overhead, simplify application maintenance, and decrease the infrastructure budget.

After reviewing this list of benefits for migrating to an AWS-based solution, migrating Amazon's Appian BPM applications to AWS became one of the organization's top priorities.

⁴ See <http://aws.amazon.com/ec2/>.

⁵ See <http://aws.amazon.com/ebs/>.

⁶ See <http://aws.amazon.com/vpc/>.

Our Migration Success Criteria

An enterprise cloud migration requires a significant investment of engineering time and resources. Having clearly defined success criteria ensures continued momentum when challenges arise. During our system assessment phase, we identified four main criteria as keys to success and summarized why the Appian BPM Suite met these success criteria. This gave us the justification to proceed with our cloud migration plan.

- **Strong executive commitment.** Top management must support the application as a viable candidate for migrating to the cloud. After reviewing our case, our senior executives were heavily in favor of investing in an Appian BPM cloud migration. They were passionate about the scalability of a cloud IT infrastructure and made our cloud migration program one of our organization's top priorities. Strong executive commitment ensured that our engineers stayed the course in the face of challenges. We knew that without such support, the migration would have been much more difficult to accomplish.
- **Motivated engineers.** The engineers who owned the application had to be excited about the promise of the cloud, and had to be willing to tackle challenging problems head on. Our BPM engineers were excited about the opportunity to blaze a trail for other teams while working with new technology. The engineers wanted to stop managing hardware infrastructure processes in general and start creating and customizing nimble enterprise applications geared toward our employees' needs. The team knew that the knowledge we gained deploying these applications in the cloud would help us discover and resolve perceived obstacles for future cloud deployments.
- **High cloud readiness and low migration effort.** The application had to lend itself to a cloud deployment. The Appian BPM Suite architecture is web-based and easily portable. The components of the architecture are not tied to any specialized hardware. Because Appian's BPM platform is 100 percent web-based, all application features can be accessed from anywhere using a web browser. As such, it could be moved easily from its existing data center to a VPC, preserving its downstream dependencies without affecting our users.
- **Strong vendor partnerships around cloud licensing and support.** Appian supported our migration plan, and took a true partnership approach to ensure that there would be no licensing or support issues in converting our deployment from on-premise to the cloud. Appian worked with us on a cloud-friendly license that would enable our team to cost-effectively scale using cloud resources. Appian also agreed to extend the existing support services agreement to the cloud-based infrastructure.

How We Migrated

Migration Support

The migration plan brought a number of enterprise and cloud technologies together for the first time in corporate IT. We opted to purchase AWS Premium Support so we could maximize pre-implementation support advice (solutions architects) and receive post-implementation assistance (support engineers). AWS solution architects helped us solve key challenges and share best practices from the field. AWS Premium Support was engaged as needed to help resolve issues and keep us on track. Having access to AWS Premium Support was critical to the project's success.

Initial Requirements Analysis

Working with AWS solutions architects, we studied the application dependencies and requirements. Together, we determined that the existing three-tier web server, engine, and NFS-mounted file server could each be moved from their current on-premise data center in the western United States to run on Amazon Virtual Private Cloud hosts without modification. Connections to system dependencies would need to be tested, but would likely remain unchanged. However, the team was motivated to eliminate a number of the architectural complexities that came with supporting multiple tiers, reattaching mounted NFS servers on reboot, and persisting the file systems between the active and standby stacks. We therefore decided to revisit the application design with an eye toward optimizing for the cloud.

Security Review

All applications undergo a risk assessment by our corporate information security team. They set the security bar for the Appian BPM Suite quite high, because they determined that the application would store highly confidential financial data. Because the Appian BPM data was classified *Amazon Confidential*, the requirements from our security review included the following:

- The Appian BPM Suite must be deployed within a VPC⁷ to restrict access to hosts only from our connected corporate network.
- All Appian BPM data must be encrypted, both at rest and in flight.
- All Appian BPM traffic in the VPC must be encrypted, in addition to traffic coming in and out of the VPC.
- Appian BPM encryption keys must be protected and cannot be stored in the same location as the data they are protecting.
- Appian BPM connections to SQL Server must be secured using SSL.
- Web service calls made from Appian's BPM applications must be secured over HTTPS.

