WHITE PAPER

INTELLIGENT BUSINESS STRATEGIES

The Business Value of Kubernetes in An Analytical Environment

By Mike Ferguson
Intelligent Business Strategies
September 2019

Prepared for:

Qlik

Table of Contents

THE INCREASING DEMAND FOR BUSINESS INTELLIGENCE AND ANALYTICS

Data, BI and analytics are seen as critical to running a successful business

In most organisations today, data, business intelligence (BI) and analytics have become strategically important with many executives now saying how critical they are to the running of their businesses. So much so that demand to consume the data, BI and analytical services is coming from every department, every core application and every business process. It seems we are now in an era where a data and analytical hub is emerging at the centre of the enterprise (see Figure 1) with data sets, queries, reports, dashboards, stories, machine learning models and analytic applications all being made available for business users to consume and use to help deliver business value more rapidly.

Data, BI and analytics have moved to the centre of the enterprise with demand coming from every part of the business

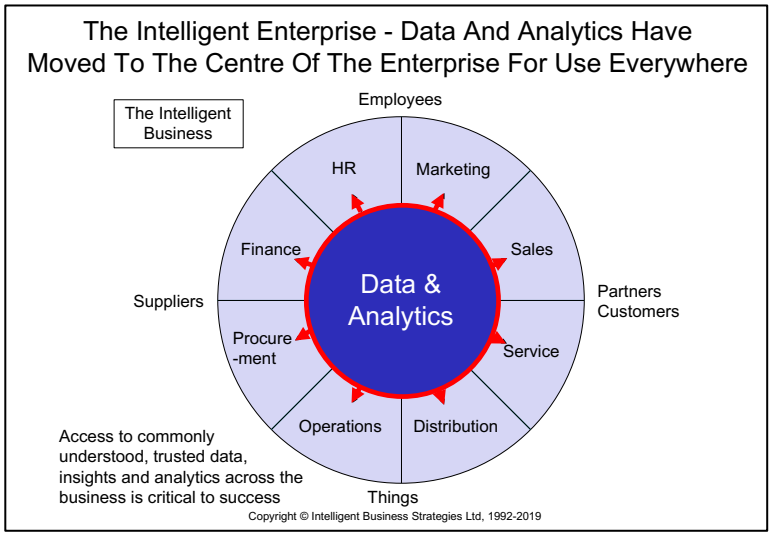

Figure 1

External users are also being given access

In some cases we are making these insights available to external users such as customers, suppliers, partners and even monetising them.

The data landscape is becoming increasingly complex with new data sources and data stores in a hybrid computing environment

Interestingly, this is happening against a backdrop of a more complex data landscape where digital transformation has caused an increased demand for more data. New internal and external machine and human generated data sources are appearing with data streaming into the enterprise in real time or being captured and stored on many different types of data store on-premises and in multiple clouds. This is giving rise to a significant growth in data volumes.

So several trends are occurring.

The trend is towards more users, more data and ready made data and analytics to shorten time to value

- Universal demand for BI and analytics across the enterprise and beyond which translates to many more internal and external users
- A huge thirst for new data to drive business value
- Increased demand for 'ready-made' consumable data, BI, predictive and prescriptive (e.g. recommendations) services to shorten time to value

However, it is not just more users who are lining up to make use of data and analytical services. Both operational transaction processing applications and managerial (e.g. planning) applications also want to utilise these services on-

Companies also want sparter processes

demand to help create smarter, more optimised business processes. Some are even going beyond this to create AI infused, self-learning applications. So, the march towards the truly intelligent business continues to progress but at a pace that, if anything, is quickening.

Given this situation, what then are the implications of these trends on BI and analytical platforms? And are there any other technologies needed to ensure that BI and analytics delivers the business value needed? Let's take a look.

