cimplifi™

Strategy for Migration to the Cloud from an On-Premises Instance

The purpose of this document is to set out relevant planning considerations and management steps involved in the migration of a data center and the DocuSign Insight application, to a cloud-based architecture. This document also addresses considerations for maintaining both on-premises instances and cloud instances and how they impact model building and instance management.

Migration to a cloud environment is a highly complex undertaking. Operational success demands equally detailed consideration of stakeholder management, legacy systems, risk management, governance, and data migration. The migration approach can be broken out into 3 phases which set out best practices for engaging each of these considerations in a timely and effective manner:

1. Pre-Migration
2. Migration Approach
3. Post Migration

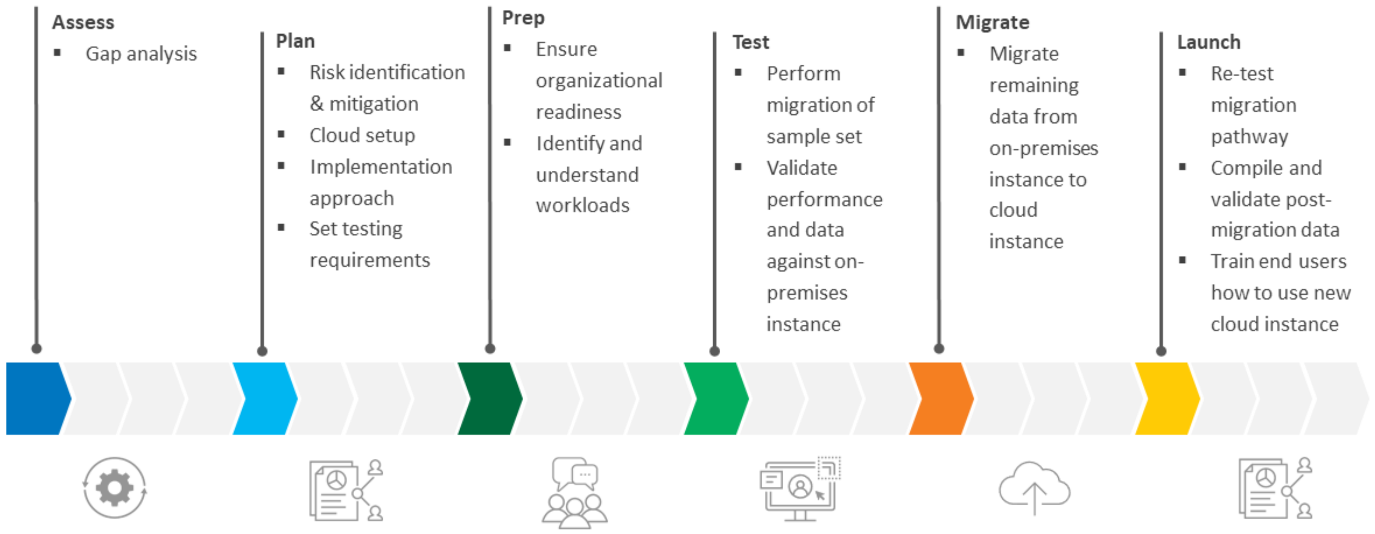

Assess
- Gap analysis

Plan
- Risk identification & mitigation
- Cloud setup
- Implementation approach
- Set testing requirements

Prep
- Ensure organizational readiness
- Identify and understand workloads

Test
- Perform migration of sample set
- Validate performance and data against on-premises instance

Migrate
- Migrate remaining data from on-premises instance to cloud instance

Launch
- Re-test migration pathway
- Compile and validate post-migration data
- Train end users how to use new cloud instance

Figure 1: Cloud migration workflow

Stage 1: Pre-Migration

Stage one involves building an intimate understanding of the existing architecture that is currently on-premises and defining the desired end state. Failure at this stage may result in increased costs, delays or operational challenges diminishing the efficacy of the new Cloud instance.

a. Gap Analysis

A gap analysis is an effective approach at initiation to lay the foundations for the creation of a comprehensive strategy that accounts for all relevant nuances of the business and its incumbent structure. This methodology requires an analysis of the ‘As-Is’ on-premises environment and proposed ‘To-Be’ Cloud environment to identify what can be retained and what must be added, deleted, amended, or replaced.

