XONA

PARTNERS

Riding the Advanced Cloud Deployment Roadmap

Creationline, Inc. Team

Dr. Riad Hartani, Rolf Lumpe (Xona Partners)

August 15th, 2014

Table of Contents

1 Synopsis

The last few years have seen a select set of large scale Internet players position their Information Technology (IT) infrastructure as their core asset over which their services would ride, and as such have, for their in-house needs developed a lot of what would form the basis of the cloud as we know it today, with most of it coming out of their labs in the form of Open Source software and solutions – This including the various aspects related to IaaS, PaaS, Big Data and Data Sciences. Most recently and something that we will very likely witness over the next few years, a progressive move towards the use of cloud solutions will be the norm for a variety of corporations and service providers. This will be gradual, this will be on a need basis and this will be function of technology maturity and foreseen returns on investments. More importantly, this will require the emergence of advanced cloud centric product and services teams that could assess such migrations, develop them, deploy them and support them.

This is exactly where, Creationline, Inc. ("Creationline"), in collaboration with Xona Partners ("Xona") have set sights in terms of putting together a cloud specific transformational information technology practice to address these upcoming challenges.

Starting from our Tokyo, Japan and San Francisco, California's Silicon Valley head offices, we set sail for a journey around what we see in terms of Cloud Infrastructure evolution challenges, and highlight our evolving contributions aiming at overcoming them. This short positioning paper, is presented as a baseline for follow up detailed discussions related to the various topics under consideration, with the goal of designing, customizing and optimizing our solutions to lead data centric organizations' needs, leveraging the broad and complementary expertise of our team.

Specifically, the paper presents a comprehensive methodology, which includes the assessment of various models for cloud migration, the design and implementation related to building IaaS/PaaS/SaaS and porting applications to these environments, as well as the operational procedures required for a successful completion of cloud design projects. Along with this, an innovative cloud monitoring and optimization solution is introduced, with the aim of dynamically adapting cloud resources, based on the processing performance requirements.

2 Rationale for a Cloud Infrastructure Rollout Revisit

The fundamental premise of the question we are addressing is fairly simple:

As a business that is built around gathering large-scale data sets, mining and learning through such data, and optimizing communication between those producing it and those using it, where shall I go from here? This would touch upon the evolution of the compute, storage and network platforms over which data sets sit and business services processed.

Most information technology players, are, or will soon be considering enhancing their IT and data management infrastructure to help build data solutions which will optimize their ability in identifying, capturing, and managing data to provide actionable, trusted insights that improve strategic and operational decision making, resulting in incremental revenues and a better customer experience.

2.1 Cloud Migration Considerations

The current challenges of the existing platforms mostly affect the operations teams' ability to provide reliable SLAs for the reporting jobs that are critical for the business, the data platforms required scale, and perform effectively as the business grows. As such, the desired goal is to create a solid foundation architecture that is able to provide these optimal functional capabilities, and a platform to overlay additional applications such as intelligent business intelligence and Data Science as a service capability.

For most players, an insider-only approach, leveraging internal resources, would only go so far in terms of architecting and designing this new cloud based architecture, given the breadth in terms areas of expertise required, and more importantly, the need for a step-back and outside the box design way of doing things. As such, our team aims to be the strategic partner bringing in such expertise, and build open models tested and validated over a large set of data platform development models.

It is worth observing that:

- The expertise in deploying large scale cloud infrastructure still sits in the hands of the large cloud players themselves, and being able to predictably design it and build it, would most likely require the expertise of teams who have done it for their own IT within these large cloud players

- The Open Source software (such as OpenStack, CloudStack, Cloud Foundry, etc.) is still in constant mutation and likely to evolve fairly rapidly over the next few years, requiring specialized teams to bring in to market, deploy it, evolve it and manage it.

- The most crucial component of cloud migration is to figure out the right ROI model, based on which applications are migrated, how they are migrated and how they are used post migration. This in turn makes it primordial to figure out the right cloud model (public, private, hybrid) as well as the right framework to monitor such cloud deployments when done.

Given the above considerations, one can see that we are in the early days of large scale cloud migration, and the specialized expertise to do so is what is needed the most at this point in time. As such, we present our views on what methodology shall be implemented and what is likely to become the cornerstone of any related future cloud rollout roadmap.

2.2 Evolution to Advanced Cloud Infrastructure – Challenges

As of today, the existing information infrastructure and analytics processes suffer from challenges we have observed and worked on in the most typical large scale data and IT projects, some of which are listed below:

- The requirement to first progress the overall virtualization of computing and storage, and increasingly network resources. The increasing tie up between the virtualization and IaaS/PaaS environments renders the virtualization strategies very much dependent on a more forward looking broader cloud migration strategy.
- Understanding the variety of software applications and service running within the IT environment and analyzing them individually, then as an aggregate as far as ways of evolving them towards a cloud model
- Understanding the still evolving toolsets for cloud services monitoring and diagnostics of distributed software applications and compute processes
- Requiring a coordination between what would run in-house as a private cloud, or externally hosted private cloud off premise, with what runs over chosen public clouds
- Understanding how internal business processes shall evolve to accommodate such cloud migrations, and in fact, having that as a constraint that would force specific directions in such migration.

