



## ***Cloud Center of Excellence***

### ***VSO Approach and Customer Case Study***

***Date: January 27, 2020***

*Prepared by:*

**VSO, LLC**

380 Broadview Ave

Warrenton, VA 20186

---

# Contents

- Executive Summary..... 3
  - VSO SQOD Approach..... 4
  - VSO Veteran Services Approach ..... 4
  - VSO Approach to a Customer’s Journey to Cloud ..... 5
  - VSO Design Thinking – The Start of the CCoE..... 6
  - Understanding Agile Concepts..... 7
- Customer Case Study ..... 9
  - MVP Approach Detail..... 10
- Phase 1 - Customer Design to Agile Backlog..... 12
- Official Customer Enterprise Cloud Strategy ..... 12
  - Identity and Data Security ..... 13
  - Resource Tagging ..... 17
  - Collaboration..... 17
  - Data Lake..... 20
  - Qlik ..... 22
- Customer Platform Governance Model..... 23
- Customer Platform Roadmap ..... 24
  - Data Catalog and Lake ..... 24
  - Add Remaining Customer Entities ..... 25
  - Moving to a Highly Collaborative Model ..... 25
  - Customer Enterprise Knowledgebase..... 28
- Case Study Conclusion ..... 29

---

## Executive Summary

The VSO Cloud Center of Excellence (CCoE) focuses on the Customer's iconic user experiences and builds a framework that aligns technology, Agile, DevOps and governance to deliver these to the business. We begin with Design Thinking, then help customers modernize and migrate, with US Veteran-based operations and performance-based feedback.

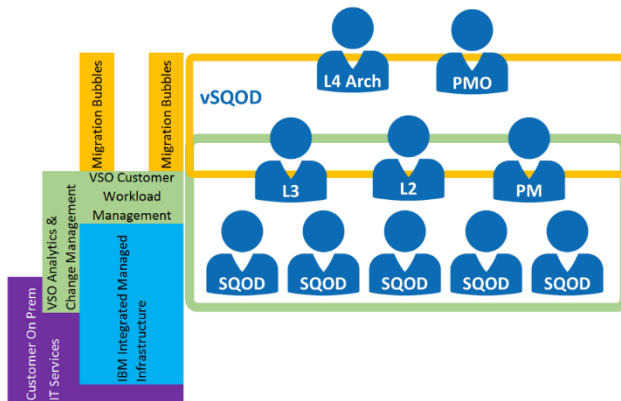
We offer several services within the CCoE and can tailor what we deliver to customer need:



- Begin with Design Thinking Workshops to define "Hills" of success with the customer. We then prioritize and build a roadmap with the customer, including capturing environmental details and building a backlog
- Run projects in Jira, breaking Hills into epics, stories and tasks, typically in a three-month release schedule of two-week sprints
- Perform discovery and use a hybrid service catalog to assess an application's suitability to be moved/modernized
- Build a Design Authority to foster and promote technical vitality for our customers
- Build an Architecture Control Board to build/review/refine architectures, solutions, standard operating procedures, pre-go-live checklists, etc.
- Define and implement best practice automation and development tools to accelerate deployment for our customers
- Lead steady state programs, such as monitoring, infrastructure code life cycle and IAC governance
- Develop client service offerings, such as VSO's FogLifter® and Managed Services
- Stand up DevOps CI/CD pipelines and Automation practices for customers
- Measure performance and provide feedback to the business via our Theta Scoring System

## VSO SQOD Approach

VSO is organized to provide services around supportable team units called Service Quality Operations Divisions, or “SQODs”.



The VSO team begins with L4 Architecture and Engineering. VSO as an organizing principle will not attempt to deliver a service for which they do not employ a leading industry expert. These Level “4” architects ensure that the very best practices are pushed down through the organization, and that every efficiency of automation and tooling is taken full

advantage of. These Level 4 architects are paired with very senior Program Managers to form the backbone of the corporation. The individual account SQODs are backed up by leveraged teams of these “Level 4” Architects and Program Managers-whose involvement scales based on the size, duration, and complexity of the migration bubbles.

## VSO Veteran Services Approach

Virtual Service Operations, or VSO, has developed a methodology to train and deploy a veteran based workforce to provide architecture, migration, and managed services. Teaming with Paralyzed Veterans of America, the US



**Paralyzed Veterans  
of America**

Department of Labor’s Apprenticeship USA program, Amazon Web Services and the Microsoft Software and Systems Academy, VSO’s leadership has successfully trained veterans to transform and manage some of the largest IT enterprise in the world including Healthcare.gov, R1, SAP Cloud, Howard University, and Kaiser Permanente.

Most veteran placement programs fail because vets are brought into a company in the same way as any other hire. Particularly, in fields where the learning cycle is more technical, vets benefit from a team environment. For this reason, VSO always recruits vets in groups. Mapping the Veteran’s expertise to civilian skills is also critical to the success of VSO’s model. VSO has seen 98% retention and success rate as well as the highest level of customer satisfaction.

VSO has developed a staffing approach that focuses on hiring veterans-particularly disabled



vets-and providing them an environment of training, automation, and “Level 4” engineering mentorship that ensures they are effective very quickly.

Training starts at the basic transitional business skills one learns at the start of a career, and then through intensely documented technical training materials, they are able to begin performing Service Assurance processes and rapidly move on to deep Technical Service Delivery. This allows VSO to keep overall costs down, while providing an effective approach to a new career for those who have fought hardest for us.

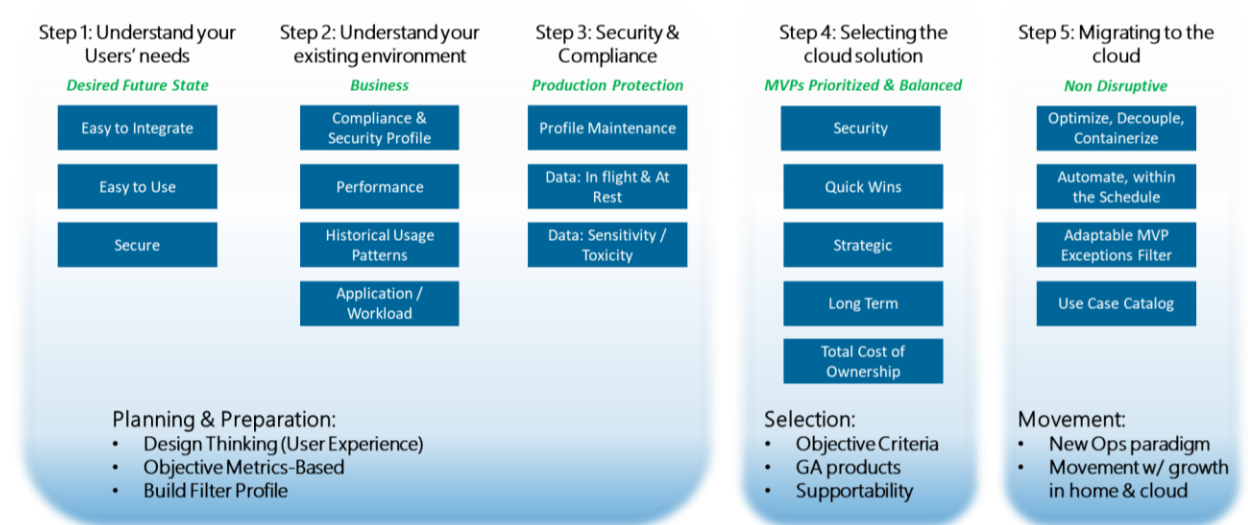
## VSO Approach to a Customer’s Journey to Cloud

VSO follows a five-step process for assisting a customer in their journey. It is VSO’s contention that most cloud migrations are conducted in incorrect order-often the reverse of these five steps-starting with cloud platform selection, then studying the legacy on-prem environment and then beginning the migrations themselves.

At best, these efforts are begun with some internal apps chosen by the customer; at worst, they are wide-ranging enterprise-wide mandated migrations to single-vendor cloud platforms.

These approaches ignore the needs of the Users, and focus instead on the TCO of on-prem vs cloud cost and support financials.

## VSO Recommended Journey to Cloud



## VSO Design Thinking – The Start of the CCoE

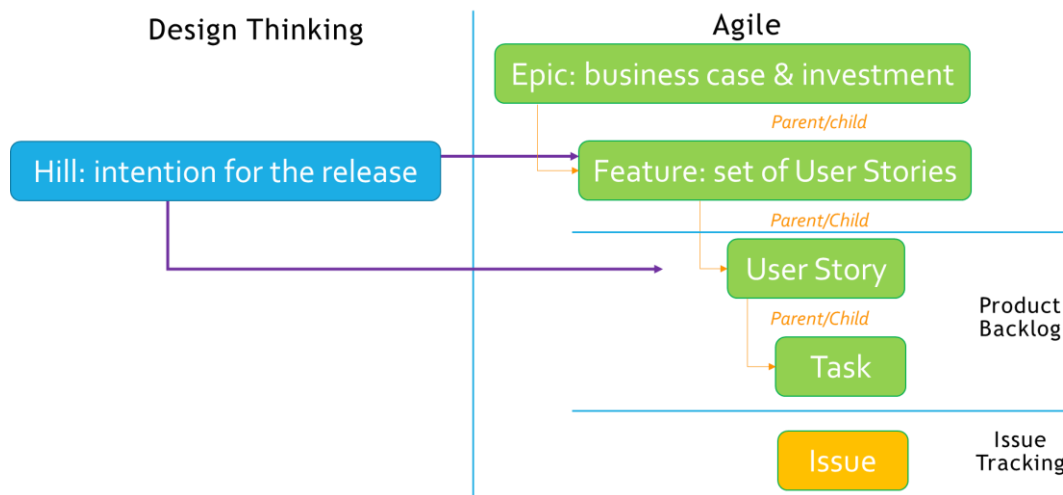
A VSO Cloud Center of Excellence engagement begins with Design Thinking Workshop. The primary goal of Design Thinking is to document key user roles and to define desired user experiences for each role. Each desired experience is captured in a Hill.