These significant requirements are due to the sensitive nature of the financial data stored in our Appian BPM applications, and were taken into consideration as the team revisited the application design.

Proof of Concept

The proof of concept was an integral step in demonstrating the applicability of the AWS platform, our new cloud-based Appian BPM Suite design, and our security requirements. We were able to prototype each element independently, and tear down the environment after each proof while spending very little money.

Amazon Virtual Private Cloud

A VPC was established in the US East region connected to the corporate network in the nearby on-premise East Coast data center. To provision capacity to the VPC, we also submitted an EC2 limit increase request to change the maximum on-demand instances limit per AWS account⁸ to the maximum number needed for our organization. This initiated a

⁷ See <http://aws.amazon.com/vpc>.

⁸ Form to increase to your Amazon EC2 instance limit: <http://aws.amazon.com/contact-us/ec2-request/>.

forecasting process with AWS, who worked with us to determine the maximum number of instances we might need in the coming year, keeping in mind other migration projects we have on our roadmap.

Application Installation on EC2

The Appian BPM web server and engine were consolidated and both installed on an extra large, high-memory EC2 instance (M2.4XL). The application functioned and required no modifications, thus proving the RHEL5 EC2 image compatibility with the Appian BPM Suite.

Encryption at Rest on EBS

An Amazon EBS volume was created and attached to a fresh EC2 host. We used this stage to determine the steps needed to encrypt the volume using the `dm-crypt` module and Linux Unified Key Setup (LUKS) extension and deploy our internal key store service on the host. We then proved and documented the encryption and key management procedure.

The Appian BPM Suite was installed on the new encrypted volume. Again, the application functioned and required no modifications. We proved the application could be installed on Amazon EBS after deploying our encryption solution on the volume.

SQL Server Database Deployment

At our request, the corporate IT database administrator team deployed an encrypted Shared SQL Farm⁹ in the cloud. This provided the Appian BPM application with local secure connections to its databases, and provided a highly secure database farm for other applications that would soon require it.

The Appian BPM vendor data was then replicated to our test SQL Server instance. Using a stored procedure scheduled to run daily as a *cron* job (or scheduled task) and a linked (trusted) SQL Server connection, we kept vendor data updates in the cloud synchronized with the vendor database in the on-premise Shared SQL Server farm.

Encryption in Flight

HTTPS was used for securing browser-based client connections to the Appian BPM web server running in the cloud. New HTTPS endpoints were created to support secure web service calls between the Appian BPM engine and the Amazon data warehouse. We used test data to prove the Appian BPM workflows could use the new secure endpoints to query and write data over the secure HTTPS web service. Finally, we tested our data replication solution using `rsync` over SSH between our Amazon EBS file system and an on-premise test file server.

⁹ See whitepaper on "Amazon's Corporate IT Launches SharePoint 2010 to the Amazon Web Services Cloud" to learn about a SQL Server 2008 Enterprise cloud deployment

Application Architecture

Following the proof of concept, we agreed that the application design (see Figure 2) would accommodate both application and security requirements. This new cloud architecture was a simpler architecture with zero impact to the application configuration. Our corporate information security team approved the proposed security design.