In the upcoming sections, we position our methodology, on how these cloud architectures should be evolving, short, mid and long terms. The analysis is presented in a generic fashion, but builds on top of a very selective and specific set of case studies that we have worked on in the real world, developed and completed. In other words, what is described is a pragmatic successfully completed case study, albeit made generic, to show broader applicability.

3 Cloud Infrastructure – The Road Ahead

Our approach to tacking evolutions towards cloud architectures is based on our understanding of the underlying business models, the existing IT and data architectures and deployment patterns and the assumptions we set, as far how data information models are structured and the underlying performance and reliability requirements. Additionally, our methodology and solutions aim at providing a capability maturity path for increasing capabilities of the system with minimal disruption to current operations, forming the basis for an evolutionary migration. As such, the architectural models we built our platforms upon, are designed to address some of the persistent problems in current infrastructure and future needs of the organization for data processing.

The figure below shows a standard component of a data center (DC) cloud infrastructure, including the compute, storage and network components.

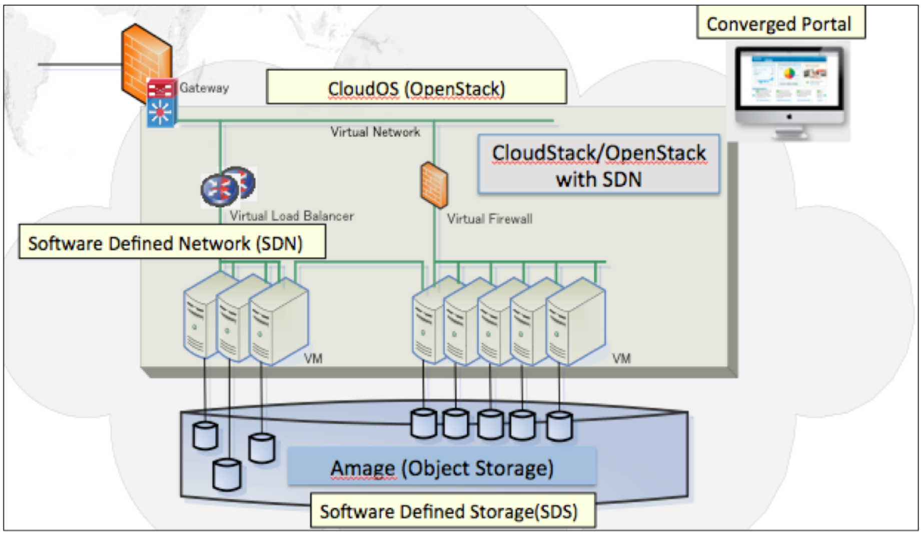

Figure 1: Standard components of DC cloud infrastructure

The figure below shows a complementary perspective, related to the various layers of the cloud implementation within an IaaS provider, and the various components required achieving it.

