SAS® Viya® and the cloud: How SAS is changing the game it invented

A dbInsight white paper for SAS

Trigger

Recent events have underscored how suddenly the competitive and customer landscape can change. Businesses require the ability to make *smart decisions* based on *trusted data* that generate *reliable outcomes.* Analytics is the key for enterprises to understand their landscape and make the best decisions to maximize business value. Recent events are also underscoring the growing adoption of cloud, not just as a deployment target, but for adoption of a new generation of cloud-native services that dissolve traditional application and data silos, while providing ready access to the right amount of compute when it is needed.

However, most enterprises are facing compute bottlenecks that affect analytics. On one hand, IT organizations are overtaxed in keeping essential systems running, with few resources to spare for deploying and maintaining software. Conversely, everyone from business analysts to data scientists has access to an unprecedented richness of data, tools, and frameworks. There is more data being generated than ever. Business analysts are empowered by self-service tools to produce their own analysis without the need for constant IT support, while their counterparts in data science are reaping the benefits of innovation – much of it coming from the open source world – to take advantage of the tools, frameworks, or algorithms of choice to develop sophisticated predictive models. Yet, all too often, the data is silo'ed and so are the models, which often fail to make the move from data scientist laptops to production.

Our Take

Data, tools, and algorithms alone won't drive success; it is the *decisions* that enterprises make, based on *trusted data* and resulting in *trusted outcomes*, that will spell the difference. That requires tooling and technology that enable enterprises to apply governance to all aspects of the analytics life cycle, from data ingest and transformation to model creation and deployment and the process of making decisions. Because of the complexity of data, rapidly changing business needs and the unpredictable computing demands involved with all aspects of modeling, organizations need ready access to compute and storage that can scale.

The cloud provides an opportunity for enterprises to gain access to the necessary computer power and scale of storage. It also provides an opportunity to direct their focus away from infrastructure to problem-solving. For enterprise software providers, the cloud offers an opportunity to modernize their portfolios, transitioning from monolithic, often silo'ed applications and tools, to end-to-end solutions that are refactored into services that are deployed as lightweight containers and continuously updated with no disruption to the customer. And, thanks to emerging standards such as Kubernetes, they can deliver solutions that fully take advantage of the flexibility, elasticity, and scalability of the cloud while avoiding lock-in to any single cloud platform. In the latest release, SAS has modernized SAS

Viya to take advantage of containers, Kubernetes and other open industry standards to deliver a cloud-native solution that will enable its customers to run on the cloud of their choice.

The analytics bottleneck

When it comes to data and analytics, enterprises may have too much of a good thing. Data, which feeds analytics, is proliferating; according to some industry estimates, the sheer volume and variety of data is doubling every year or two. Meanwhile, access is becoming democratized as, at one end of the spectrum, line-of-business users and business analysts are gaining access to self-service tools empowering them to generate analysis without having to rely on IT. At the other end, data scientists and business analysts are gaining access to a wealth of models, frameworks, and algorithms, many of which are coming from the open source community, and the emergence of AutoML tools that are demystifying machine learning.

The question is, are enterprises realizing better business value? All too often, teams in different parts of the organization are gaining access to data on an ad hoc basis, resulting in islands of analytics. Data and analytics are often maintained by individual departments or lines of businesses and are not governed. Models that are being developed by data scientists are not successfully moving from the sandboxes on their laptops into production. When models do enter production, efforts to track alignment of data and algorithms are frequently conducted manually. And if implemented on-premises, the burdens of keeping software versions up to date and patched falls on IT, which in most organizations is backlogged. The result is that the value that enterprises are gaining is not consistent.

Decisions, not data alone, drive success

The challenge is that data, tools, and algorithms alone will not deliver value; it is the *decisions* enterprises make that will spell the difference. That requires *trust* in the data and models, and the decisions that are made from them.