**Business goals for the release, framed around user experience**

Design Thinking Hills are aspirational and define the mission and scope of the functions required. They serve to focus the design and development work on desired, measurable outcomes. Who, What, and Wow. Design Hills are broken into multiple User Stories that can, in turn, be grouped as Features for a release.

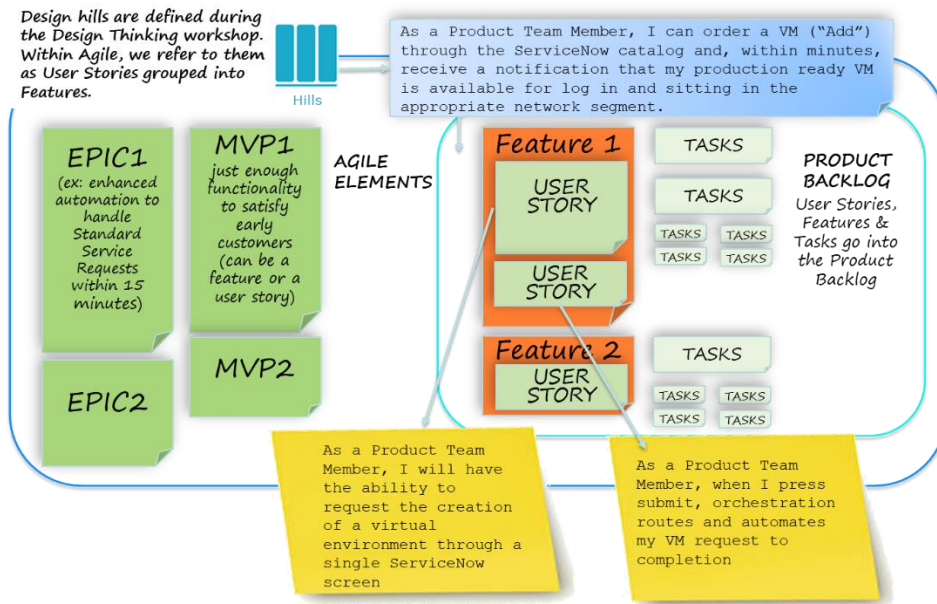
**VSO combines Design Thinking & Agile to achieve real and measurable outcomes:**



**Step 1. Define success with Design Thinking:** everyone conceptually understands hills are to be climbed or taken. From a marketing perspective, it is clear what we are aiming at, hill has both a wow factor and specific metrics-based outcomes.

**Step 2. Achieve success with Agile:** while hills can be reviewed monthly to make sure the team is taking the hills, for the purpose of achieving functionality, in Agile, we start referring to them as User Stories grouped into Features. User Stories, Features and Tasks are managed and refined in the Product Backlog.

# Understanding Agile Concepts



There are many flavors of Agile methodologies, and, adding to the market confusion, Design Thinking itself tends to try to define a comprehensive approach. This can lead to terms being confusingly conflated and customer left wondering how what they are being asked to participate in relates to the larger picture of accomplishing a cloud migration.


VSO has chosen to take the most clear and useful of the terms and methodologies and define them for the customer.

Using only these terms, a project can be defined in terms of a final product backlog for an MVP. ***These elements can be predicted for achievement in a fixed period of time at a fixed price.***



▶ Epics represent a significant program investment delivered by multiple teams and can span teams and/or releases. Business cases are hallmarks of Epics.



▶ Epics are broken down into related, but separate User Stories that can be developed independently. User Stories are capabilities that are released by a team in an iteration or sprint. Sprints are comprised of Tasks → 



▶ Feature is a *releasable* set of User Stories, and these features are organized around the Design Hills defined during the Design Thinking workshop.



▶ An iteration represents a single development cycle, called a Sprint. A daily scrum is performed during a Sprint.





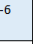



▶ Features, User Stories, and Tasks are maintained via a Release Schedule comprised of multiple Sprints in a Product Backlog.



▶ Playbacks are demos of end-to-end user stories.

▶ MVP releases include just enough functionality to satisfy early customers.

A sample of a resulting MVP Product Backlog is illustrated below:

Notes	Release / Feature	Sprint	Hill / User Story / Task	Source Group	Acceptance Criteria	Priority	ID	Tech Design / Solution / Deliverable	Est. Story Points
MVP1 MVP Minimum Viable Product	TBD	1 - 8	As a Product Team Member, I can order a VM ("Add") through the ServiceNow catalog and, within minutes, receive a notification that my production ready VM is available for log in and sitting in the appropriate network segment.	Fulfillment Orchestration / Automation	TBD	High	FOA1		10
		1 - 4	As a Product Team Member, I will have the ability to request the creation of a virtual environment through a single ServiceNow screen						
			Task #1						
			Task #2						
		5-6	As a Product Team Member, when I press submit, orchestration routes and automates my VM request to completion						
			Task #1						
			Task #2						

VSO can usually produce such a deliverable, following a week-long engagement. The length of the engagement adjusts based upon the number of applications being considered and the number of stakeholders that need to participate.



---

## Customer Case Study

The Customer, a large US insurance company, including several affiliated Corporate Business Units (BUs) other Customer Entities, hereinafter referred to as “Customer”, requested that VSO provide a design and implementation of an enterprise Proof of Concept (POC) for a cloud solution. This effort included cloud provider recommendations, design and implementation of initial POC. VSO provided a team both onsite and remote for the implementation phase.

The Customer engaged VSO to assist them in moving to their organization to cloud. This involved the consolidation of nine completely separate corporate entities onto one shared platform. The goals were to build a platform that would assist these separate Business Units (BUs) with collaboration, file sharing, data sharing, and the use of common tools. The organization was a very data-driven organization and their ultimate goal was a data lake that would allow not just predictive, but proscriptive data analytics-with the aid of machine learning and artificial intelligence tools.

The first part of the effort was to conduct a series of VSO’s customized Design Thinking Workshop with the various BU stakeholders. The “Who” were the various stakeholders of the nine entities, the “What” that came out of those workshops outlined the goal of a solution that included:

- Unified authentication
- File sharing
- Data sharing
- Shared tools (Qlik selected as first shared tool)

The resulting “Wow” of the first phase of the effort was a set of business metrics and KPIs that against which any technical solution could be measured and validated. A product backlog was created to break down Design Thinking Hills into actionable user stories for development. Details articulated in specific stories and tasks ensured the various entities would be happy with-and therefore comply with migration to-the shared platform. These user stories delivered the following aspirational Minimum Viable Product (MVP) Feature Set:

### **Self-Service - Easy to Integrate, to Use, Secure**

- Integrated, self-service portal experience
- Easy search capability
- Upload and download capabilities
- Basic file security
- Create metadata & make specified tags
- Versioning, curation, and stewardship function will be required for Production, but not strictly part of MVP1

### **Heterogenous Data Sources (in native format)**

- Twitter

- Bureau of Labor Statistics
- Weather
- Other key web data (expect 5 in MVP)

#### **Enterprise Access to the Lake**

- Native Qlik access
- Portal, CLI, PowerShell
- REST API over HTTPS

VSO then engaged in a two-month project to select the right cloud platforms and tools to meet the functions and user stories articulated. This collaborative process led to a multi-cloud solution with multiple API tool interfaces and shared dashboards. VSO then assisted with build out of the solution and the migration of target applications and data sets onto the solution.

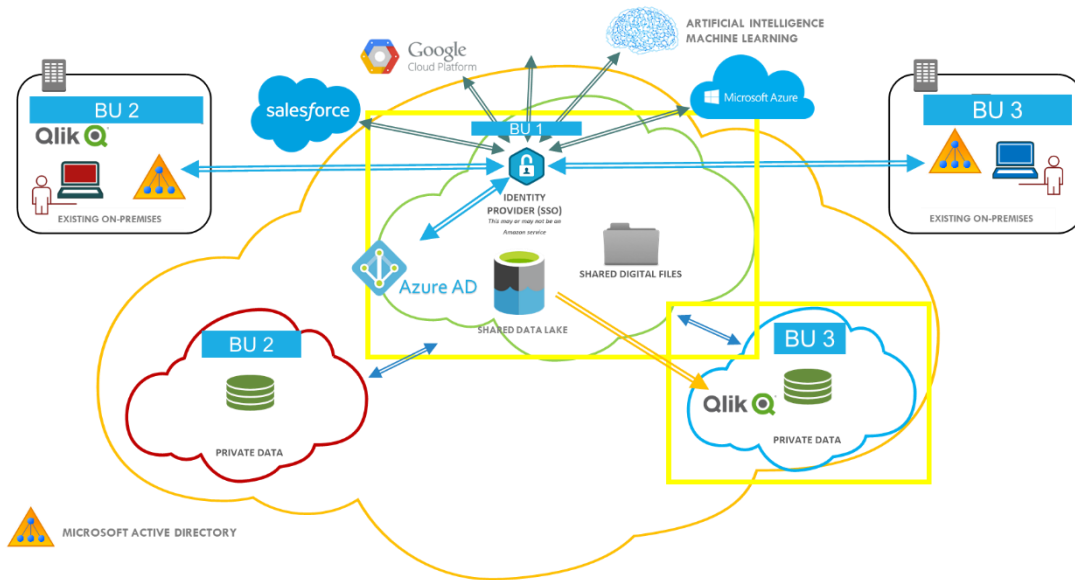
Finally, the Customer requested that VSO provide goal state recommendations for Customer's continued journey to cloud geared to enhance affiliate user experience, accelerate innovation, and enhance Customer's competitive edge.