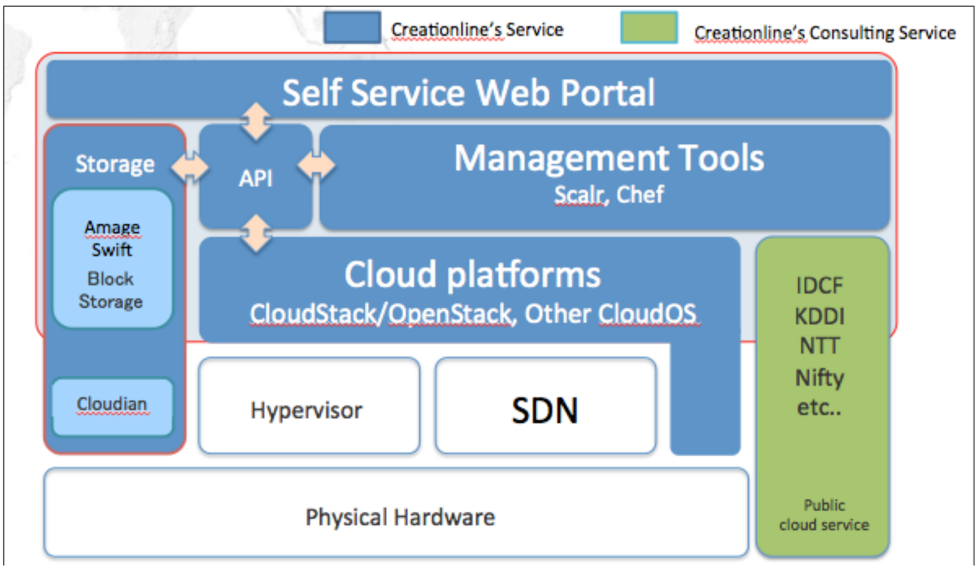

Figure 2: Cloud implementation for IaaS provider

Our methodology is based on four distinct steps. The first step is to have the target cloud infrastructure architected and implemented to allow any migration towards such target. Such design, which would form the basis of the IaaS and PaaS components, would be done with the right public or private cloud provider. The second step would approach the problem from the customer side of the cloud infrastructure, in other words, the corporation owning the IT and services that would want to run them over a cloud model. This would lead to the assessment of what would migrate, how and when based on the understanding of the IaaS and PaaS designs done prior. The third step would focus on supporting the migration and post migration, which would include performance and fault analysis and diagnostic. The last step would focus on overall optimization via appropriate orchestration and automation, and leveraging arbitrage models to select what to migrate on which cloud over time based on ROI dynamics.

The figure below shows the components of the methodology. It includes a non exhaustive list of technology partners with whom we are working to put such methodology into practice. Such eco-system is rapidly evolving and would need to cover the various options required to satisfy the diverse needs of cloud migration initiatives.

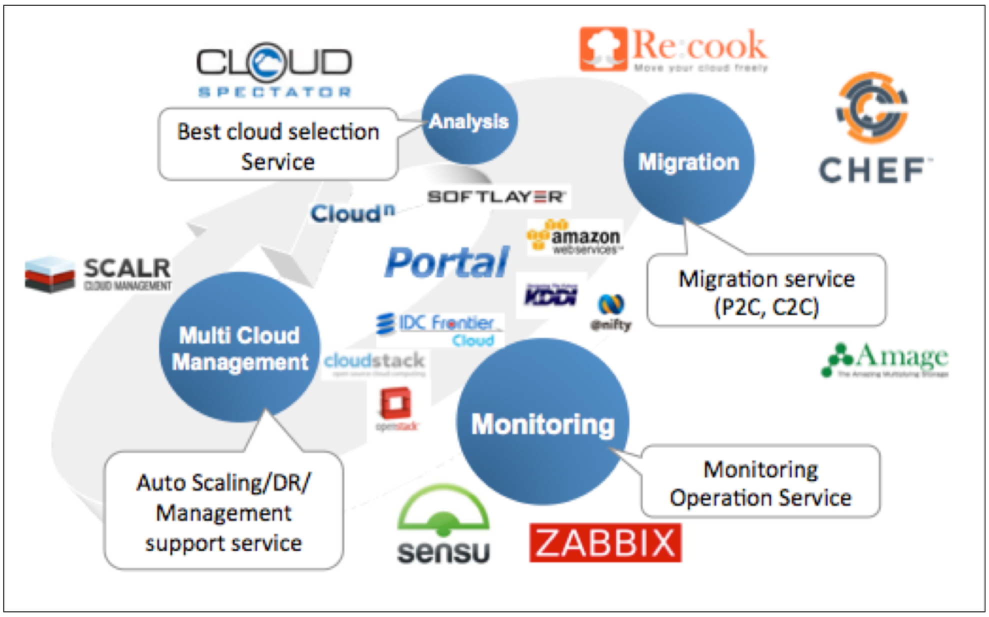

Figure 3: Components & methodology

We briefly describe these 4 steps and highlight the key aspects to look into while achieving them.

3.1 Understanding the design of the IaaS / PaaS component, as a leverage into migration trade-offs analysis.

Three pre-requisites are required to implement an adequate cloud migration strategy. First is a comprehensive understanding of the target virtualization, IaaS / PaaS environments, second is a tight working relationship with the cloud eco-system players into this environment and third an understanding of the IT and application software environment that would be candidate for migration and specifically its big data management component, as this would, in most cases, form the basis of the business ROI when migrating application and data management to the cloud.

These three aspects are briefly described.

Understanding The Virtualization, IaaS and PaaS designs and engineering considerations

As applications migrate to the cloud, the first thing to do is to figure out what to migrate, to what and how. Our thesis is that the best way to provide the best answer is to first have a detailed understanding of the design of IaaS and PaaS platforms and overall virtualized environments in the DC. Such knowledge can only be achieved through having led the design of these virtualization, IaaS and PaaS deployments.

Our experience having led the rollout of large IaaS solutions, with both CloudStack and OpenStack is what we build on to provide insight into what target IaaS and PaaS models would be optimal for various migrations, as well as laying out the right engineering and implementation models.

The figure below describe a high-level implementation example, where a CloudStack based IaaS and PaaS have been designed from the ground up to host enterprise applications in a telecom provider public cloud infrastructure.