Definitions

- As-Is Architecture – The current structure of the on-premises instance considering data processes, governance, and user organization.
- To-Be Architecture – The end state represented by the cloud instance on which DocuSign Insight will continue to exist. This view will be considerate of the inherent differences that characterize a cloud instance.

This analysis should take account of the following key areas:

i. Existing Data Processes
Engage in a thorough analysis of legacy operating systems and databases that are currently related to the on-premises Insight and Relativity instances. Beyond identifying relevant system components present it is also necessary to identify all related APIs that connect both Insight to initial data lakes and contract repositories as well as Relativity which is leveraged for manual document review. To achieve this a deep application interdependency analysis must be performed.

ii. User Impact
Following migration, it will be necessary to train resources on how to interact with the new cloud instance. Management of this change and minimizing user impact can be achieved by having a comprehensive understanding of how users interact with different elements of the on-premises instance. This will assist in the migration strategy to:

1. Determine optimal time to migrate specific use cases/workloads accounting for business usage.
2. Provide insight into effective communication with relevant stakeholders in regard to potential service disruptions and relevant training for use of the new cloud instance.
3. Identify differences between existing and new processes helping inform the creation of relevant and targeted training for those affected.

iii. Governance Structure
In the context of a today's focus on corporate accountability, stringent regulatory requirements have resulted in complex governance structures to ensure compliance with such. Given the use cases for which Insight may be leveraged related governance around the creation and supervision of data is expected to be significant.

Moving to a Cloud instance will result in organizational changes relating to management of data for the Insight instance. Migration to the Cloud brings potential for Cloud service providers to aid in areas which previously would have been managed by internal IT teams. Incorporation of this third-party would require consideration as part of the migration strategy as it relates to the new governance structure.

If the Cloud service provider were not to be engaged in this way, relevant responsibilities and processes remaining internally would still have to be reassessed and documentation updated to reflect changes to workflows and maintenance activities.

b. Risk Identification and Mitigation

Identifying and creating mitigating strategies as it relates to both the implementation and operational running of the cloud instance is key to understand prior to engaging in a migration.

A risk management approach should make considerations of the following:

> ***i. Integrating cloud technology with legacy systems.*** Exisiting data likely originates off the cloud from pre-defined data lakes and it is assumed that data created on the cloud (within Insight) may also be transported back to legacy systems. There is inherent risk associated with this communication of data and the potential for failure in either direction.
> ***ii. Operational risks*** that may adversely impact the capability of the cloud instance once operational.
> ***iii. Testing requirements*** for the cloud instance to ensure the instance is operating as required and expected.
> ***iv. Risk mitigation approach.*** Requires running both the cloud instance and on-premises instance in parallel for up to two weeks and compare results daily before spinning down the on-premises instance.

c. Cloud Setup

As part of the gap analysis, it is necessary to decide on a ‘To-Be’ architecture as it relates to the cloud setup. There are four options in relation to the setup of the cloud instance and how it exists in relation to on-premises instances, the majority of which leverage a hybrid approach:

i. Hybrid Cloud with Workload Separation
Hybrid cloud with workload separation is a common approach whereby specific use cases are hosted either on-premises or in the cloud, depending on business requirements. Usually static or legacy workflows will remain on-premises while dynamic workloads are moved to the cloud. This allows for a level of risk mitigation as not all workflows are dependent on the cloud instance, however there are inherent on-going maintenance costs for the on-premises instance and data centers.