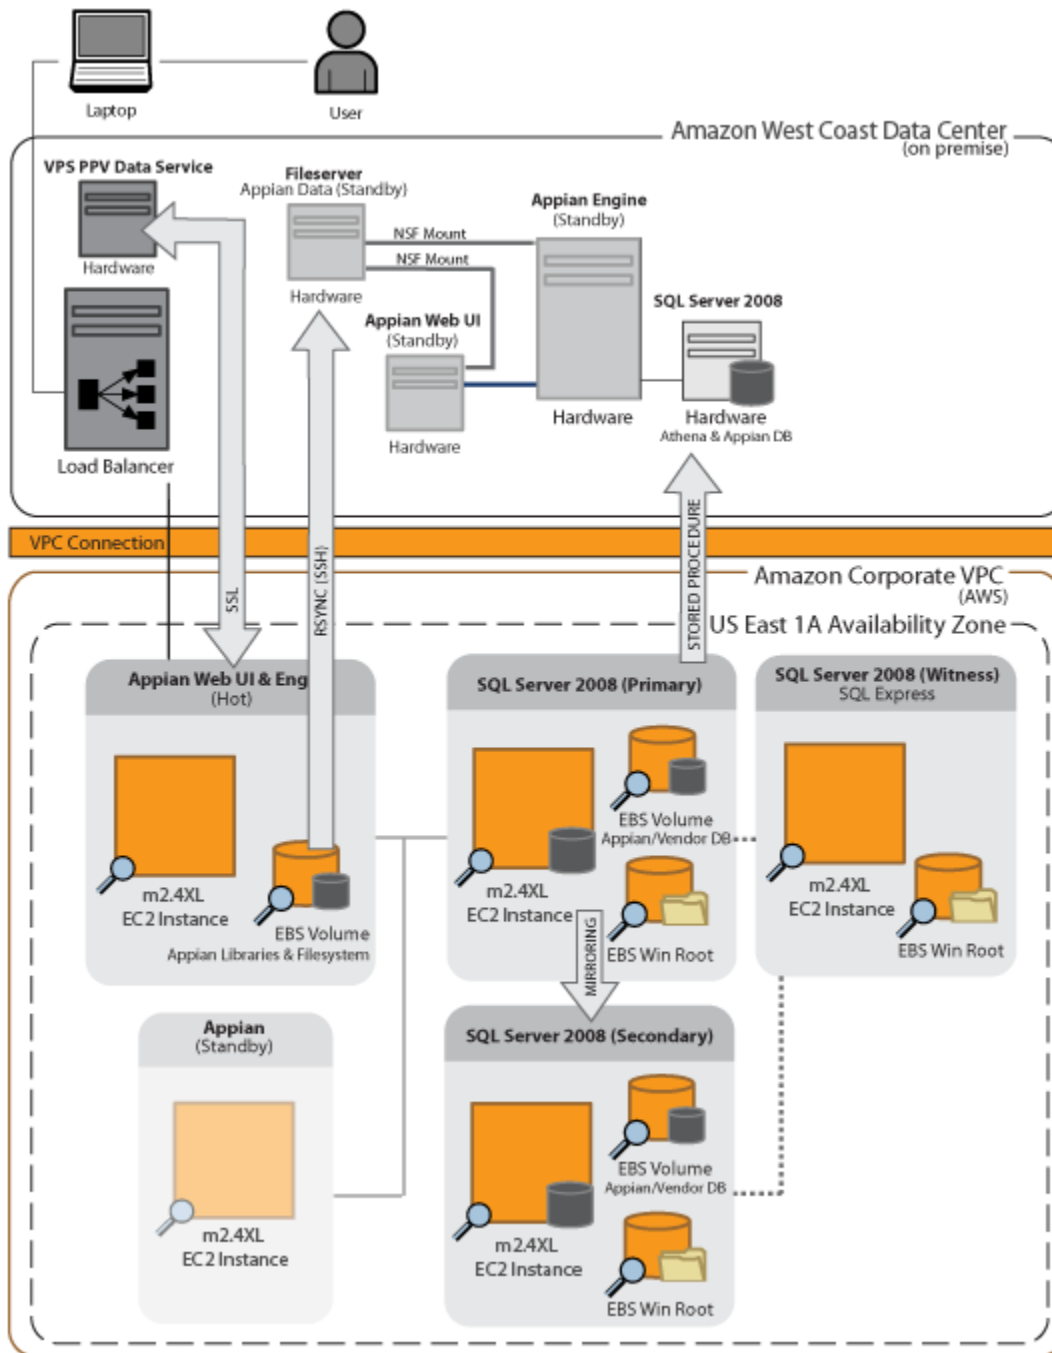


Figure 2. The Appian BPM Suite Component Breakdown after Migration

Migrating the Core Application

The migration progressed through the following stages:

- Application component migration
- Test data migration
- Application functionality testing

Application Component Migration

Custom Appian BPM components developed for specific business processes were migrated from the on-premise West Coast data center to the new cloud instance running in our VPC. These components were all stored on the same Amazon EBS volume where the Appian BPM libraries were installed.