THE IMPLICATION OF GREATER DEMAND ON BI AND ANALYTICS PLATFORMS

Pressure is growing on BI and analytical platforms to be able to support these trends

There are a number of implications on BI and analytics platforms that need to be considered if they are to become core technologies in a data and analytics hub at the centre of the data driven enterprise (as shown in Figure 1). These are:

- With demand coming from right across the enterprise and beyond, simpler user interfaces are needed to broaden the usage of BI and analytics so lesser skilled workers can also contribute to delivering value

More users as BI gets easier to use

- As this happens, more users and applications will start to leverage data and BI causing the number of requests for BI queries, reports and dashboards to grow rapidly. It will grow even faster if external users and applications are also granted access to these services. The implication here is that BI platforms have to scale to support a major growth in concurrent users so that they can deliver value across the enterprise

More applications want to embed BI or invoke BI and analytics on demand

- The same applies to a major growth in on-demand invocations of machine learning models (both predictive and prescriptive) which means that multiple instances of machine learning models will likely be needed along with load balancing of invocation requests for these services to handle concurrent usage

- In addition, the rapid growth in data science means that more predictive and prescriptive models will be deployed as services throughout the enterprise which in turn will fuel even more growth in concurrent usage

BI platforms need to be able to scale to handle more users and more data volumes

- The growth in data also means BI servers will have to scale to enable to load more data into memory for analysis in addition to scaling to handle more concurrent users. Execution of queries and predictive models also needs to scale to perform well on larger data volumes

BI platforms need to run both on-premises and on multiple clouds

- With analytical systems now running both on premises and in the cloud, BI and analytical servers need to run in both environments and be capable of porting artefacts between them. They also need to support connectivity to both cloud and on-premises data sources

Highavailability is needed

- The expectation is now that BI servers must support high availability both on-premises and in the cloud. This is especially the case when BI and analytics are integrated with customer facing operational transaction processing systems that could be 24 x 365 and when external users have access to BI on-demand from browsers and mobile devices

Internal and external user access

- The need to accommodate access to BI by internal and external users as well as cloud and on-premises applications means that different users

and applications will likely need to connect to BI and analytical servers via different identity providers

Privacy is becoming more important

- Internal and external user access also elevates the importance of data privacy once users have been authenticated. Protecting privacy means on BI servers only serve up what authorized users are allowed to see and that they mask or prevent access to what they are not allowed to see

Continuous integration / continuous delivery is needed to drive value

- A collaborative and continuous integration / continuous delivery (CI/CD) DataOps approach to developing queries, reports, dashboards and analytical applications is needed on BI platforms to significantly reduce time to value and continuously deliver new value

Leverage the cloud for consumption of BI and analytics

- Flexibility is needed to potentially develop on-premises and deploy on the cloud in as simple a manner as possible
- BI platform APIs are needed to easily integrate with multiple applications running in multiple environments either side of the firewall
- As reliance on BI and analytics grows and continues to be integrated into more applications, BI software platforms need to be easily upgradable and extensible without interrupting service availability

WHAT'S HAPPENING IN ANALYTICAL ENVIRONMENTS TO PROVIDE FLEXIBILITY, EXTENSIBILITY AND PRODUCTIVITY

BI software is changing to improve flexibility and productivity

Given these implications, a number of things are happening with respect to BI software that help improve flexibility, productivity and extensibility as well as shorten time to value. These include:

Easier to use natural language and voice interfaces

- The introduction of augmented intelligence (artificial intelligence built-in to BI tools) to opening up new natural language and voice interfaces so that business users can just ask questions and get answers. The result is even more people can start to make use of BI even if they are not sure about how to write SQL or how to navigate rich graphical user interfaces

- Collaboration support to improve productivity and speed up development

A more modular micro services architecture

- Support for a micro services architecture where the BI platform is made up of a number of discrete software components that interact to provide the functionality required. Each microservice is capable of performing a specific function in isolation and can interact with other microservices in the BI platform via internal APIs. The result is a platform that:

    - Is much more modular and easier to maintain

    - Can be quickly updated by replacing existing micro service with a new version to improve functionality without impacting on other services within the same platform

    - Can be easily extended by adding new microservices without impacting on the rest of the platform. This not only provides flexibility but also enables both vendors and certified partners to add services to extend the platform capability

This makes BI platforms easier to maintain, upgrade and extend

A modern microservices architecture also potentially opens up the platform to other developers to utilise specific analytical services e.g. to embed BI and analytical functionality when building applications.