On-premises instance hosting

Cloud instance hosting

Analysis: If applied, workload separation should be reflected at a use case level. As referenced, this approach allows a level of risk mitigation resulting from spreading dependencies across different instances and will reduce the burden of migration planning as less use cases are being migrated. When engaging in this process, considerations should be made with regards to optimization of this setup, including the CLM module of the DocuSign Agreement Cloud and the implication of migrating existing models from one instance to another. Refer to the appendix for a detailed analysis of how to optimize a hybrid cloud setup with workload separation including these considerations.

ii. Hybrid Cloud with Workload Balancing

With hybrid cloud workload balancing, all relevant data is cloned between both the cloud and on-premises instance. This option is chosen by businesses where significant spikes are seen in activity at certain times putting pressure on either instance alone.

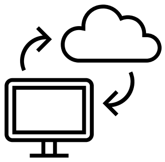

Workload balancing between
on-premises and cloud instance

Analysis: Given the potential vertical/horizontal scalability of a Cloud instance it is not advised to attribute significant value to workload balancing as an offsetting rationale for the resulting on-premises maintenance costs. vIt is also strongly advised that maintaining the same use case across both on-premises and cloud instances in tandem will be a duplicative effort that serves little purpose in the context of Insight, resulting in a significant burden on both model building and IT teams. When considering a hybrid setup use case, separation must always be applied which does not occur in the case of workload balancing. Refer to the appendix for further considerations surrounding use case separation as the best method to optimize the hybrid setup.

iii. Hybrid Cloud for Disaster Recovery

The hybrid cloud for disaster recovery is leveraging an on-premises instance for disaster recovery in case of failure of the cloud instance. Automation plays a key role in the success of this approach because when a disaster hits, shorter recovery time is critical. This disaster recovery approach involves precisely integrating several components like infrastructure, application and data synchronization, routing, security, and compliance and often can be expedited with automating your infrastructure and processes.

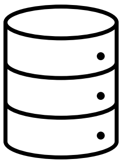

Disaster recovery instance

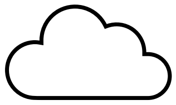

Cloud instance

Analysis: This approach involves on-premises maintenance costs where a secondary cloud instance could be used instead for disaster recovery purposes similar to the existing disaster recovery instance that exists currently on-premises. This approach also involves critical and detailed design activities on the entire infrastructure and process automation which also increases the time and cost of both implementation and maintenance.

iv. All-in on the Cloud
All-in on the cloud is the most straightforward approach whereby everything is moved to the cloud and on-premises instance is no longer leveraged. This approach presents less cost than the other options for this reason, though requires detailed understanding of the technology and associated processes.

Cloud instance hosting all use cases

Analysis: This approach appears to be the most cost effective because it is the only option that excludes the ongoing on-premises expenses. Given the complete dependency on the cloud, it is key that prior steps relating to gap analysis and risk management are carried out prudently to ensure the knowledge of existing processes is sufficiently accounted for. DocuSign as a Cloud service provider will be able to provide support in relation to Cloud maintenance assisting on the technology side.

Stage 2: Migration Approach

Stage two involves leveraging information gathered at stage one to make an informed decision as to how to implement the migration. Particular attention must be paid to the migration of data to ensure it is both successful and secure.

a. Full or Phased Implementation

It must firstly be decided how the migration will occur. Namely, will this be a full implementation (implement a full migration in one go) or a phased implementation (implement an incremental migration based on cost, benefits, and risk analysis). Given the potential for lessons learned is inherent in either phased approach, this is a preferable approach for migration to a full implementation provided there are no immediate time constraints.

A phased implementation has two approaches:

i. Priority approach – implement one use case at a time based on the immediate operational impact.
ii. Risk approach – migrate low risk data first to learn lessons and refine the approach for the next implementation where greater risk associated.

When considering the phased approach, migration of each use case will represent a phase, meaning, that at any one time any use case should only exist on one instance (either on-premises or cloud) following initial post migration checks. Attempts to continue to maintain a use case across two instances and sync separate environments would be a manual, highly labor and process intensive exercise that should be avoided.

b. Replatforming as the Migration Strategy

Insight via the cloud presents as Software as a Service (SaaS), meaning it is provided over the internet rather than via installation of software.

