

Data Architecture at a Global Insurance Company

Background and Business Problem

A global insurance company had achieved success in operationalizing early data use cases, notably in its data innovation lab and operating entities. However, they realized that they were suffering from an inability to dynamically scale their new enterprise-wide shared service infrastructure to support a multitude of emerging real time use cases.

Specifically, our client wanted a blueprint for a data platform that allowed for scalability, data residency in specific named countries, and incorporation of additional advanced functionality over time (e.g. data streaming and real time processing). They also wanted to take advantage of cloud services so that they would not need to custom develop platform capabilities where possible. Our client needed specific recommendations to make this possible.

A global insurance company needed a design to easily, quickly, and dynamically scale a service infrastructure to support real time data science workloads up to 10 PB.

Silicon Valley Data Science advised the client on how to move big data workloads into the public cloud.

The Challenge


Needed the ability to easily scale their data science platform to enable emerging use cases

Wanted to alleviate the operational burdens of a custom deployment while maintaining flexibility and avoiding vendor lock-in

Didn't have a target architecture that would satisfy their key use cases and constraints in the cloud

UNDERSTANDING READINESS TO MOVE TO THE CLOUD

- **Use Case Applicability** - by use case, are there any specific implementation details/risks that must be addressed prior to migration?
- **Data** - due to the nature of specific types of data, are there any specific security restrictions that must be implemented?
- **Geography** - are there any specific geography restrictions that prohibit migration to the cloud?
- **Provider** - does the cloud provider (e.g. Amazon AWS, Microsoft Azure) meet key requirements for consideration for the cloud?
- **Global** - questions that pertain to any use case or cloud provider, and are **essential must haves** prior to moving to the cloud



© 2017 SILICON VALLEY DATA SCIENCE LLC. ALL RIGHTS RESERVED.

@SVDataScience



Solution

SVDS assessed multiple cloud providers across dimensions such as scalability, cost, security, maturity, and lock-in against our client's specific use cases and primary requirements for the cloud. To do so, we needed to first define a global reference architecture that met their use cases, was available quickly, was capital efficient, and was future-proof. All this while maintaining stringent standards around security and compliance.

The global reference architecture provided a blueprint for scoring each dimension against cloud provider specific "infrastructure as a service" (IaaS) and "platform as a service" (PaaS) capabilities. It also provided a roadmap for the client to know how to migrate their use cases to the cloud.

We provided the client with overall confidence, and with detailed recommendations on how to move big data workloads into the public cloud, in a way that met the client's security, compliance, and operational requirements on a worldwide basis.

Our Approach

SVDS developed a generalized cloud architecture based on an assessment of their data science use cases

We scored each Cloudera provider's IaaS and PaaS offerings against the reference architecture

We recommended which cloud provider's offerings best met our clients' use cases

New Capabilities

An architecture that would allow them to quickly scale up to petabytes in the cloud for real time data science workloads

An understanding of the costs, risks, and trade-offs associated with the target architecture

Ability to flexibly swap in and out cloud technologies and service offerings against the target architecture