The first step is taking a full lifecycle approach to managing models, from creation to development, and monitoring of performance and reliability whole in production. The final piece is the decision; ideally, decisions should not be made on an ad hoc basis, but should be driven by analytics-led insights and in a manner that is consistent and supports the rules, policies, and goals of the organization.

Automation is critical for productivity

In most organizations, practitioners spend more of their time cleansing and preparing data than analyzing it; they require automation that can make the process less time-draining, comprehensive, and efficient. Another large chunk of time is spent with the mechanics of the

modeling lifecycle, including feature engineering, algorithm matching, model training, and monitoring. This area of automation is often called to as "AutoML," where the system helps the data scientist become more productive, allowing them to spend more time on the higher-value tasking of developing, identifying, and continuously refining the best models. Automation can also aid the interpretation of results where results from queries can be communicated in plain, easy-to-understand business terms to a broader audience through natural language processing (NLP).

Governance is becoming paramount

Given the power of analytic models, the entire process-- selection- preparation, and consumption of data development of models and choice of decisions should be tracked from cradle to grave. That is the essence of applying governance to the analytics process, which provides accountability.

Governance starts with analytics teams should understand the reliability of the data. That requires tracking data lineage to identify its source and currency and to track how the data has been transformed and consumed, whether there is bias, and by which people and which models. This is critical given existing – and emerging -- mandates for data privacy, confidentiality, and retention. Today, more and more organizations are becoming subject to strict guidelines on who can see the data and whether that data should be anonymized and/or obfuscated through masking or encryption. Models should be governed through the entire lifecycle, from development to deployment, for parameters such as bias and policy compliance.

Deployment is a huge obstacle

While many organizations build powerful analytic models, only about half of them ever make it into production. That is a huge wasted effort by data scientists. And when models actually make it into production, all too often the entire process drags on for months. And in many cases, the results may not be useful as the model fails to predict the outcomes that address the business problem. Not surprisingly, enterprises are not getting the right results, and consequently, are not getting the full value from their investments in time and talent. There are numerous reasons why models might not see the light of day. For instance, data scientists usually don't have the skills of data engineers, DevOps specialists, or developers. Not surprisingly, they may lack full understanding of the data sets or infrastructure, leading to algorithms that are misaligned to the way that data sets or compute clusters are laid out. Furthermore, there may be problems when developers port models, written in languages such as Python or R, to established production languages such as Java or C++. These are problems that often come as unpleasant surprises to data scientists who see the results of

their work literally lost in translation. And even if models make it to production, performance degradation may occur over time, especially if the characteristics of the data sets change.

Go where the developers are

Everyone, from business analysts to developers and data scientists, has their own preferences for tooling languages and/or algorithms, and for good reason. The SAS programming language has attracted a large skills base over the years because of its functionality and access to a rich and deep set of libraries and modeling capabilities. In recent years, the open source world has spawned significant innovation in areas such as neural network frameworks, with languages such as Python providing its own sets of libraries for building machine learning algorithms. In many cases, data scientists may build sophisticated machine or deep learning programs through mixing and matching languages, including SAS, Python, Java, and others, based on the libraries or frameworks they support.

Practitioners also have their own preferences when it comes to workspaces, as many are embracing notebooks such as Jupyter and Zeppelin for organizing their work. Analytic solutions today must enable data scientists, analytic solutions to work with the languages, APIs, and frameworks with which they work.