Test Data Migration

Sample application data was replicated from the on-premise environment to the Appian BPM file system on EBS. The Appian BPM configuration information and sample vendor data were copied from the SQL Server on-premise server over to the new Amazon VPC SQL Server database.

Application Functionality Test

With application components migrated and test data in place, all features of the entire application were exercised to ensure all the key steps in the workflows worked as expected.

Preproduction Readiness

The next phase involved steps to get our capacity, availability, data redundancy, backup and recovery, and monitoring solutions in place for our production launch on the cloud.

EC2 Capacity

As we moved into production, we were in a position to pay for reserved EC2 instances.¹⁰ This ensured that we got more favorable reserved pricing for *always on* Appian BPM capacity. This pricing option for EC2 allowed us to make a low one-time payment to further reduce hourly usage charges. Reserved Instances complement our existing EC2 on-demand instances to reduce our EC2 costs.

Availability

We deployed standby EC2 hosts with preconfigured Appian BPM Suite licenses and simulated an outage by terminating the live EC2 host. Their failover solution involved salvaging the Appian BPM libraries, application components, and sample data all stored on a persisted Amazon EBS volume. We then tested unlocking the encrypted volume, attaching it to the standby host, and activating the Appian BPM application to run through the full functional test plan. The functional test plan successfully confirmed Appian's BPM application availability had been restored.

¹⁰ To learn more about Reserved Instances, go to <http://aws.amazon.com/ec2/>.

Redundancy and Backup

To create redundancy for our VPC deployment, we decided to both back up and replicate the Appian BPM file system. Our primary backup solution uses a scheduled job that creates Amazon EBS snapshots of the file system every 15 minutes. Using the Amazon EBS snapshots, we can quickly recover the Appian BPM application libraries, custom components, and all corresponding workflow data. The file system was also replicated to an on-premise standby environment to be available for disaster recovery.

Monitoring

Our existing application-monitoring mechanisms were extended to the cloud. The application metrics produce an alarm when a fatal error occurs in the log files. In addition to application monitoring, host-monitoring agents were also initiated on the active EC2 hosts, measuring the CPU, RAM and swap usage. If any metrics fall outside a defined range, the support team is notified and deployed to diagnose and resolve the problem.

Migrating the Production Data

The final step migrated the Appian BPM production data from the on-premise environment to the VPC. This required a scheduled eight-hour application freeze during a weekend when global users would not be affected. With production at a halt, the team moved the data using a secure FTP connection from the on-premise West Coast data center to our VPC and into its new file system on a production-ready Amazon EBS volume.

When the transfer was complete, the production load balancer was pointed to the Appian BPM's EC2 instance. By the next working day, the Appian BPM was taking production traffic on EC2.

Lessons We Learned

As with any major enterprise migration, we encountered a few challenges during the deployment. We will apply the knowledge gained from these lessons to drive future cloud migration projects and share them with other teams.

Licensing

Many vendor licensing agreements are not yet written to anticipate the use of a cloud-based infrastructure. This will change, but in the meantime be sure to engage your vendor early in the project, as we did with our Appian account representatives to understand how their licensing and support agreements will be interpreted on AWS. We learned that it is important to educate the vendor about the cloud architectural patterns and invent new licensing models to support the cloud. Although our Appian BPM Suite license lets us scale without incurring additional application costs, a fully qualified domain name is required for each new license. For failover protection, we must maintain a standby EC2 host and domain name to which the license can be associated. Appian is actively working on an alternative solution for this current technical licensing requirement.

Forecast VPC Resources

The goal for our corporate IT was to use a single AWS account for all EC2 capacity in our VPC. As a large enterprise customer, we therefore went through a forecasting process with AWS to fulfill our VPC capacity requirements. This process began with a request to increase the EC2 limit for our AWS account¹¹. We then worked with AWS sales and solution architects to determine how many reserved vs. on-demand instances we would need over three years.