Continuous integration / continuous delivery enables faster introduction of new functionality

- The adoption of a continuous integration / continuous delivery (CI / CD) approach to BI software development which enables vendors to release new functionality much more rapidly. Companies therefore get the chance to quickly take advantage of new capabilities, especially in private and public cloud computing environments, without having to go through major re-installation of new releases of software

DevOps technologies to run on-premises and in a multi-cloud environment

- The exploitation of DevOps technologies to seamlessly deploy BI / analytical software in a number of operating environments whether that be on-premises and on one or more clouds, automatically manage scalability, high availability, rolling upgrades and movement between development, test and production environments

New Analytical Operating Requirements Now Needed To Support Demand And Deliver Value

New requirements have emerged to enable business users to deliver vakue

With demand for both data, insights and analytical services rapidly increasing, there are a number of new operating requirements that need to be supported now that most companies are operating in a hybrid computing environment to enable business users to deliver value. These are as shown below.

It should be possible to:

Run on-premises and in multiple clouds

- Easily deploy, manage and operate data, BI and analytical platforms with minimal effort on-premises, in any cloud or both
- Declaratively configure and manage underlying infrastructure resources associated with a BI platform to
  - Dynamically scale the platform to handle more concurrent users and more data as demand grows and more data is accessed

Dynamic scalability

  - Scale specific BI platform microservices independently of others to cater for increases in demand for specific functionality e.g. security, data connectivity, analytical engines

Support continuous deployment using DevOps technologies

- Support continuous development / deployment of BI platform software using DevOps technologies (e.g. Docker) to place specific BI platform microservices in software containers (more later) for easy deployment of BI platform functionality on-premises or on the cloud. With this capability, new functionality can be provided in isolation at a much more granular level without the need to upgrade the entire platform

Leverage conrainer management software to automatically deploy and scale your BI platform on-premises or in any cloud

- Exploit container management software (e.g. Kubernetes) to automatically:
  - Manage the deployment of BI platform software on any operating environment on premises or in any cloud
  - Scale a BI platform to run on a cluster of servers on any operating environment on premises or in any cloud
  - Load balance requests across containers to avoid bottlenecks
  - Provide high availability to avoid outage of BI platform software services running in any environment
  - Manage rolling upgrades to BI platform software without disrupting usage
  - Manage the movement of analytical artifacts between development, test and production environments

Continuous development and deployment of reports, dashboards and models anywhere

- Support continuous development / deployment of data sets, queries, reports, dashboards, stories and machine learning models akin to a publish / subscribe style data and analytics hub and seamlessly port BI and analytical artefacts across cloud and on-premises environments

WHAT ARE CONTAINERS AND WHAT VALUE DO THEY BRING?

DevOps is key to meeting new operation requirements for BI software

These new operational requirements need the use of new technologies that underpin BI servers to support a multi-cloud hybrid computing environment. One of the key technologies required is containerisation technology.

Docker packages up software as a container image that can run in any environment

A good example of this is Docker. Docker can package up an application such as a BI tool or a set of microservices representing specific functionality within it, along with all its dependencies and create a Docker container image. This contains everything needed to run the application or specific application functionality on a particular operating system including code, runtime, system libraries and any settings.

A Docker container is a running instance of a container image

The Docker Container Engine enables containers to run regardles of underlying infrastructure and operating system

A container image requires a Container Engine to run the application or specific application functionality on a particular operating system. In the case of Docker, this is the Docker Engine. When an application runs it becomes a docker container which is a running instance of a docker container image. Each container runs as an isolated process (see Figure 2) on a Docker Engine which itself runs on a range of operating systems (e.g. Windows and Linux) both on premises and on multiple clouds. The advantage of this is that containerised software runs the same *regardless* of the underlying infrastructure or operating system. This makes applications portable across on-premises and multiple cloud environments. Containers therefore enable applications to run quickly and reliably in a number of computing environments. They are also lightweight and multiple containers can run on the same machine sharing the operating system with other containers. They therefore use less computing resources than virtual machines (see Figure 2) and so make better use of infrastructure.