### ***MVP Approach Detail***

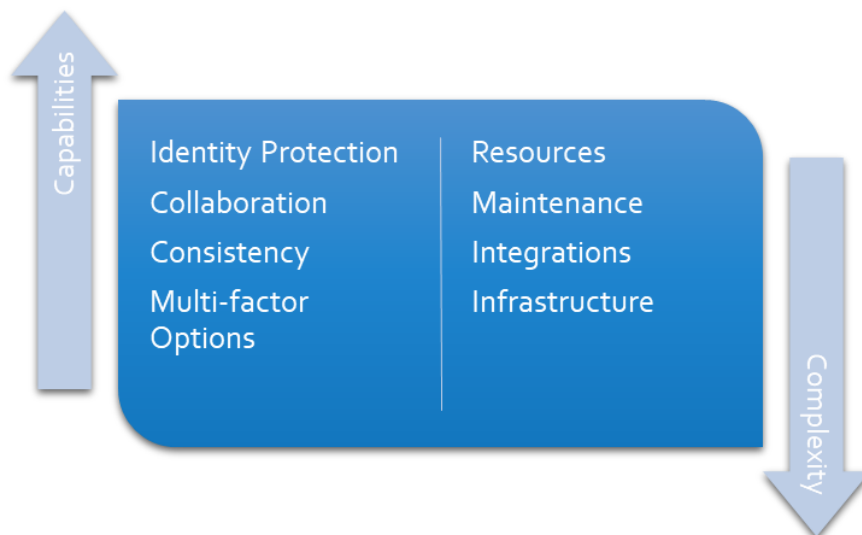
The cloud has a diverse number of providers. VSO is organized in a unique position in the industry as it provides services across many cloud providers such as Amazon Web Services (AWS), Microsoft Azure, IBM Cloud, Google Cloud Platform, and Oracle. We feel that no one provider is best in all cases, and that use cases matter.

Customer's approach aligned with VSO's recommended cloud best practices, in that Customer insists on a data centric solution that balanced best-fit-for-purpose with cost efficiency for each use case. This is particularly important with respect to ongoing management, maintenance, and user training and support.

In MVP1, BU 1, BU 2 Assurance (BU 2a), BU 3 Corporate (BU 3c), BU 3, and VSO designed and implemented a Customer platform tenant, a BU 3 tenant, and a Qlik instance within the BU 3 tenant. Using data collected during an initial workshop with key stakeholders, the multi-entity team implemented a federated identity solution, a Customer platform tenant, single sign-on, a data lake, a collaboration/file share solution within the Customer platform, and an BU 3 Cloud tenant with an instance of Qlik (in yellow below).



At VSO's recommendation, the team chose common services whenever possible, to increase security and interactive capabilities while decreasing complexity (chart below). The choice for common services guaranteed streamlined user experiences even across entities and, at the same time, deduplicated management, maintenance, user training and support efforts. As a result, precious time has been freed up that can now be used for innovation.



---

## Phase 1 - Customer Design to Agile Backlog

With Customer key stakeholders, including two stakeholders from BU 1, two stakeholders from BU2, and three stakeholders from BU 3, VSO conducted an Design Thinking workshop to identify key user experiences desired by Customer for unified authentication, file sharing, data sharing, and shared tools such as Qlik. With the information gathered in this workshop, VSO proposed a Cloud Center of Excellence (CCoE) for Customer where VSO designed and implemented a Customer tenant and an BU 3 Cloud tenant, to include:

- a branded portal for Customer single sign-on;
- a Customer data lake;
- a collaboration/file share solution within the Customer; and
- an instance of Qlik within the BU 3 tenant.

Finally, VSO provided detailed recommendations for technology and service alternatives including goal state recommendations and a VSO proposal for short term hybrid managed services to operate during Minimal Viable Product (MVP)-based wave expansion.

Customer's expected outcomes focus on a data-centric solution that balances best-fit-for-purpose with cost efficiency for each use case, particularly with respect to ongoing management and maintenance. These outcomes were VSO's guiding principles throughout this CCoE. VSO's efforts enabled the Customer to be set up for:

- self-service and self-training for affiliates;
- a shared portal experience for Customer and BU 3 affiliates;
- real-time Customer (Customer's Governance Board) collaboration including co-authoring and versioning;
- a data lake for Customer that includes a feed to BU 3's Qlik instance; and
- a Qlik experience for BU 3 that is repeatable.

---

## Official Customer Enterprise Cloud Strategy

Customer is a Microsoft Windows shop and has had a great deal of experience in the Microsoft tool area. This is one of the main reasons VSO feels that Azure should be a cloud choice. This is reflected in what has been done in MVP1 and the roadmap as well.

To the point earlier that use cases matter, our research for a data lake clearly found that Amazon Web Services (AWS) had the most maturity in this area, with several CloudFormation stacks to quickly stand up and tear down data lakes and catalogs.

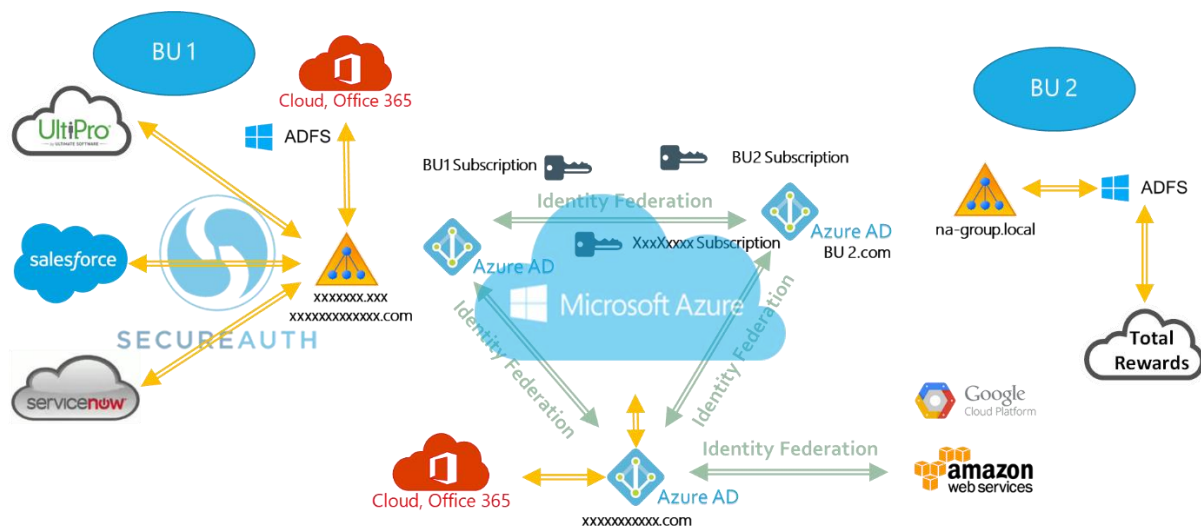
We also recommended, in the use case of production load balancing that may occur in the future, it is highly advisable to use a provider that offers local access. AWS, Azure, and IBM all have Dallas as the closest entry point. Google, however, has an Oklahoma data presence in Tulsa, which is an important aspect for those applications that will require lowest latency response times.

Strategically, VSO proposes a combined Azure, AWS, and Google Cloud presence to benefit Customer the best. In each case, the solution has been designed from the ground up to handle sensitive data. It is recommended that the MVP1 start on Azure and AWS, and Google should be introduced as production requirements ramp up.

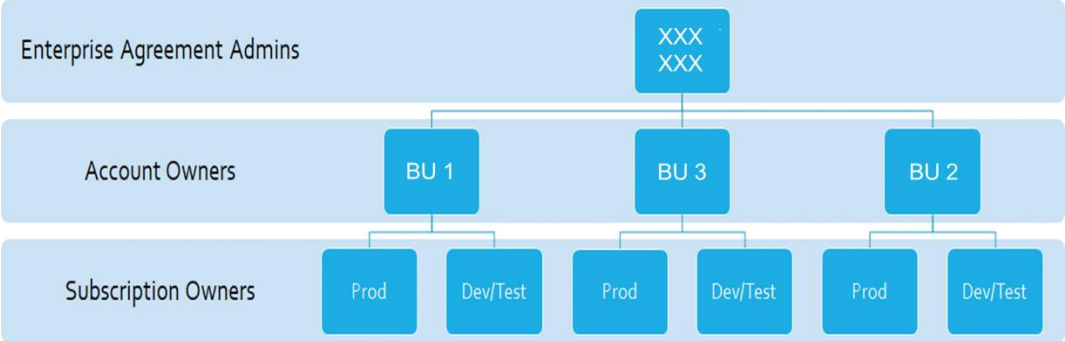
## Identity and Data Security

Identity and access management is crucial to a highly collaborative digital ecosystem. Ideally Customer affiliates should be able to login once to the Customer platform and, through identity federation, be able to access and collaborate on all common services and apps. To achieve this, the team designed a unified authentication service provider (SP) & identity provider (IdP) architecture (below). Because of Customer’s experience with Microsoft, Azure was the cloud service provider selected, and Azure AD identity provider was chosen as the trusted provider for single sign-on (SSO) to access other services and apps.