Use Standard Scalability and High Availability Practices

Corporate IT best practices state that we should deploy applications to at least two data centers for high availability. We therefore deployed our active Appian BPM application in the VPC, while maintaining our standby application stack in an on-premise data center to protect against an Availability Zone or VPC outage. We leveraged our existing on-premise load balancers to direct traffic to the correct Appian BPM web server. Should an outage of our VPC or Availability Zone occur, all traffic will be directed to the on-premise web front-end server by updating our internal load balancing infrastructure. In the future, we look forward to expanding our cloud footprint to a second Availability Zone in our VPC, which would allow us to decommission the on-premise servers and use virtual load balancers in our VPC to direct traffic.

Plan an Access Control Strategy for Cloud Resources

With the power and flexibility of having on-demand hosts comes ability to terminate hosts just as easily. We needed to think through the processes and controls needed to protect these resources from accidental deletion or misidentification. Our approach relied on the use of AWS Identity and Access Management (IAM)¹² policies and tagging for resource identification. We assigned AWS user policies to engineers across teams where limiting access to various AWS APIs helped us avoid accidental deletions. Defining a tagging convention also helped us to easily identify EC2 and Amazon EBS resources. We gave AWS feedback on how resource-level access controls could simplify and improve our IAM strategy.

Security Is a Shared Responsibility

Our corporate information security team held AWS to the same standards as any external vendor. AWS security features (especially those relating to VPC) protect instances, data volumes, and network traffic, but the team learned that it remains their responsibility to protect the corporate data that is stored and running in the cloud. Although the VPC controls traffic to and from corporate IT's connected network, the cloud still runs on a shared network. The corporate information security team therefore required us to encrypt all data at rest and in flight. From these requirements, the team learned about the architectural challenges and operational overhead that comes with supporting a robust data encryption and key management solution. We overcame these challenges by investing the time and resources needed to incorporate these critical security components into our cloud deployment. These are now core framework components which we will use in other similar migration projects.

¹¹ Form to increase to your AWS EC2 instance limit: <http://aws.amazon.com/contact-us/ec2-request/>.

¹² See <http://aws.amazon.com/documentation/iam/>.

Get Customer Buy-In Early

We involved AWS early in the process of getting buy-in from our Appian BPM stakeholders and users. It was important for our customers to hear from AWS, as they would from any cloud vendor, that architecting availability into the AWS infrastructure was no different than its on-premise equivalent. After our Appian BPM stakeholders and customers understood how to leverage the cloud to architect a highly available solution, they supported and anticipated our decision to migrate.

What Benefits Emerged

Migration of our Appian BPM application to the AWS cloud was completed quickly and efficiently, with no disruption to application end-users. The move has extended the BPM application's overall business value. Corporate IT is now freed from the cost and time required to maintain mission-critical financial process applications, while data security and application reliability are ensured through the proven AWS cloud infrastructure.

Specific benefits include the following:

- **Infrastructure Procurement Time Reduced** from four - six weeks down to minutes
- **Eliminating Lease Return Process** freeing 16 labor hours per server
- **Risk of Data Loss Reduced** by doubling data replication frequency using Amazon EBS snapshots
- **100 percent Data Persistence** on Amazon EBS that can be attached to a new instance within minutes of an EC2 failure
- **Annual Infrastructure Costs Cut** by 86 percent for the active Appian BPM Suite application stack

Where We Go From Here

With Appian BPM Suite now running on AWS in a VPC, and key lessons learned in the process, we are sharing our story with internal peers and external customers. We solved a number of key availability, performance, and security concerns that can be leveraged within the company and drive other internal teams to adopt the cloud. This will build visibility and separate the *truths* from *myths* that are prevalent within many enterprise IT organizations.

We are now in a position to make more cloud optimizations. We plan to reduce operational overhead even further by leveraging AWS for automating failover and recovery. We are exploring how to expand our cloud footprint into other geographies to meet stakeholder demand for an Appian BPM presence in Europe. Lastly, we plan to decommission all on-premise standby hardware in favor of a fully redundant and highly available Appian BPM presence on the cloud.