Running multiple containers on hardware infrastructure uses a lot less computing resources than virtual machines

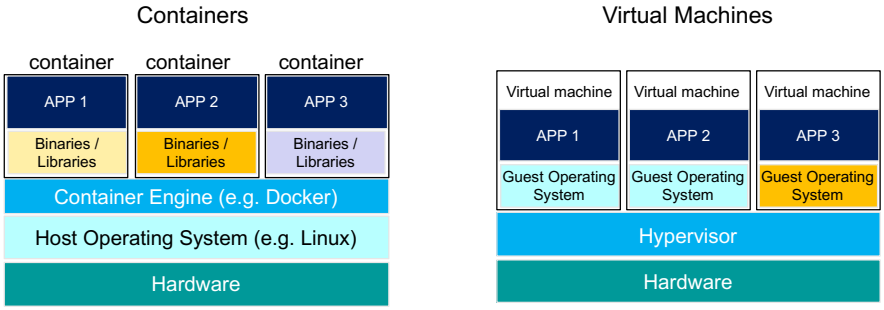

Figure 2

Running containerised software on the cloud reduces costs and makes better utilisation of resources

Different functional components of a microservices based BI platform can run in different containers

When deploying containerised software on the cloud, that translates to saving money. Also because containers run in isolation, failure of one is not going to impact on others running on the same machine. Another benefit is that different functionality in the form of microservices can be deployed in different containers allowing components of the same application or software product to be managed separately. In the case of a BI platform that means that new versions of functionality can be separately controlled and deployed rapidly and you can trust that a new build will run in production. Also, functionality can be substituted by other containerised services (e.g. different security identity managers) and it becomes easy to extend the platform without impacting on any other microservices within it.

WHAT IS KUBERNETES AND WHY DOES IT MATTER?

Kubernetes is open source container management software that can manage containers in a cluster of servers on premises or in any cloud

Another key DevOps technology is Kubernetes. This is open source container management software for running, managing and orchestrating containers in a cluster of servers.

Kubernetes can run on-premises, in the cloud or across multiple clouds. Many cloud providers offer Kubernetes as a service in their cloud environments e.g. Amazon Elastic Kubernetes Service, Microsoft Azure Kubernetes Service or IBM Cloud Kubernetes Service.

You can define how you want your Kubernetes cluster to be configured using a language call YAML

A number of Kubernetes objects can be used to configure a cluster

Kubernetes lets you define how you want a cluster to be configured. It does this by supporting a declarative language called YAML[1] which lets you define what you want in the cluster. Kubernetes then makes sure the system matches the configuration you have defined in your YAML file(s). Using YAML you can define a number of Kubernetes objects. These include

- Namespaces
- PODs
- Services
- Replication Controllers
- Deployments
- StatefulSets, Volumes and Claims

Namespaces are used to group together other objects so that they can be easily managed

A namespace is a logical grouping (a label) that all the other objects are associated with for a particular application. So for example, you could define a 'BIPlatform' namespace or a 'Customer_BI' namespace with all associated objects using that namespace.

One or more software containers runs in a Kubernetes POD

If a POD fails Kubernetes autimatically starts up a new one to replace it

PODs are the basic object in a Kubernetes cluster. They consist of 1 or more software containers. Normally only one container runs in a POD but it can run more. Each POD gets its own internal IP address and you can group PODs together using labels (i.e. all PODs with the same label). In the event of a POD failing, Kubernetes will automatically spin-up a new one to replace the faulty one. You define PODs in YAML and specify what containers you want to run in each specific POD. Kubernetes then does the rest.