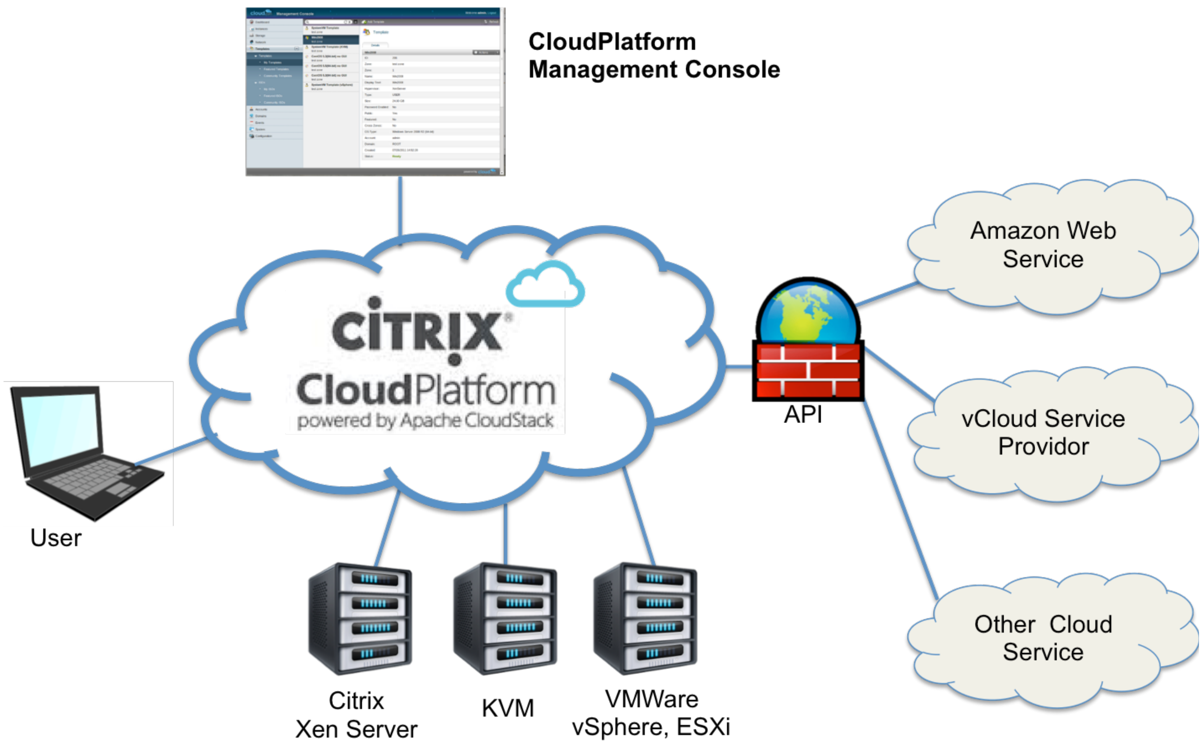

Figure 4: High-level CloudStack implementation (host enterprise applications)

In a similar way, below describes a high level OpenStack IaaS/PaaS implementation over which a high availability storage solution, complemented by a hybrid cloud disaster recovery solution as a service has been implemented in a Tier cloud services provider.

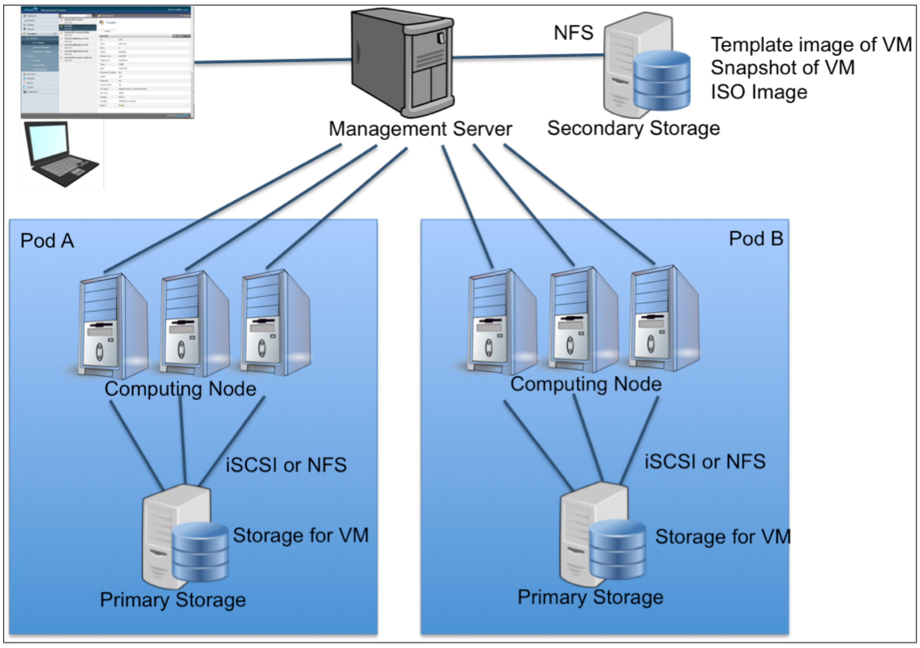

Figure 5: High-level CloudStack implementation (high availability storage)

Understanding the evolution of the partner eco-system

Understanding the various options and trade-offs in building the cloud infrastructure would require an eco-system of partners, where insights and rollout experiences are shared and contrasted. Below is a description of the partner solutions hierarchy, as well as illustrative parties within such eco-system that we have been working with. It is worth noting that most of these partner solutions' focus is in deploying, supporting and commercializing open source solutions. This list is fast growing with most partners being less than few years old, which is a testament of the novelty of the whole industry.