The domain Customercloud.com was purchased and tenants were set up for the Customer platform and BU 3 within Microsoft Azure.



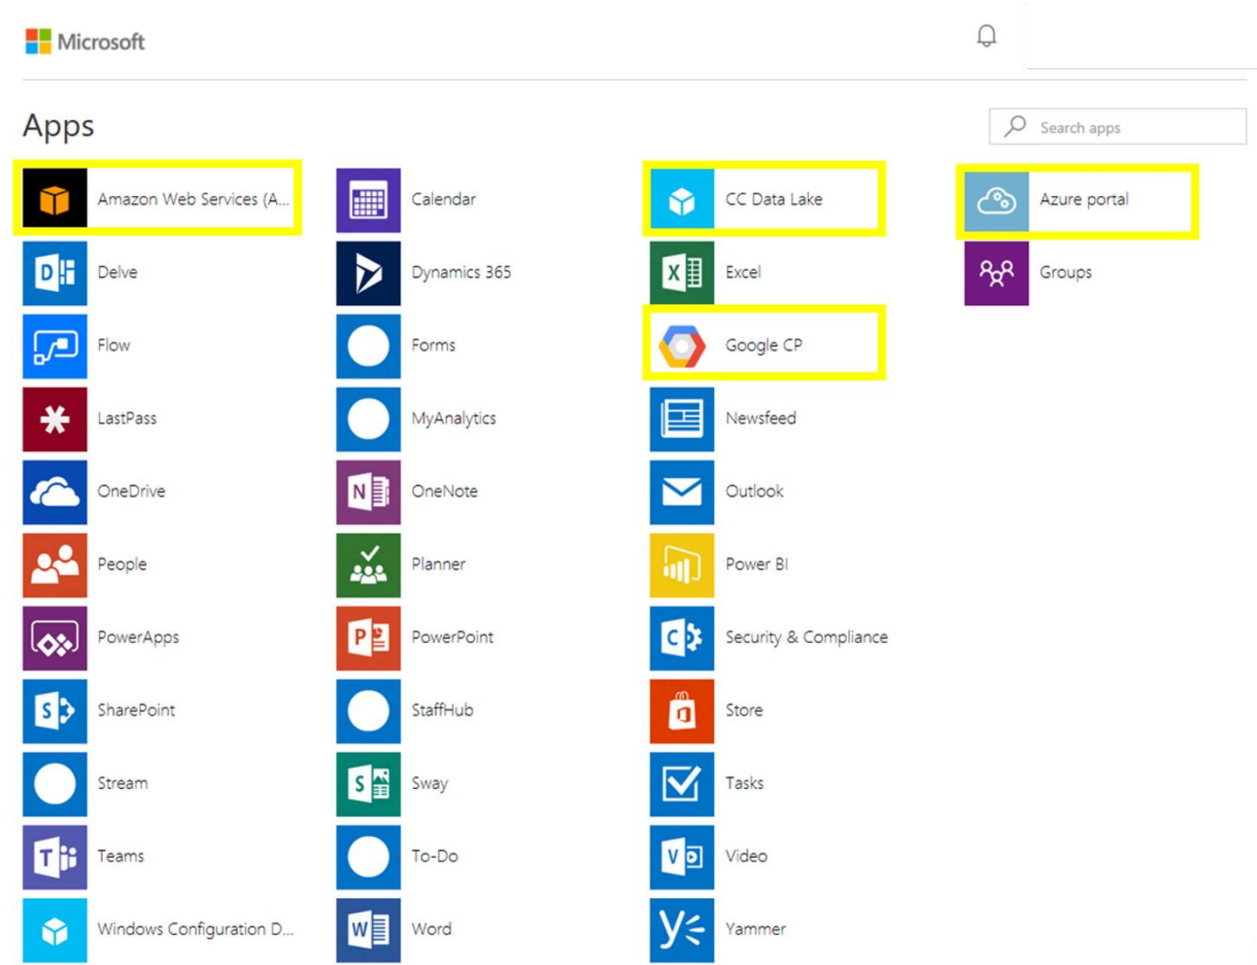
Azure Enterprise Agreement (EA) structure with Production and Dev/Test subscriptions was designed and documented (below), in a manner that can be easily replicated for each Customer affiliate. Azure subscriptions were delegated to BU 3 and BU 2a Assurance administrators to give each control of their Azure AD.



For BU 3, who initially did not have Azure AD, the team evaluated and selected BU 3 IdP and management model that fit BU 3's requirements. Then BU 3 provisioned an on-prem server for Azure Active Directory (AD) Connect, and the team installed BU 3 Azure AD Connect and set up BU 3 AD Sync so that BU 3 users/password hashes now sync up to BU 3 Azure AD directory.

This unified and federated model also made it easy for the team to create Amazon Web Services (AWS) and Google Cloud Platform (GCP) tenants and federate securely into both. Single Sign-on was enabled to provide seamless access to apps at each step in the project (see

below).

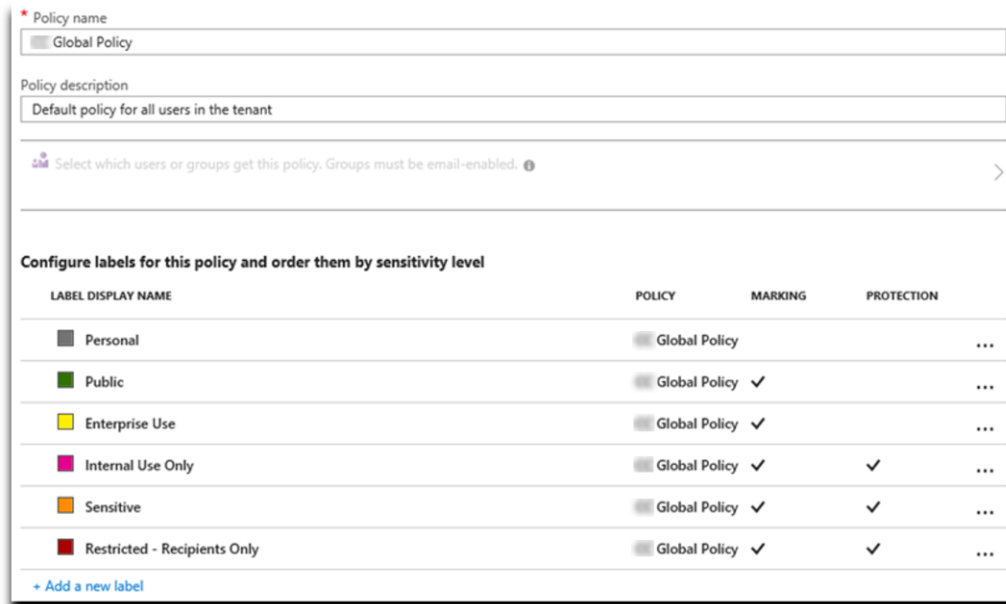


To enhance identity protection, multi-factor authentication (MFA) was enabled for all users. For Customer data security, a strategy for file/data labels & data loss prevention (DLP) in Office 365 was designed and executed. This strategy is currently under Customer platform Governance review.

Label	Description
Personal	Non-business data, for personal use only. Examples include personal documents, schoolwork, or personal receipts. Limit the amount of personal data you keep on company assets to minimize its risk of accidental disclosure.
Public	Business data that is specifically prepared and approved for public consumption.
Enterprise Use	Business data that is not intended for public consumption. However, this can be shared with other Cameron entities, and business partners as required. Examples include a company internal telephone directory, organizational charts, internal standards, and most internal communication. This label does not add protection to the file.
Internal Use Only	Data that is sensitive and poses some risk to the organization if it were breached. This label applies to the individual organization using it (e.g. AFC, AFA, Insurica). Data that is appropriate for all Cameron entities to use should be labeled "Enterprise Use". This label adds protection to the file.
Sensitive	Sensitive business data that could cause damage to the business if shared with unauthorized people. Examples include individual customer HIPAA data, PII, contracts, security reports, forecast summaries, and sales account data. This label adds protection to the file and is automatically added when low counts of sensitive data types are detected.
Restricted – Recipients Only	Very sensitive business data that would cause damage to the business if it was shared with unauthorized people. Examples include bulk employee and customer information, passwords, source code, and pre-announced financial reports. This label adds protection to the file and is automatically added when high counts of sensitive data types are detected.