Services are the interface to PODs and are used to load balance requests across multiple instances of a POD

Services are the interface to PODs and act as internal load balancers between POD instances. So instead of pointing to the POD directly, you issue requests to the service associated with a POD. The service knows which POD to send the request to (via the labels) and will keep listening for new PODs that use the labels you configured. You can also make POD endpoints available to the outside world by binding an *external* load balancer to your services. An example of why you would do this is to enable a web application (e.g. a BI tool) to be accessed by external users via a browser. You define services in YAML and associate them with a POD.

Kubernetes PODs can be replicated for high availability

Replication Controllers (RCs) run multiple instances of a POD on a cluster and so provide both scalability and high availability. An RC uses a POD template to

[1] Yet Another Mark-up Language

Each POD instance has a different IP address

understand which PODs to replicate. Each POD instance has a different IP address. You use YAML to define how many replicas you want and can change it via Kubectl commands.

Kubernetes Deployments enable automatic roll out of software updates without bringing down existing PODs

Deployments are another way of handling the scaling of PODs. This is a type of replication that provides automatic roll out of updates without bringing down existing PODs and the applications or application functionality that run in the containers associated with them. Every time the application code changes, a new version of your container will be built. Kubernetes will then start terminating PODs with the old version as it starts to spin up the new PODs with the updated container(s).

StatefulSets and Volumes allow you to persist and remember state even if a container crashes or is updates

Finally there are StatefulSets, Volumes, and Claims. While containers do not rely on state (which enables high end scalability), it is often the case that certain services require some kind of state to be persisted. For example a database or file may need to hold some data to remember state. In order to make this possible, you can create volumes and allocate them to a StatefulSet to maintain files in the volume even if a container crashes or if you update the software image in a container. Therefore a volume will remain in existence even if a container has to be restarted or updated inside a POD but when a POD is shut down the volume will cease to exist.

Kubernetes automatically creates a cluster of servers on any cloud to run containerised software

Kubernetes matters because of its ability to automatically:

- Create a cluster on premises or on any cloud to run and scale your containerized software on a cluster of servers
- Manage container execution on individual cluster nodes
- Scale containerised software by running containers in PODs, replicating PODs and load balancing requests across PODs using Kubernetes services to manage concurrent usage
- Recover from container failure by restarting them
- Manage rolling updates to software running in containers in any POD
- Maintain state if data is being stored in files or databases accessed by containerized software running in any PODs
- Create multiple cluster instances which are isolated from each other

It provides automatic elastic scalability, load balancing, high availability and roll out of updates with no downtime

Also multiple clusters with different configurations can be created running the same software

A Kubernetes cluster is made up of a master and nodes (see Figure 3).

A Kubernetes cluster has a master and multiple nodes

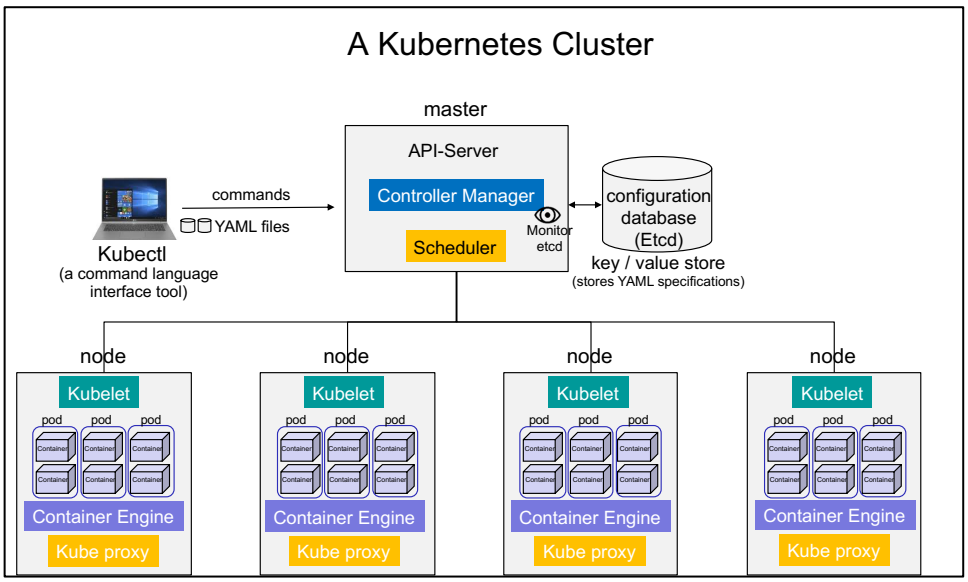

Figure 3

You can define how you want you cluster to be configured in any environment on-premises or in the cloud