Benefits of Replatforming:

- *Cloud-native functionality* – better leverages the potential of the cloud and takes full advantage of features such as auto-scaling and infrastructure as code (IaC).
- *Cost effective* – as software is provided over the internet by the provider (DocuSign), no significant development work required internally.
- *Complements a conservative phased implementation* – facilitates a phased implementation allowing for testing and experimentation in the cloud instance before committing to a full migration effort.
- *Scalability* – load balancing service is included by the provider (DocuSign). This service improves resource utilization, facilitates scaling, and helps ensure high availability.
- *Automatic software update* – New versions of the software are updated by the service provider (DocuSign) regularly. That allows users to benefit from the latest upgraded features and new functions in a timely manner at no extra cost.
- *Disaster recovery and backup* – leveraging the provider's (DocuSign) standard SaaS offering, off-site backup is included, and faster disaster recovery can be expected with no extra cost.

c. Data Migration

Existing data must be migrated from the on-premises platform to the cloud instance. When engaging in this process, three key areas must be considered:

i. Initial Movement of Data to the Cloud from On-premises Data Center

a. The business must plan for transferring data to the cloud. This can be done via two methods, across a network or physical migration (i.e., moving storage devices). If delivery is to be via a network, then the business must ensure sufficient bandwidth is provisioned to ensure migration occurs within designated time constraints (which would have been understood during the gap analysis during stage one).

Either option must be complimented by relevant security checks to ensure compliance with relevant data governance requirements. Specifically, security controls should be put in place to protect data in transit via encryption and mitigation strategies should be in place in the case of accidental or intentional data loss.

b. Data storage capacity in the cloud instance must be assessed ahead of migration to ensure sufficient capacity exists for all incoming data as well as the ability for scaling to accommodate for creation of new data.

c. Data should be validated following migration to confirm the accuracy and completeness of the migration process via a comparison between the on-premises instance and the new cloud instance.

ii. Transport of Data Between On-premises Data Lake and the Cloud During Normal Operations

a. Plan to ensure the security of data both in transit and at rest during normal operations. Given insecure interfaces/APIs are seen as the biggest security threat to a cloud instance, this presents as an area of real concern.

b. Bandwidth and capacity management is critical for transport and storage of data to ensure adequate support for BAU requirements and future projected requirements (including unforeseen/unknown requirements).

iii.Ownership of the Data

a. If data resides on DocuSign servers, attention must be paid to the legal nuances surrounding ownership of the data. Consideration must be given to the legal implications resulting from either termination/expiration of the service contract or unforeseen events such as insolvency of the provider.

Stage 3: Post Migration

Following initial migration to the cloud instance, the final stage is to validate that the migration has been successful and then to train relevant stakeholders on the nuances of the new instance.

a. Re-test the Migration Pathway

Following migration of the data to the cloud, prior to rolling out the instance to the wider business, first, optimize cloud data lake settings and verify accessibility, permissions, and integrity of the instance.

This can be done through the following process:

i. Select a large data sample that is representative of different time ranges.

ii. Engage competent backend users who are familiar with relevant data to validate those historical values match results seen in the new instance before decommissioning the on-premises instance.

b. Temporarily Maintain On-premises Data Lake Company-wide Approval of Cloud Data Lake

Engage all relevant stakeholders involved in those related use cases to provide sign off on the validity of the new cloud data lake. The time this will take is variable and dependent on the quantity of related data and time taken to validate.

c. Educate Relevant Users and Stakeholders on the New Cloud Instance

i. Share regular updates around progress of the initial migration with external stakeholders to help socialize the move to the cloud which will help build confidence in the new instance.

ii. Following initial migration to the cloud, training relevant internal resources how to leverage the cloud is key to ensure a successful adoption as well as protecting the business from potential security risks derived from human error.