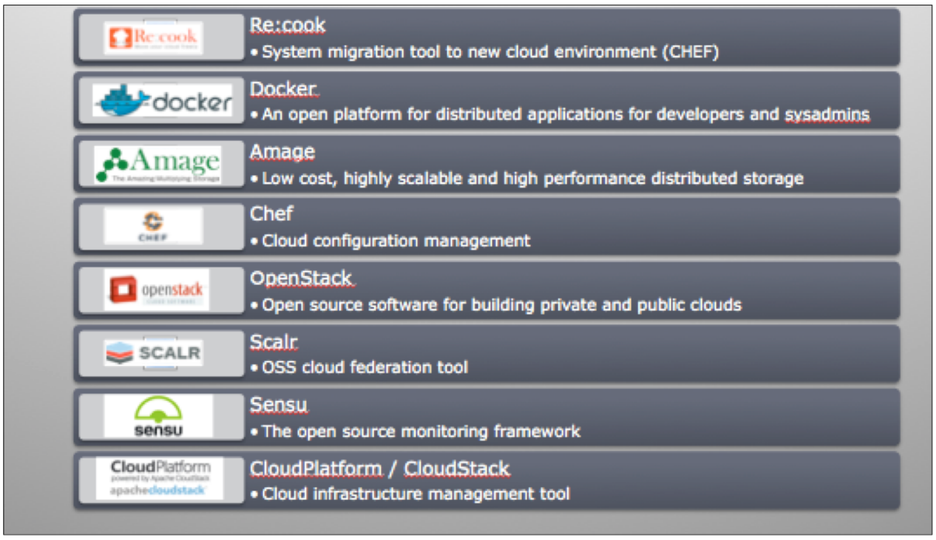

Figure 6: Partner Solutions

Understanding the Big Data Sciences Angle

Big Data Sciences is a combination of technology, business and mathematics that increasingly impacts every facet of daily life. The combination of traditional disciplines of data extraction, data intelligence, data analytics, data modeling, data warehousing, and reporting along with statistics and predictive analytics can be referred to as Data Sciences as illustrated in the diagram below.

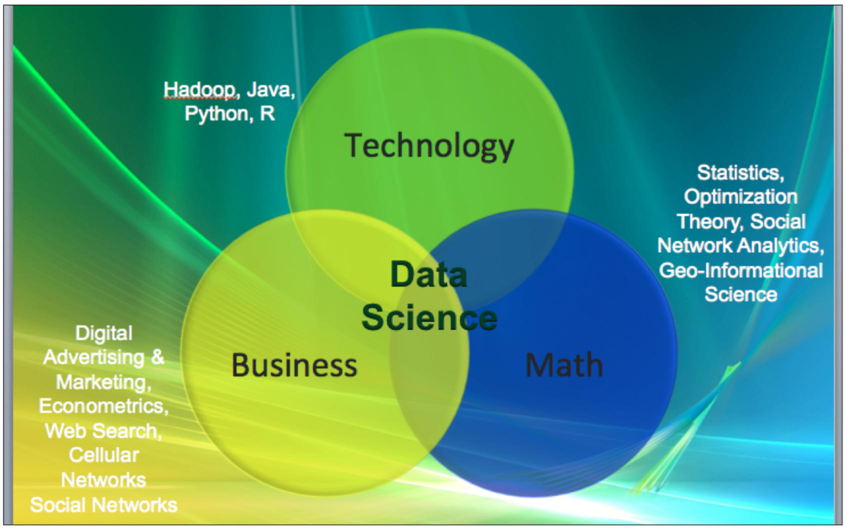

Figure 7: Big Data science angle

Overall, the Data Science requirements would direct link to the data management (based on open source or commercial frameworks such as Hadoop or the various SQL variations) and analytics derived from such data sets, which in sits on top of the cloud infrastructure. This in turn sets a lot of the requirements that one would have tackle in designing such infrastructure, based on the recursive logic of first understanding the end user application, the underlying data management with impact on the cloud implementation. A natural way of visualizing the various components of the Data Sciences hierarchy is shown below, taking it from data extraction at the bottom all the way to applications to specific verticals at the top.