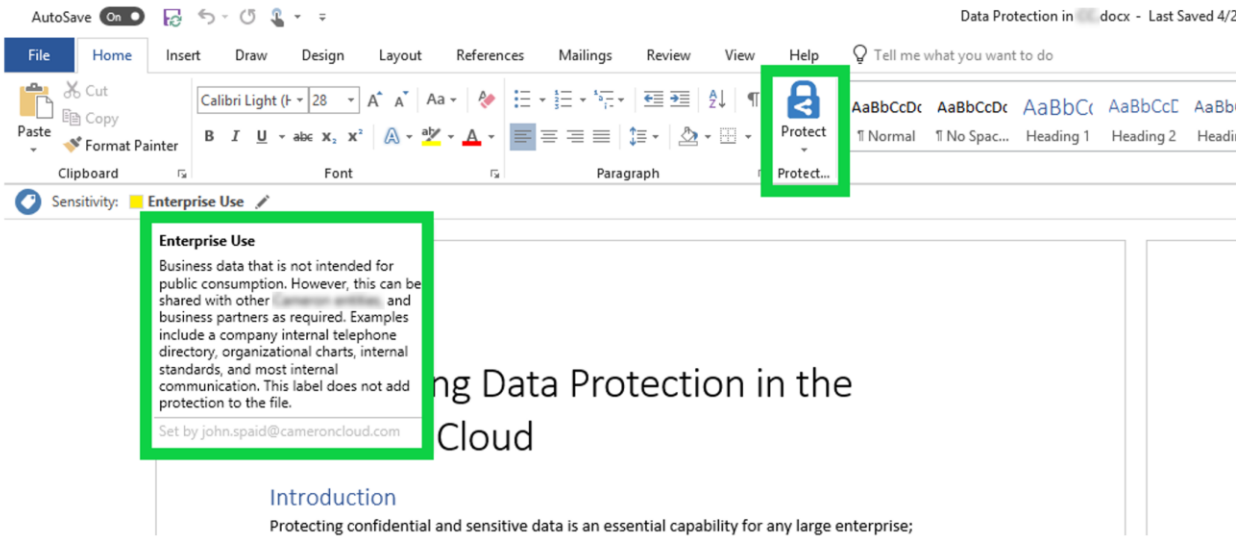
First, data classification and retention policies were codified for Personal, Public, Enterprise Use, Internal Use Only, Sensitive and Restricted labels (above).

The team added and activated Customer platform Global Policy in Azure Information Protection.



Classification is configured to be automatic and user managed (see automatic example below).





## Resource Tagging

To track and manage costs across entities, the team designed a resource tagging strategy in Azure that combines BU 1 recommended resource tagging for affiliate, application owner, and application custodian, with Azure tagging best practices recommendations. After reviewing Azure best practices (<http://www.azurefieldnotes.com/2016/07/18/azure-resource-tagging-best-practices/>), the team recommends Azure's tags, adjusted slightly for required vs optional, for Customer platform (below).

### Azure Best Practices for Resource Tagging

- AppTaxonomy (required)
- EnvironmentType (optional)
- MaintenanceWindow (optional)
- BillingIdentifier (optional)
- ExpirationDate (optional)

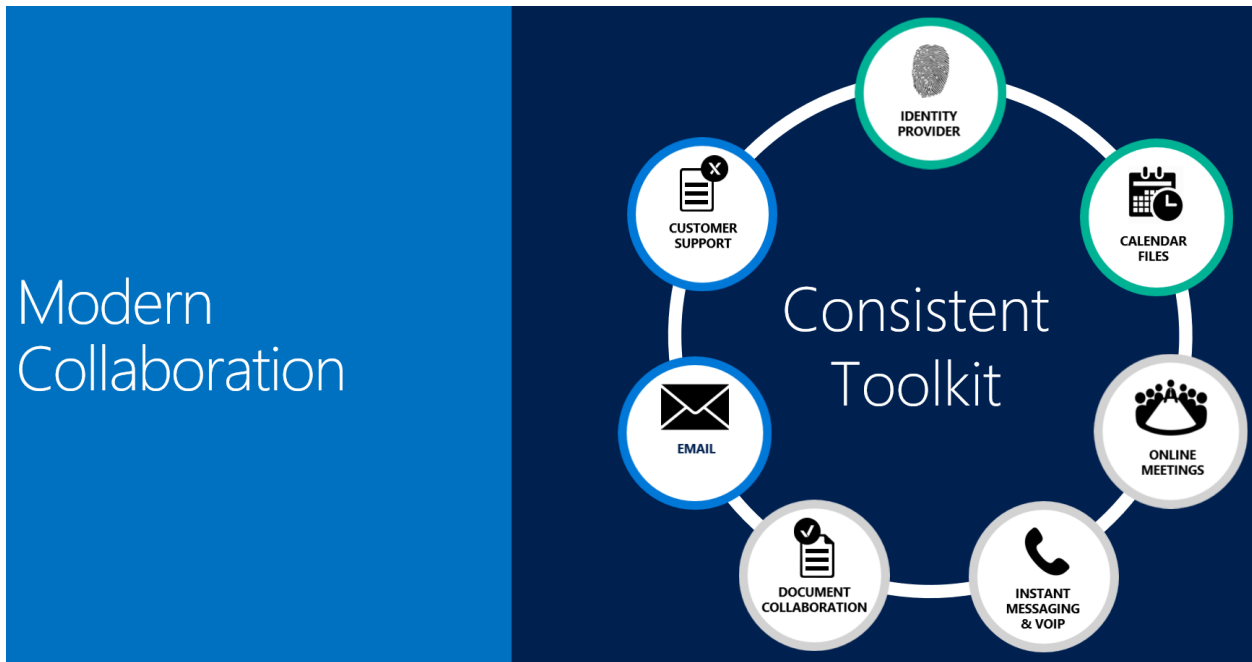
## Collaboration

For BU 1 MVP1 collaboration and file sharing, the team designed a two-pronged approach. The first being Office 365, an instance of which was stood up in the Customer platform. Throughout the project, the VSO, BU 1, BU 2, BU 3a, and BU 3 team collaborated via Office 365's SharePoint Teams in the Customer platform. Modern collaboration has evolved from being siloed for Customers, Partners and Internal affiliates (below):

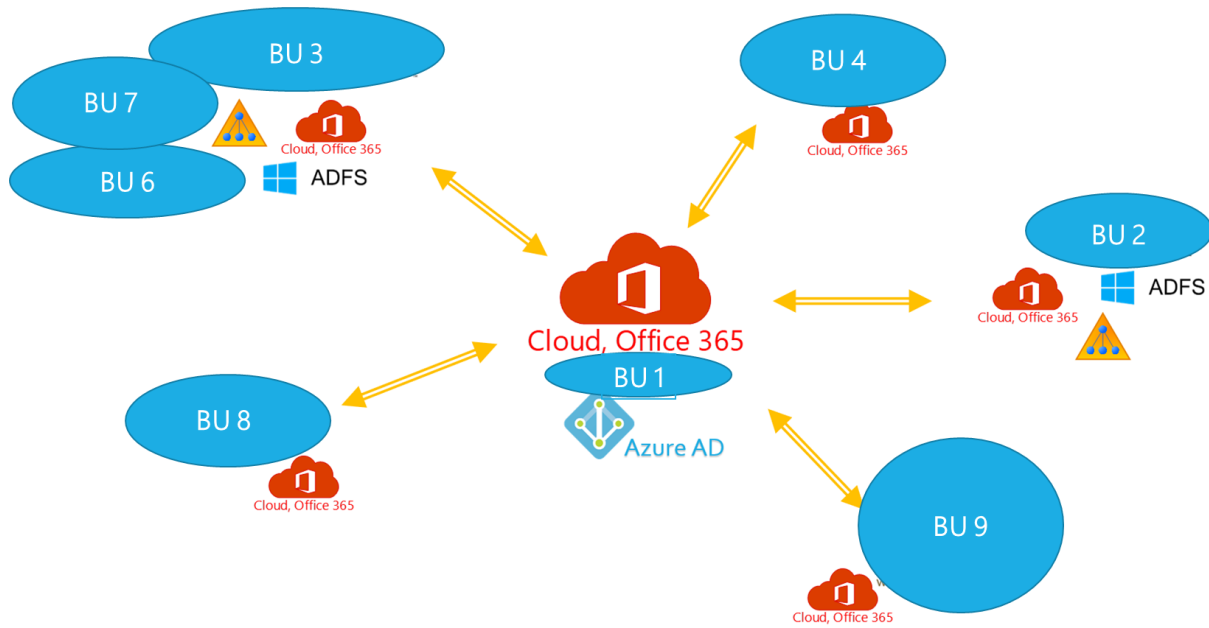
## The evolution of Collaboration



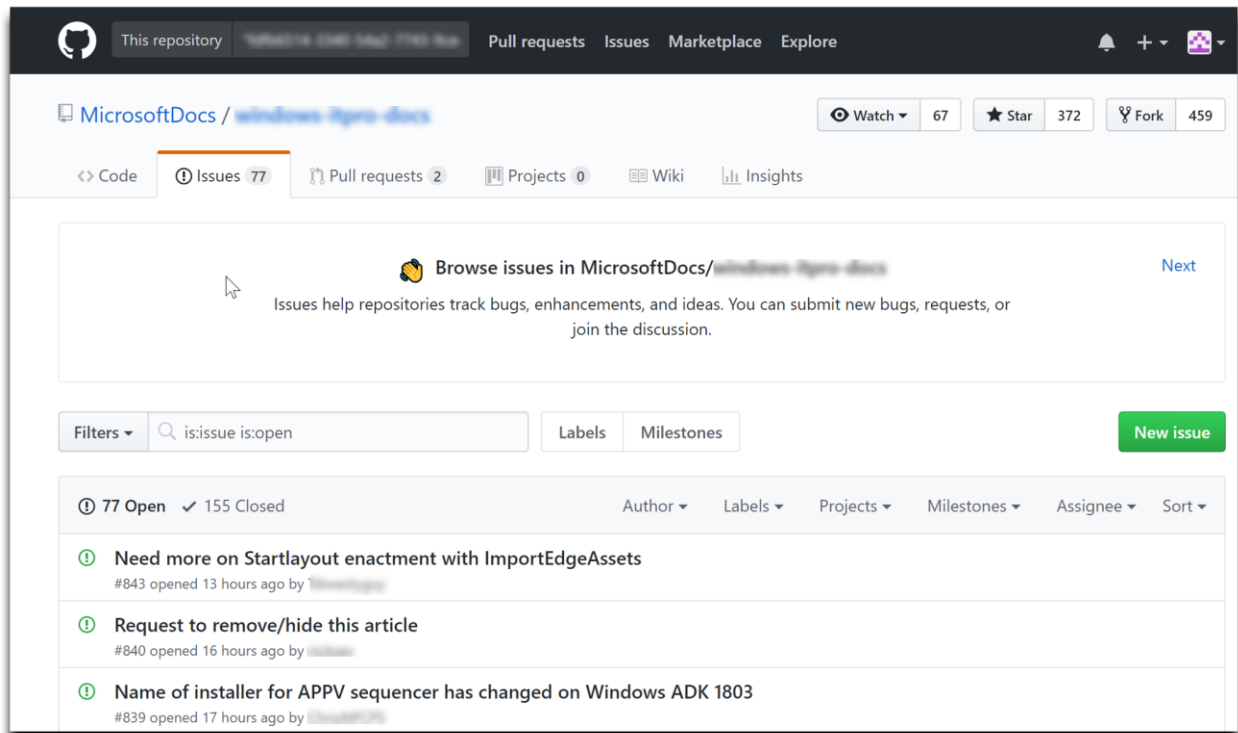
Strategically, Office 365 provides highly interactive and modern collaboration, with its ability to provide a consistent toolset for meetings, instant messaging & voice over IP, document collaboration, email, calendar, and mobile endpoint management.