Figure 1. Evolution of cloud deployment

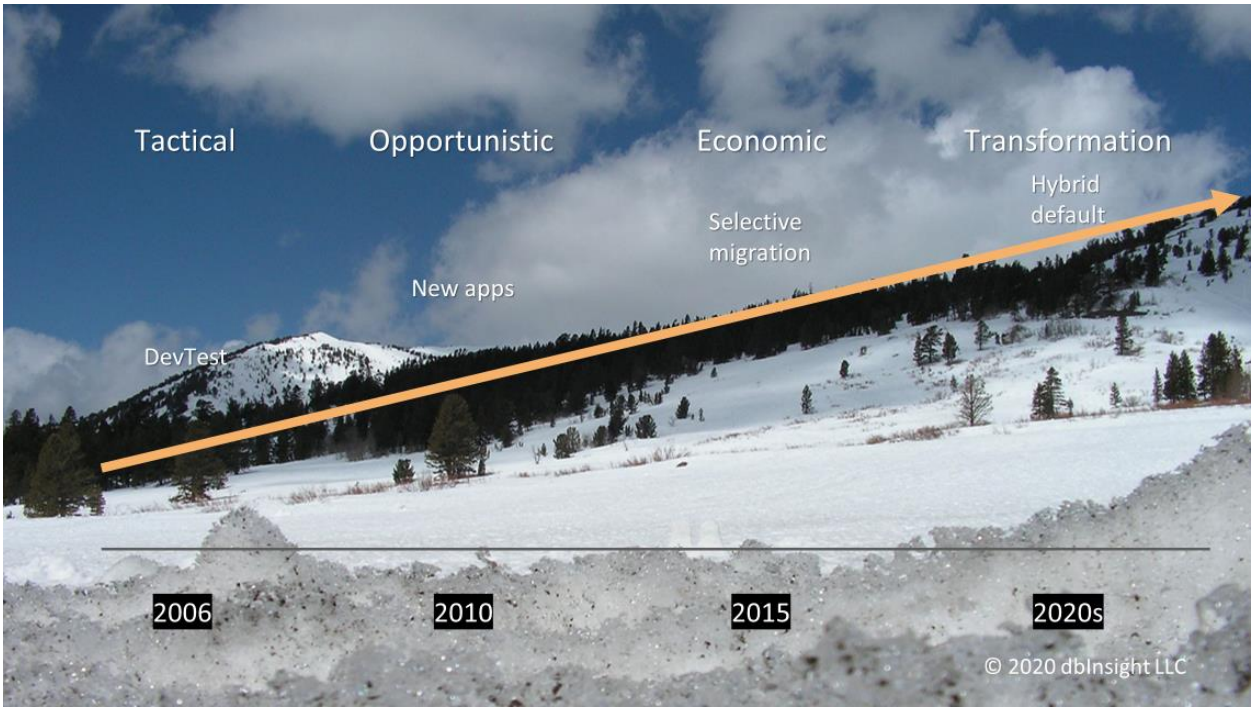

Source: dbInsight

Enterprises are embracing the cloud

Welcome to the Hybrid Default

The COVID-19 pandemic has accelerated a trend that was already well in progress: the growing embrace of the cloud. The rapid upheaval in the economic landscape and the expectation of virtual work becoming the norm have further accentuated demand for the flexibility and reliability of cloud-based solutions.

As shown in Figure 1, most enterprises already have had experience using the cloud. According to Flexera's ninth annual survey of cloud use, 98% of respondents report using at least one public or private cloud.[1] After getting introduced through tactical and opportunistic use cases such as TestDev or development of new mobile apps, enterprises are now looking at the cloud to simplify IT and apply new flexibility to addressing their core business-critical systems. Enterprises are seeking to take advantage of the operational simplicity, flexibility, agility, and fast time-to-benefit that cloud-native deployment enables. They are also capitalizing on the leading-edge tools- such as streaming, data pipeline, or machine learning services -that are only available in the cloud.

Consequently, decision-making for deployment of enterprise systems has changed. Until recently, the default criteria for deployment of enterprise systems was on-premises, with cloud deployment requiring justification (e.g., data gravity, the availability of services such as AutoML, etc.). In the 2020s, we see this situation reversing: The default assumption for new or modernized systems will involve cloud deployment unless a special exception is made. This does not necessarily mean that all new or modernized systems will run in a public cloud. However, regardless of whether enterprises deploy their systems in a public cloud, or in a hybrid or private cloud on premises, there is one common thread: they expect to gain the operational simplicity, flexibility, and agility made possible by the cloud control plane. Welcome to the new era of "The Hybrid Default."