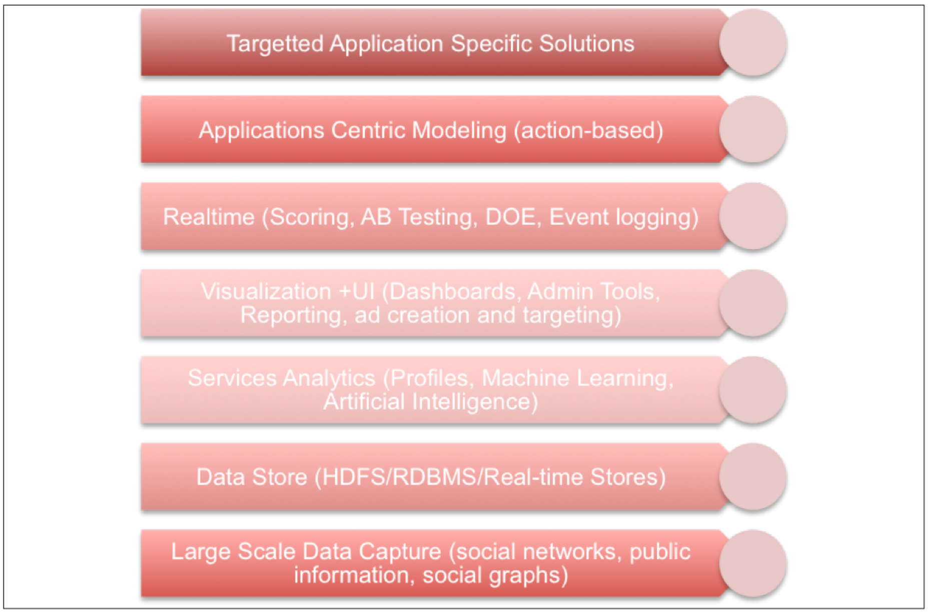

Figure 8: Big Data science angle

Based on the framework, which we have developed and put into practice on real case studies, specific credit transformational models will be highlighted. Below is an illustration of such cloud migration in the online advertising space where data is aggregated, processed and exposed by real time bidders for online ads (as Demand Side Platforms, Sell Side Platforms or Data Management Platforms). This is implemented over multiple steps, including the porting of data management tools into a Hadoop based management platform, which would sit within a private/public cloud environment.

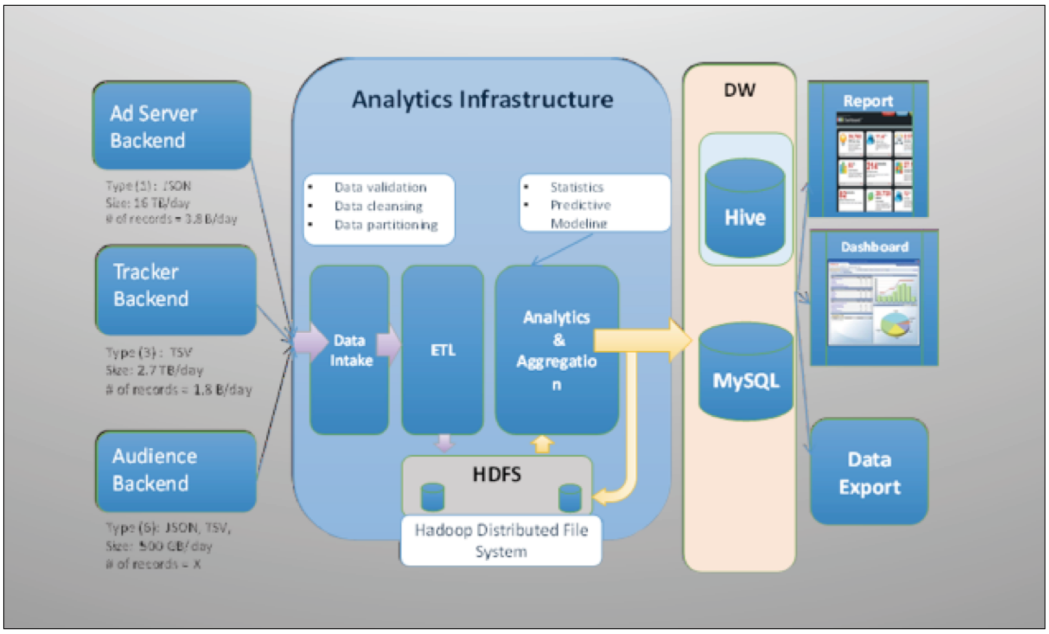

Figure 9: Analytics Infrastructure

It is this understanding of (a) the IaaS / PaaS environment, (b) the cloud eco-system players into this environment and how they are evolving, and (c) the big data and data science angle that would form the basis of the business ROI when migrating application and data management to the cloud that would form the cornerstone of the overall analysis.