It is recommended that Office 365 become a standard across Customer entities and that it be fully federated so that affiliates from each entity can collaborate with all others. The initial high-level architecture is shown below:

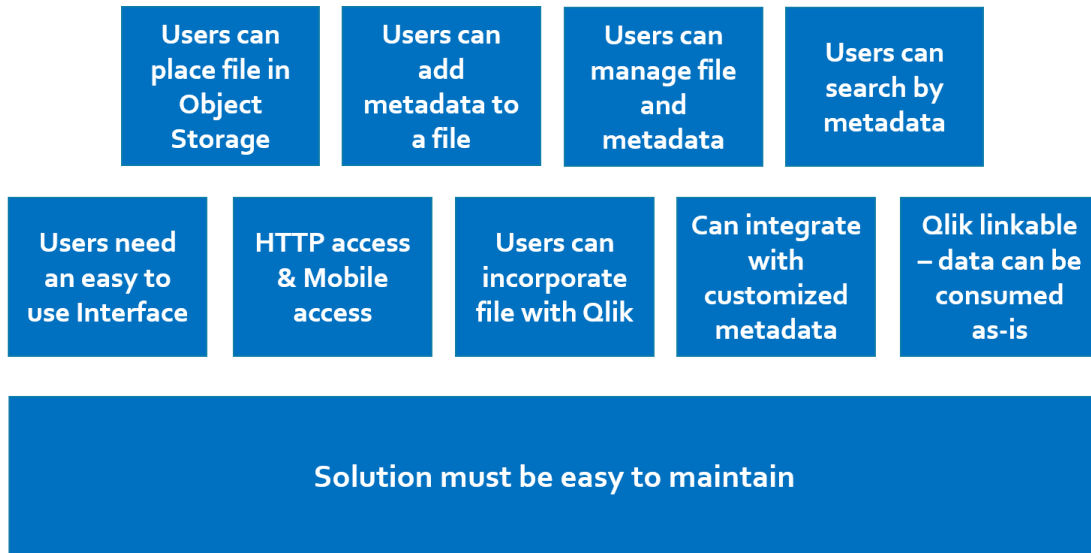


For the second prong, the team design and created a Git Repository cross-entity knowledgebase that will enable affiliates to self-manage technical and instructional documentation, updates and issues, through a self-service Git interface/wiki. The framework has been laid (an example below shown using Microsoft Docs) and the initial population of this repository will begin in Phase 2 of the roadmap.



## Data Lake

For the Customer platform data lake, VSO researched BU 1's functional requirements, including:



Specifically, the Customer team required the data lake to be self-service for affiliates, and it needed to accommodate data in its raw format, whether structured, unstructured or semi-structured. It must support Qlik access for data visualization. Below is the detailed list of requirements:

### Self-Service - Easy to Integrate, to Use, Secure

- Integrated, self-service portal experience
- Easy search capability
- Upload and download capabilities
- Basic file security
- Create metadata & make specified tags
- Versioning, curation, and stewardship function will be required for Production, but not strictly part of MVP

### Heterogenous Data Sources (in native format)

- Twitter
- Bureau of Labor Statistics
- Weather
- Other key web data (expect 5 in MVP)

### Enterprise Access to the Lake

- Native Qlik access
- Portal, CLI, PowerShell
- REST API over HTTPS

Given these, VSO recommended Amazon Web Services (AWS) for the Customer platform data lake. Amazon is a leader in this space with deployment guides, best practices and CloudFormation templates to made stack deployment and deletion quick and easy. In addition to meeting Customer's requirements, AWS provides the following:

- native security, including encryption and versioning

- highly durable, available & scalable
- low cost (Amazon Simple Scalable Storage (S3) with lifecycle policies to Glacier)
- more robust with a roadmap that continues to grow
- more mature environment for big data
- natively supported by frameworks like – Spark, Hive, Presto, etc.

The team deployed two data lakes in AWS. The first most closely fits the MVP1 requirements, with a self-service portal experience, easy file upload, easy search, ability update tags and metadata and native Qlik access (see pros and cons below).

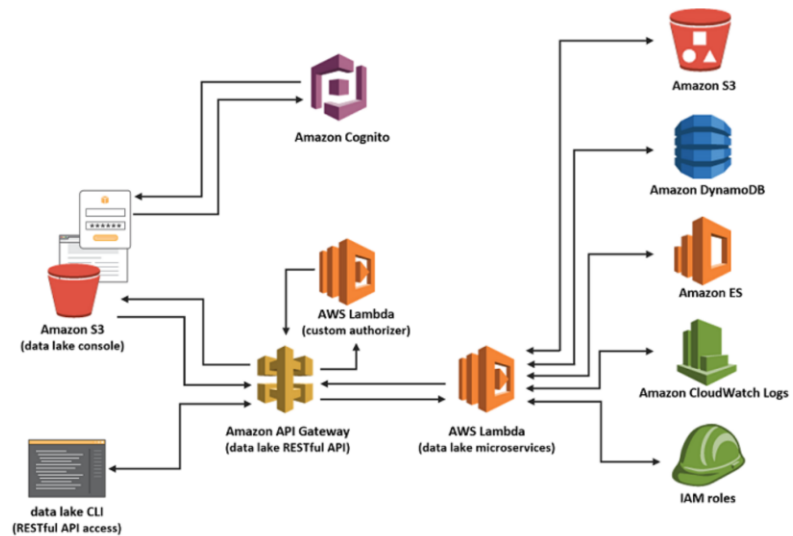
### Data Lake Solution #1

**Pros:**

- Self-service portal experience
- Easy to upload files
- Easy to search, and update tags and metadata
- Native Qlik access

**Cons:**

- No way to curate data
- Downloading a file requires url lookup
- Tags were manually entered (human error prone)
- Complex management
- Development team required to run and update code base
- Manual security configuration required for each service
- No AngularJS expertise



Out of the box, this solution did not offer a way to curate the data, but it did meet MVP1 requirements for tagging, search, upload and download via a self-service portal. The team worked with AWS to deploy another more robust data lake solution with data curation, and with Microsoft, however, these fell short of the MVP1 requirements. Solution #2 offered curation but did not have a self-service portal for affiliates (see pros and cons below).

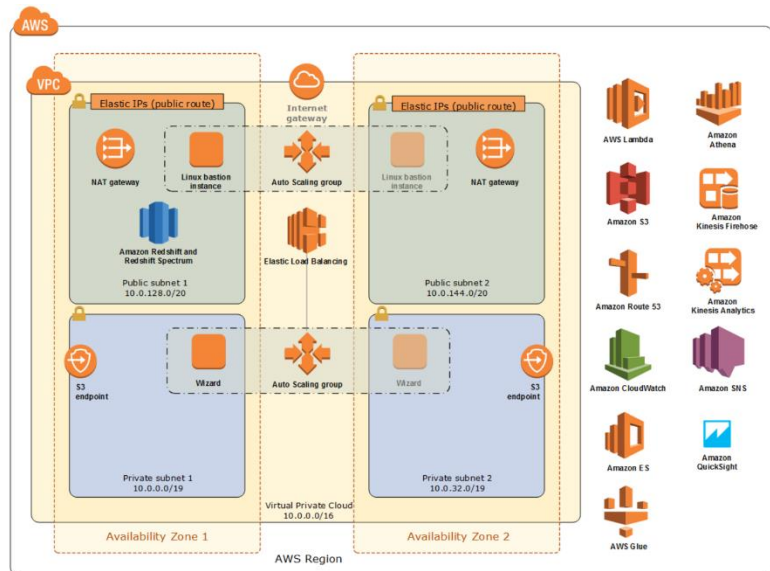
## Data Lake Solution #2

### Pros:

- Robust ingest processing, analytics, visualization and search
- Push button data curation - transformation, aggregation, and analysis
- Native Qlik access

### Cons:

- No self-service portal
- Complex management
- Development team required to run and update code base
- Manual security configuration required for each service
- No AngularJS expertise



Solution #3 was an agile team decision to test drive the Azure Data Lake solution. It met nearly all of MVP1's requirements, but it broke the use case because it did not allow guests to use the Customer platform Azure Data Catalog; instead it redirected them to their home Azure AD tenants (see pros and cons below).