Why cloud-native is critical

Not all cloud deployments are alike, and neither are the benefits. There are two pathways to cloud deployment including:

- Lift and shift, where existing workloads are moved as-is to a public cloud. Here the benefits are budgetary (line items costs are moved from capital to operating expenses, simplifying budgeting), and elimination of the need to acquire and maintain physical infrastructure.

[1] Flexera 2020 State of the Cloud Report

- Lift and transform, where workloads are reimagined and implemented in a cloud-native architecture that separates compute from storage; monolithic applications are refactored into services that are deployed in lightweight containers (instead of virtual machines); teams embrace DevOps practices, including continuous integration and continuous deployment (CI/CD); tasks such as cluster orchestration can automate placement of containerized workloads.

In a cloud-native architecture, organizations can exert better control over their costs by taking full advantage of the ability to ramp up and ramp down compute and storage based on actual usage and SLA requirements), while streamlining software maintenance and upgrading.

For solution providers, cloud-native architecture offers the chance to reimagine their applications. No longer tied to monolithic application portfolios, where new capabilities often surface as separate programs or tools, cloud-native architectures enable them to refactor their portfolios as microservices that can break down the functional barriers, and containers that provide the runtime. And, as they transform their applications into microservices, solution providers can more rapidly update their offerings without the headaches of breaking the complex web on interdependencies endemic to traditional monolithic software.

Kubernetes is king

Kubernetes (K8s) is an open source technology that has emerged as the de facto standard for managing and orchestrating containers, having been adopted by virtually every major cloud platform provider. While containers can modularize control of the processes and virtualize them from the underlying infrastructure and operating system, K8s provides the means for automating the operation of containers. K8s provides the infrastructure for services to talk to and discover each other, arranges access to storage, and streamlines the ramping up and down of the use of compute resources. While each major cloud platform may implement Kubernetes differently, for third parties such as SAS, support of Kubernetes means that they do not have to reinvent the infrastructure for managing the control plane for deploying and managing the lifecycle of containerized microservices.

What cloud-native means to SAS Viya

Addressing the analytics bottleneck

SAS Viya was designed to address the analytics bottlenecks that are all too often the reality for most large enterprises. It addresses data and analytic silos that have prevented departments and lines of businesses from collaborating and/or leveraging each other's' data and analytic models. It manages deployment, smoothing the path for models to make it from laptop to production. With SAS Viya, analytics can be embedded into your applications using

public APIs that are called only when needed, saving you time, resources, and reducing the tradeoff between speed and development excellence.

SAS Viya automates managing of the modeling lifecycle. It automates tasks ranging from data cleansing to data transformation, feature engineering, algorithm matching, model training and ongoing governance (often referred to as Automated Machine Learning or AutoML). The automation provided by SAS Viya makes it easier for both business analysts and data scientists to use SAS software. Results from queries are explained through NLP so data analyses are communicated in plain, easy to understand business terms.

SAS Viya was designed to meet all users where they live. It empowers business analysts, who expect self-service capabilities, not only to visualize data but to perform essential tasks such as data preparation, decisioning, generating forecasts, performing text analytics, and other functions. SAS Viya also accommodates the expectations of data scientists to get the best of both worlds: the ability to work with the tools, languages, frameworks, and algorithms of choice through open APIs, while getting access to the management and automation of SAS Viya with capabilities such as automation of pipeline building, feature engineering, simulation, and operationalization. SAS Viya was a rethink of the SAS portfolio, providing a new functional approach that integrated capabilities from multiple products into a seamless workflow; one of the earliest results was SAS Visual investigator, which utilized capabilities from SAS Decision Manager, SAS® Enterprise Miner™, and SAS In-Memory Statistics.

It offers these capabilities in a fully governed environment for delivering trusted outcomes through comprehensive data lineage and auditability, along with model lifecycle management, bias detection, and explainability.