3.2 Migration

Taking into account the specifics of the various private and public cloud models that one could migrate to, the next step is to understand the migration of the various applications and analyzing the underlying trade-offs, as described below:

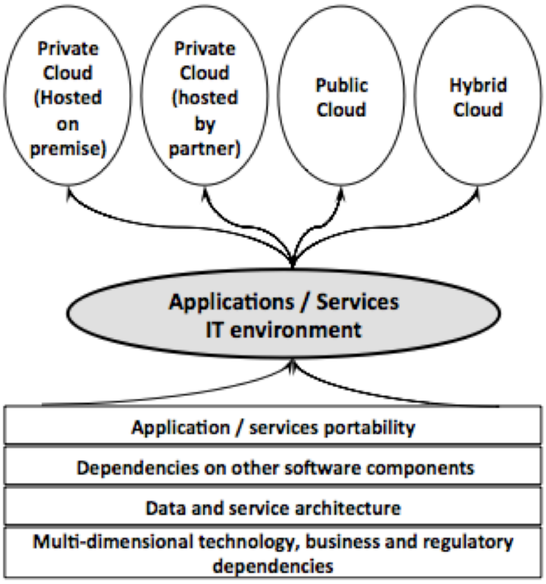

Figure 10: Migration methodology

This methodology shall result into a recommendation model in terms of migrating various applications based on their intrinsic requirements, as illustrated below:

Evaluation Criteria	Private Cloud (Premise Hosted)	Private Cloud (Partner Hosted)	Public Cloud	Hybrid Cloud
Agility / Migration	Medium	Medium	Medium	Low
CAPEX / OPEX	Medium	High	Low	Medium
Application Portability	Low	Low	High	Low
Security / Privacy	Low	Medium	High	Medium
Operational Complexity	Low	High	High	Medium
Other Criteria (function of context)	Ranking	Ranking	Ranking	Ranking

Table 1: Migration model

3.3 Monitoring, Diagnostic and Action models

Once migration is executed, the next step is to monitor its evolution and analyze the live performance, reliability and quality of service requirements. This would include the deployment of monitoring tools, interfacing with the resource management and orchestrator tools and having access to preventive and corrective action models, primarily via the automated provisioning tools, as illustrated below.

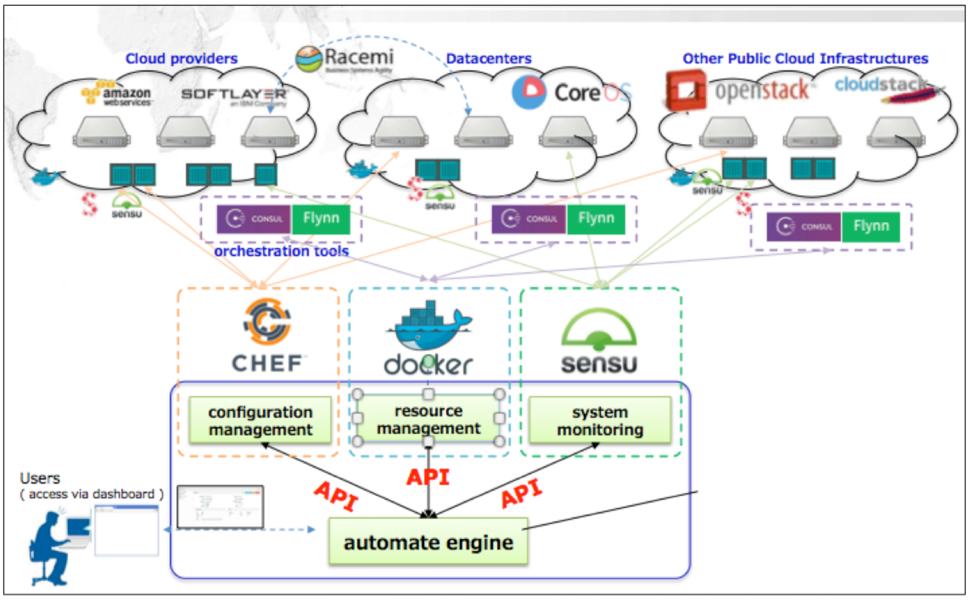

Figure 11: Cloud intelligent monitoring engine

The monitoring functionality is an essential feature for any system relying on various degrees of automation. The various system components work together utilizing the monitoring and automation APIs. Given that the various cloud infrastructure platforms and tools support the RESTful API, The proposed includes a comprehensive monitoring system, which dynamically interfaces with the various other cloud management tools. This intelligent monitoring and automation tool, which has been developed, based on the most stringent reliability and performance requirements of some of the most advanced cloud environments is described in the following section.