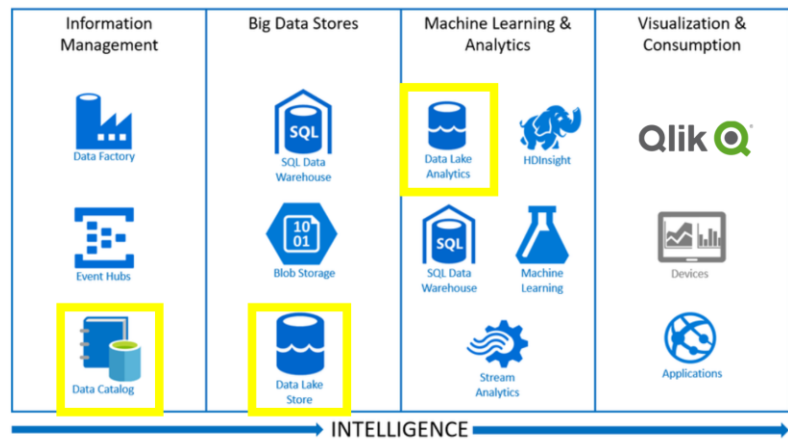
## Data Lake Solution #3

### Pros:

- Self-service portal - easy to upload, tag, download, search
- Nice abstraction of end-to-end processing
- Easy to manage
- No Development team required
- Built-in security (MFA, AIP, ....)

### Cons:

- Qlik access to be proven
- No guest access to Data Catalog

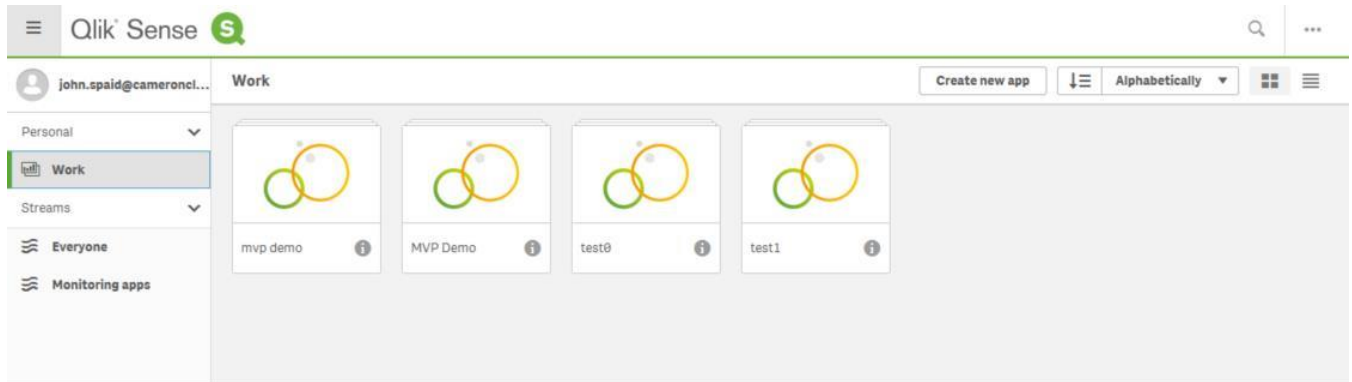


## Qlik

Customer has selected Qlik for its first shared tool in the Customer platform. Qlik lets affiliates combine any number of data sources so they can explore across all data sets and instantly adjust based upon what is seen. For MVP1, an BU 3 installed Qlik instance needed to successfully load a Customer platform data source from within BU 2a and BU 3.

The team built an Azure Resource Management (ARM) template to make the BU 3 installation as automated as possible for any Azure deployment, though based on some Qlik issues experienced during installation, there may be some manual work to be done by administrators.

Below is the Qlik application in the BU 3 Azure tenant, with data pulled from the data lake Simple Storage Service (S3) bucket in Amazon Web Services.



---

## Customer Platform Governance Model

VSO recommends the ideal Cloud Governance Model for Customer be an agile one that brings together the organization stakeholders and technical leads to resolve issues of technical control and accountability. It can provide the cross-organizational forum for decisions based on recommendations from the affiliate technical teams. It is recommended to use metrics to show organizational maturity along with schedules, red/yellow/green status, etc.

Customer's cloud governance aligns with VSO's approach, in that Customer insists on a data centric solution that balances best-fit-for-purpose with cost efficiency for each use case. As mentioned previously, this is particularly important with respect to ongoing management, maintenance, and user training and support.

VSO begins every cloud use case with the user's desired experience or future state. We believe there are three elements that are key to success. First, a solution must be easy to integrate. This is particularly important because of the impact it can have on users' time, efficiency and ability to innovate. Tools not only need to work, but they need to talk to each other, and should never slow users down. Tools that cannot integrate easily should be eliminated. Busy development communities are hallmarks of robust integration capabilities.

A solution must also be easy to use. The Customer team summed this up perfectly in the workshop when indicating the ideal user experience should mirror that of a mobile app user, with apps up and running in minutes, self-service support, low touch management, automated (invisible) maintenance, and just in time provisioning.

Finally, as cloud capabilities expose more surfaces to protect, security must be easily enforceable, either push-button or fully automated. Solutions such as auto-filtering for HIPAA



and PCI information and auto-tagging of files and data need to deliver on their promises to protect the enterprise, especially when human efforts fail.


## Customer Platform Roadmap

### Data Catalog and Lake

With MVP1 data lake in place, work can be done to make the data lake production ready, including adding curation, versioning, and stewardship. VSO

recommends getting data classifications into the lake, in a centralized metadata catalog (benefits shown right):

- Create a centralized management system / Data Catalog to manage all object metadata.
- Allow access to metadata across multiple zones
- Allow tracking object through zones and refinement cycles
- Enable more robust metadata than standard metadata for content
- Increase ability to correlate content



In MVP1, the team

developed an initial metadata catalog or glossary and VSO recommends continuing this effort across the entities. Below is VSO's best practice metadata catalog, with MVP1 types highlighted in green. To the right, are the benefits that the technical, business and operational types provide.

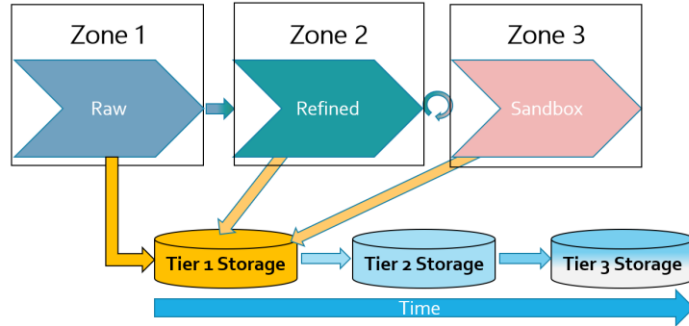
MVP1 Data Catalogue		Type	Delivers	
Metadata	Technical	<ul style="list-style-type: none"> <li>• Data Source</li> <li>• Lineage</li> <li>• Data Type</li> <li>• Location of Data</li> <li>• File Name/ Table Name</li> <li>• File Size</li> <li>• Creation Timestamp</li> <li>• Owner</li> </ul>	Searchability	<ul style="list-style-type: none"> <li>• Traceability</li> <li>• Data Quality</li> <li>• Growth</li> <li>• Automation</li> <li>• Data Flow</li> </ul>
	Business	<ul style="list-style-type: none"> <li>• Sensitivity</li> <li>• Compliance Data</li> <li>• Company</li> <li>• Department</li> <li>• Data Classification/Category</li> <li>• Business Relations</li> </ul>		<ul style="list-style-type: none"> <li>• Accountability</li> <li>• Data Governance</li> <li>• Regulatory Compliance</li> <li>• Auditability</li> </ul>
	Operational	<ul style="list-style-type: none"> <li>• Last Modified Date</li> <li>• Last Accessed Date</li> <li>• Number of Times Accessed</li> <li>• Access Security</li> <li>• Retention Period</li> <li>• Modelling Data</li> </ul>		<ul style="list-style-type: none"> <li>• Change Tracking</li> <li>• Information Lifecycle Management</li> <li>• Data Tiering</li> </ul>

The team can continue to feed common data and begin optimization, such as configuring lifecycle management policies to keep costs in check (see tiered storage model below):



Managing metadata in the data lake is important to ensure that information derived from the data lake is trusted and quantifiable.

It is important to derive a systematic approach to defining elements. This will ensure consistency in a complex heterogeneous data environment as data progresses through the data lake.



## Add Remaining Customer Entities

As the MVP waves move toward production, VSO recommends adding the next set of Customer entities into the Cloud, in the following order:

1. BU 7
2. BU 8
3. BU 5

Each affiliate may require a new tenant, production and dev/test subscriptions, and an owner/administrator. This team can assist each entity with making IdP and management decisions. For example, with BU 3 we worked through the pros, cons and risks for each option and BU 3 chose the self-managed Identity Provider solution (center, below).