How containers, microservices, and Kubernetes make SAS Viya more flexible

By moving to a fully cloud-native architecture, SAS Viya is on the path to becoming fully containerized, supporting industry standards. It has been redesigned for running on each of the major public and hybrid cloud platforms: AWS, Azure, and Google Cloud Platform, plus support for hybrid and private cloud via Red Hat OpenShift. SAS Viya will run across a variety of cloud implementations, from multi-tenant to single tenant bare metal. The service will be continuously updated by SAS, ensuring that customers always have access to the latest functionality without disruption. And unlike traditional on-premises enterprise software, users can trigger SAS Viya services without the need for separate logins or switching to different programs.

By moving to the Kubernetes ecosystem, SAS Viya can take advantage of numerous tools and plug-ins that add administrative functionality (e.g. logging, monitoring, security, etc.) outside

of core SAS technology. Furthermore, Kubernetes support enables SAS Viya customers to integrate with other commodity cloud offerings like security, identity, and compliance; databases; development and integration tools; storage; container management and orchestration; management consoles so they can leverage the core services that they may be already using.

The result is that customers will gain fine-grained control across cloud spend and service levels and have access to scalable compute at predictable service levels, without having to be tied to any single cloud platform.

Takeaways

Before the latest upheavals to the economic landscape, enterprises were already on a path to embracing the cloud. The trigger has been the need to refocus on their core competitive challenges; with the changes to the landscape, the path to the cloud will accelerate. For most enterprises, the cloud will present a transformational opportunity to rethink how they manage their own business, and how they use technology. The cloud will not be business as usual; while enterprises could simply “lift and shift” their current systems to the cloud, we believe that most will embark on an approach to steadily transform their operations. They are seeking the operational simplicity and agility that the cloud control plane promises, regardless of whether it is implemented in a public cloud or in a hybrid or private cloud on-premises. And in so doing, they have similar expectations for their business-critical applications. The move to the cloud is not just about infrastructure, but about how they run their business.

SAS as a business has taken this transformation to heart. The pivot to cloud-native development for SAS Viya is more than a technology change, but a change to the way it conducts business and serves customers. In place of the traditional, more waterfall-like releases, SAS is embracing practices of continuous integration and continuous deployment of the DevOps world. In this way, SAS can more readily respond to the needs of its customers for new functionality or new integrations. Because SAS Viya software is transitioning to a container-based, microservices architecture, SAS can make these changes non-disruptively and roll them out continually.

Cloud-native architectures, where applications are reconceived as containerized services, can deliver on the full promise of the cloud's operational simplicity, agility, and ability to scale. The emergence of K8s, as an open source technology that has become the de facto industry standard for cloud-native containerized environments, provides a roadmap for solution providers to deliver their offerings that will run across multiple clouds – providing freedom to the customer on the cloud to use. With the latest release, SAS has put its next-generation analytics platform, SAS Viya, on a fully cloud-native path that is based on standards and open

APIs. By moving to a cloud-native implementation, SAS Viya will take full advantage of the cloud's scalability, providing a solution that can readily be integrated into other cloud-based solutions, while delivering the assurance that customers are always using the latest, up to date capabilities.

Author

Tony Baer, Principal, dbInsight

tony@dbinsight.io

Twitter @TonyBaer

About dbInsight

dbInsight LLC® provides an independent view on the database and analytics technology ecosystem. dbInsight publishes independent research, and from our research, distills insights to help data and analytics technology providers understand their competitive positioning and sharpen their message.

Tony Baer, the founder and principal of dbInsight, is a recognized industry expert on data-driven transformation. *Onalytica* named him as one of its Top 100 influencers lists for data and cloud in 2019 and 2020. *Analytics Insight* named him one of the 2019 Top 100 Artificial Intelligence and Big Data Influencers. His combined expertise in both legacy database technologies and emerging cloud and analytics technologies shapes how technology providers go to market in an industry undergoing significant transformation. His regular ZDnet *"Big on Data"* posts are read 25,000 – 30,000 times monthly.

dbInsight® is a registered trademark of dbInsight LLC.

111484_G129210.0720