3.4 Optimization

As the application are ported in the cloud, and the multi-layer monitoring, including the IaaS and PaaS environment, as well as the application software and services environment, the last step is to put in place a dynamic optimization model that would analyze next actions to implement. This would include actions for optimizing reliability and predictability, as well as overall cost dynamics leveraging the choices offered by the various players in the eco-system, including data centers and cloud providers, systems integrators, support partners and the remaining players. This is mostly done via a cost arbitrage function that would dynamically recommend ongoing application migration options to select cloud environments, as described in the figure below:

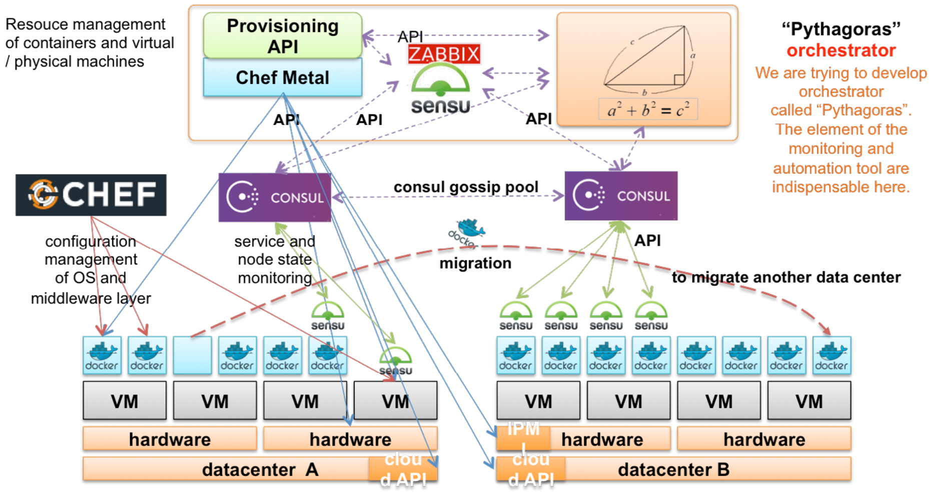

Figure 12: Dynamic optimization model

The above illustration describes an architecture where the cloud compute, network and storage resources are managed by the intelligent orchestrator, names "Pythagoras". This not only includes the datacenter underlying infrastructure but also the hardware resources, virtual machines' resources and "Docker" container resources. The orchestrator integrates "Chef Metal" and "Consul" for resource management and execution, based on the active monitoring and the respective data acquired by "Sense agent and APIs". The "Pythagoras" orchestrator enables a highly reliable and a fully optimized automated cloud operation.

4 Conclusions & Call for Partnership

Following some successful validation over the last few years, where some of the largest cloud and big data transformation projects have been conducted, involving designs with some of the most aggressive scaling, reliability and manageability deployment requirements, we are now in the process of taking in-house development and deployment methodologies to the larger market, and would welcome discussing specific requirements with key IT and cloud solutions architects having for a mission to lead their IT transformation architectures, as well as with managed services players wanting to build on their existing IT and big data capabilities and augment it with specific cloud based data management platforms.

Specifically, the analysis models developed for cloud migration assessment, the software tools and processes in use for the IaaS/PaaS and SaaS design and operations, as well as the intelligent monitoring, automation and orchestration tools for an optimal cloud operation would provide us with an optimal starting point in analyzing the specifics of the cloud architecture, and ensuring its successful deployment.

We believe that the next generation cloud and big data platform architectures will be evolving in the direction we have been highlighting, and hence, encourage various players to speed up such evolution, for the common interest of the various eco-system players.

5 Acronyms

DC	Data Center
HSFS	High Sierra File System
HDFX	Hadoop Distributed File System
IaaS	Infrastructure as a Service
IT	Information Technology
PaaS	Platform as a Service
RDBMS	Rational Data Base Management System
ROI	Return On Investment
SaaS	Software as a Service
SLA	Service Level Agreement
UI	User Interface

6 The Creationline Team

Creationline was founded in 2006 and developed over the past 8 years as Japan's most advanced cloud enabling and professional service company. Initially providing consultation services to major Japanese telecom carriers and IDC firms building cloud service infrastructure based on open architecture Cloud such as CloudStack and OpenStack, as well as big data architectures such as Hadoop, and large scale PaaS infrastructures such as Cloud Foundry. Services include Proof Of Concept (POC), design, implementation and support. Nowadays Creationline service set includes migration (P2C, C2C), monitoring & operations, multi-cloud management and cloud building services.

Sample relevant globally recognized projects completion includes Softbank Japan Cloud Services Design, KDDI public and private hosted cloud services design, NTT Cloud Services architecture as well as various cloud & Big data engagement with select data center providers.

7 The Xona Partners Team

Xona Partners (Xona) is a boutique advisory services firm specialized in technology, media and telecommunications. Xona was founded in 2012 by a team of seasoned technologists and startup founders, managing directors in global ventures, and investment advisors. Drawing on its founders' cross functional expertise, Xona offers a unique multi-disciplinary integrative technology and investment advisory service to private equity and venture funds, technology corporations, as well as regulators and public sector organizations. We help our clients in pre-investment due diligence, post investment life-cycle management, and strategic technology management to develop new sources of revenue. The firm operates out of four regional hubs which include San Francisco, Paris, Dubai, and Singapore.

Xona Partners

www.xonapartners.com

advisors@xonapartners.com