Existing AD FS - BU 2-managed	BU 2-managed Azure AD	BU 3-managed Azure AD
<p><b>Pros</b></p> <ul style="list-style-type: none"> <li>• Maximum BU 2 control</li> </ul> <p><b>Cons</b></p> <ul style="list-style-type: none"> <li>• Highest human resources requirement</li> <li>• Server overhead</li> <li>• Human error prone</li> <li>• Low availability</li> <li>• Governance outside of BU 3 strategy</li> </ul> <p><b>Risk</b></p> <ul style="list-style-type: none"> <li>• Highest risk</li> <li>• Least mature analytics</li> </ul> <p><b>Questions</b></p> <ul style="list-style-type: none"> <li>• Where are the servers now?</li> <li>• How many are there?</li> <li>• Disaster recovery plan?</li> </ul>	<p><b>Pros</b></p> <ul style="list-style-type: none"> <li>• Simplified management</li> <li>• Maintains BU 2 control</li> <li>• No server overhead</li> <li>• Wide array of built-in apps</li> <li>• High-availability built-in</li> <li>• Mature analytics</li> </ul> <p><b>Cons</b></p> <ul style="list-style-type: none"> <li>• Small human resource requirement</li> <li>• Governance outside of BU 3 strategy</li> </ul> <p><b>Risk</b></p> <ul style="list-style-type: none"> <li>• Low Risk</li> </ul> <p><b>Questions</b></p> <ul style="list-style-type: none"> <li>• Does BU 2 want to manage their IdP?</li> </ul>	<p><b>Pros</b></p> <ul style="list-style-type: none"> <li>• Simplified management</li> <li>• No human resource requirement (BU 3 manages)</li> <li>• No server overhead</li> <li>• Wide array of built-in apps</li> <li>• High-availability built-in</li> <li>• Governance within BU 3 strategy</li> <li>• Mature analytics</li> </ul> <p><b>Cons</b></p> <ul style="list-style-type: none"> <li>• Least control for BU 2</li> </ul> <p><b>Risk</b></p> <ul style="list-style-type: none"> <li>• Low Risk</li> </ul> <p><b>Questions</b></p> <ul style="list-style-type: none"> <li>• Does BU 2 want a customer/service-provider relationship with BU 3?</li> </ul>

## Moving to a Highly Collaborative Model


Customer is taking a smart approach to cloud adoption. By identifying what is most important to BU 1 and the Customer affiliates, and insisting solutions meet their needs, Customer is

putting the guardrails in place to ensure the journey to cloud is both cost efficient and best fit-for-purpose. Below is a current affiliate technology summary:

Technology Platform	BU 1		BU 8	BU 3	BU 5	BU 7
Identity Provider	<a href="#">SecureAuth</a>	AD FS	Azure AD	Azure AD	None	Azure AD
	Azure AD	Okta		AD FS		
Device & Endpoint Management	AirWatch		None	None	None	None
Email & Calendar	Exchange On-Prem		Exchange On-Prem	Exchange On-Prem?	???	Exchange Online
Instant Messaging & Chat	SfB On-Prem		None	None	Cisco Jabber	<a href="#">SfBO</a>
Online Meetings	SfB On-Prem		None	None	None	<a href="#">SfBO</a> /MSFT Teams
	MSFT Teams					
Telephony	SfB On-Prem w/ Zeacom		???	???	???	<a href="#">SfBO</a> w/ PSTN
Data Classification	Manual		Manual	Manual	Manual	Manual
Document Collaboration	SharePoint On-Prem					SharePoint Online
	SharePoint Online					
Personal File Storage	OneDrive		On-Prem File Shares	On-Prem File Shares		OneDrive
	On-Prem File Shares					
Task Automation	None		None	None	None	Flow

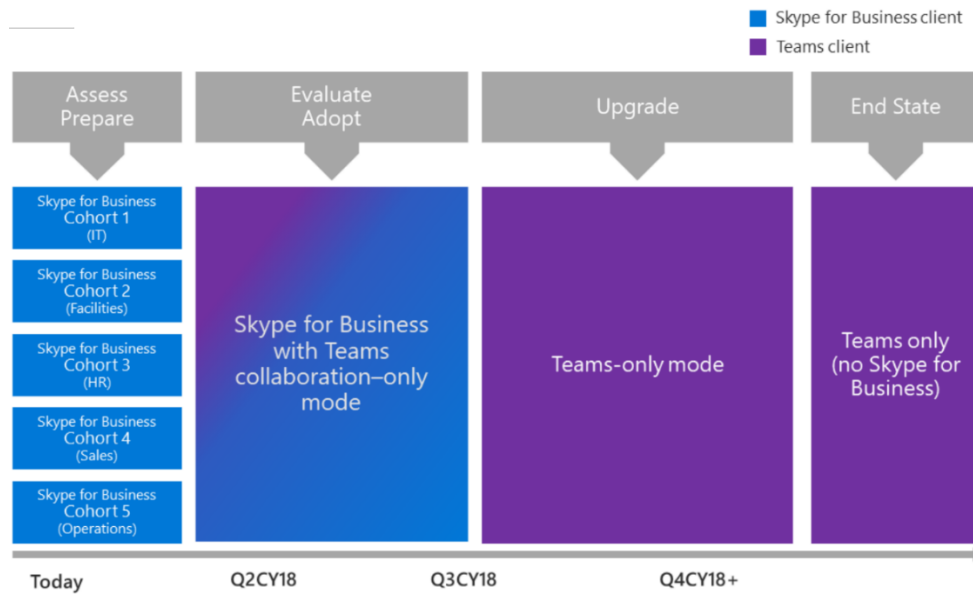
VSO recommends beginning Phase 2 by implementing a common Identity Provider across the entities, and architecting/implementing a device and endpoint management that will automatically protect all Customer affiliate data on company and non-company devices. Data classification and data loss prevention policies should be standardized for each entity and be implemented upon BU 1 Governance Board approval as early as possible. Although each entity has their own compliance and regulatory requirements, there exists an initial set of classifications and policies developed by BU 2a that may serve to expedite this process for the other entities.

With data and identities security ensured, enabling collaboration apps should be the next area of focus. Customer entities have chosen Office 365 for their online collaboration suite. Each Office 365 should be federated to the other instances within the Customer enterprise as shown below.



For email and calendar, a migration method, plan and timeline to move mailboxes into Exchange Online should be developed and executed. This plan should include a review of limits and best practices for Exchange Online to make sure Customer affiliates will get the expected performance and behavior after migration.

Skype for Business on-premise for online meetings and telephony can be migrated to Skype for Business Online (SfBO). It should be noted that Microsoft Teams is scheduled to replace Skype for Business and this should be factored into each migration strategy. Microsoft has provided multiple upgrade and co-existence models, building blocks, and journeys (simple journey shown below).



Once a journey path is chosen, VSO recommends a pilot set of users per entity be moved first to confirm that the new environment is correctly configured. This way, Customer can verify features and services function as expected before attempting to move more affiliates.

For SharePoint on-premise servers, VSO recommends an assessment to assist in migration to SharePoint online (SPO). The method selected to move data can depend on factors including the amount of data and the size, what data needs to move and where content is currently stored. Taking the time to plan the migration, including taking an inventory and assessment of each entity's data, will ensure a successful migration.

VSO recommends affiliate personal file storage be moved from on-premise file servers to OneDrive for Business in Office 365. OneDrive is great for storing and sharing files and Office 365 Enterprise E3 and E5 plans both give 5 TB of initial storage per user. Migrations to OneDrive for Business and SharePoint Online can be simplified with Microsoft's SharePoint Online Migration API. The API integrates with Microsoft Azure to minimize the number of calls required to complete the migration.

Task automation has always been key to freeing up time for new ideas and innovation. Microsoft Flow is a service that helps affiliates create automated workflows between favorite apps and services to synchronize files, get notifications, collect data, and more. This is a newer feature of Office 365 that can be powerful for affiliate users.

## ***Customer Enterprise Knowledgebase***

Though the framework has been created for Git Repository-based cross-entity knowledgebase, significant work remains to be done to populate it with technical instructions that are Customer focused and to make it a one-stop-shop for affiliate learning.

An initial list of architectural documents and instructions to be created and uploaded for affiliate users includes:

- Best practices and how to use Office 365 collaboration in Skype for Business Online and in Teams, tailored for affiliates
- How to use single-sign on across the Customer entities, including self -help
- How to use device/endpoint management software features
- How to use manual and automatic data classification for common Office documents
- How to share files (i.e. how to share from Office 365, from SharePoint, from OneDrive for Business, from the Git Customer Enterprise Knowledgebase, to name a few)
- How to interact with the Data Lake, including upload, search, tagging and download
- How to interact with Qlik, including admin and normal user support and advanced user experience
- How to use Microsoft Flow for key Customer tasks

As each Customer platform component is rolled out, training should be developed to enable users to adopt technologies quickly, ideally at the speed with which a mobile app user learns a new app. VSO recommends short videos and step-by-step instructions for accelerated understanding and make the learning experience quick and painless.

---

## Case Study Conclusion

In MVP1, BU 1, BU 2a Assurance (BU 2a), BU 2, BU 3 Corporate (BU 3c), BU 3, and VSO successfully designed and implemented a Customer platform tenant, an BU 3 tenant, and a Qlik instance within the BU 3 tenant. Using data collected during an initial workshop with key stakeholders, the multi-entity team implemented a federated identity solution, a Customer platform tenant, single sign-on, a data lake, a collaboration/file share solution within the Customer platform, and an BU 3 Cloud tenant with an instance of Qlik.

We will use the momentum of MVP1 to continue the journey to Customer platform. We recommend completing the readiness of the data lake for production, adding the balance of the Customer entities to the Cloud, federating Office 365, centralizing Office 365 tools such as Exchange, and developing the Customer platform Enterprise Knowledgebase. These recommendations are mapped to a timeline below.