References

• NetApp, ‘3 Cloud Migration Approaches and Their Pros and Cons’ - https://cloud.netapp.com/blog/cvo-blg-cloud-migration-approach-rehost-refactor-or-replatform#H_H2

• ‘Migrating to the Cloud, How to leverage commerce tools and the Google Cloud infrastructure to increase agility and drive business’ - https://f.hubspotusercontent30.net/hubfs/4784080/Statamic/Files/commercetools-wp-migration-to-the-cloud.pdf

• Dynatrace, ‘Plan, execute and monitor your cloud migration for sustained success’ https://www.dynatrace.com/resources/ebooks/cloud-migration-aws/

• Cisco, ‘What Is a Cloud Migration Strategy?’ https://www.cisco.com/c/en_uk/solutions/cloud/what-is-a-cloud-migration-strategy.html

• ‘What is cloud migration? An introduction to moving to the cloud’ - https://searchcloudcomputing.techtarget.com/definition/cloud-migration

• Compliance, ‘On-Premises versus Cloud Computing’

• Cloud Security: Cloud Security Alliance, “Security Guidance for Critical Areas of Focus in Cloud Computing v4.0,” 2017 - https://cloudsecurityalliance.org/download/securityguidance-v4/

• NIST: Cloud Reference Architecture (NIST SP 500-292) - http://ws680.nist.gov/publication/get_pdf.cfm?pub_id=909505

Appendix

Q. How should SLA's and resulting user expectations be managed if operating a hybrid cloud as it relates to latency issues?

A. Network latency issues are a significant legacy hinderance on the existing on-premises instance which impedes the ability of internal teams to build models and extract data. The impact of such manifests in both a compromised model building process and a heavily burdened IT team (tasked with addressing the latency issue). Latency can be attributed to two factors, model building best practices and overburdened on-premises servers.
Moving to the cloud can alleviate both the latency issues as well as the resulting burden that is placed on internal IT teams. As the cloud service provider, DocuSign will manage the cloud instance and be best placed to rapidly address any latency issues as per the singular Service Level Agreement (SLA) that they will accede to.

When considering a hybrid setup whereby both on-premises and cloud instances are maintained, this will achieve the benefits set out above for those use cases migrated to the cloud but will continue to perpetuate such for any kept on-premises.

Applying a hybrid setup will result in further administrative difficulties in managing user expectations as it will be necessary to negotiate a combination of SLAs between both DocuSign and the User as well as internally between different areas of business. It would also likely increase the time associated with diagnosing where latency issues are occurring resulting in longer troubleshooting times. As a result, if such a setup is applied it is advised that special consideration be applied when defining remediation times within the relevant SLA representing a hybrid setup.

Q. When operating a hybrid cloud, what is the best way to optimize resulting capabilities?

A. If a hybrid cloud setup is pursued, to ensure the best outcome it is necessary to use a use case driven approach. Each use case will have its own unique requirements and resulting models, and these should be grouped as such.

Following such grouping, it can then be decided which use cases to migrate and in what order When deciding which use cases to migrate the following considerations should be made:First, at what stage is model building for the related use case as it is advised that model building should occur only in the environment that the documents are housed. In considering both migration and maintenance of the output of models from one single use case it is a best practice that model building and tuning ideally occur only within one instance for the following reasons:

1. **UDML models can only be tuned in the environment in which they are created.** Though UDML

models can be migrated between environments, once this occurs, they can no longer be tuned regardless of the perceived performance of the model. It is advised that in these cases it is better to rebuild any UDML in the new instance. Though there is an associated time and effort to achieve such this will allow for tuning to address unforeseen performance issues of the model.

2. **Migrated models may not perform the same in a different environment.** The cloud instance may be a later released version and DocuSign advises that the output of models moved between versions may not be alike. Further, documents represented in the cloud instance may not mirror the on-premises instance which may result in drift. This impacts the latent semantic scoring of models in the environment meaning extraction will differ from the original instance.

Second, to understand whether the output of the data extraction can be leveraged for document generation, we must understand DocuSign's product offerings. The CLM module of the DocuSign Agreement Cloud (DAC) only exists on the cloud as a SaaS offering, therefore related data extraction should occur in the cloud instance to facilitate the efficacy of the solution.