The master decides which software containers run on which nodes

The master receives YAML files via REST calls to the API Server that come in via the Kubectl command line interface tool or from applications using the APIs. It then validates these files and stores the YAML specifications in its configuration data store. These files define the aforementioned Kubernetes artefacts (i.e. namespace, PODs, services, replication controllers, deployments, statefulsets and volumes) that together make up how the cluster should be configured to run the containerised software. The master then decides which PODs run on which nodes of the cluster and assigns them to nodes using the scheduler. The controller manager looks out for changes to the cluster that come it via Kubectl commands and new versions of YAML files. When this happens, it then manages automatic updates.

New containers, container instances and new nodes can all be added dynamically on the fly

Each node runs a piece of software called a Kubelet that registers the node in the cluster and looks out for commands from the master scheduler to instantiate PODs. Kubelets communicate with the container engine runtime (e.g. Docker) on a node in order to run the software containers within the PODs.

By using continuous integration / continuous delivery it becomes possible to automatically detect and roll out software upgrades across a Kubernetes cluster

Finally it is possible to automatically operate and manage changes to a Kubernetes cluster using continuous integration / continuous delivery software (e.g. Jenkins), that watches for newly committed changes to an application code repository, create new docker containers for your application or software product, updates the YAML infrastructure files with the new container versions and then triggers automatic rollout of upgrades across your Kubernetes cluster. It's a powerful story and is an increasingly well-trodden path in software development and operations. So much so, that it is becoming a de-facto standard method of managing and operating in a modern 'always on' hybrid computing environment.

WHAT BUSINESS VALUE DOES KUBERNETES BRING TO AN ANALYTICAL ENVIRONMENT?

Having identified the above requirements and obtained a deeper understanding of DevOps technologies, the question is how does this apply to an analytical environment?

Kubernetes enables your BI platform to be deployed on-premises or in any cloud

Kubernetes provides a number of benefits. For example, you can manage and automate the deployment of software (e.g. your BI platform or a containerised microservices based subset of its functionality) on any operating environment on premises or in any cloud. You can also dynamically configure the scalability of a BI platform and/or any individual containerised microservices within it to run on a cluster of servers on any operating environment on premises or in any cloud. This is important for scaling BI platforms to handle more concurrent users and analyse more data as described earlier. The same is true for any services developed on a BI Platform. For example if you create new queries, reports, dashboards, predictive models etc., and publish them as analytical services on your BI platform, then Kubernetes, by underpinning the running of your BI Platform, can effectively scale these services to support a growing number of user and application requests to consume the insights they produce.

BI platform and / or any of its individual containerised microservices can be dynamically scaled in any environment

Configure different instances of a BI Platform running in a Kubernetes cluster for development, test and production

In addition you could create multiple instances of a Kubernetes cluster (with different configurations) for development, test and production. You could also do this to isolate different types or analytical workloads or different types of users e.g. internal users and external users.

Automatically scale your BI platform using Kubernetes to handle more users, more data and to enable high availability for use by applications

Kubernetes can also replicate containers in a cluster and load balance requests across them (using PODs, services and RCs) to scale specific BI platform microservices functionality, avoid bottlenecks and provide high availability. Therefore individual containers can be scaled independently of others which may be needed as analytical workloads vary or as concurrent user numbers grow. User-defined configuration of replica PODs running containers allows you to enable and automatically manage high availability of your BI platform to avoid service outage in any environment. This is particularly important as BI and analytical services are made available to global and external users and integrated into 'always on' externally facing web applications.

Use Kubernetes to automatically roll out updates to your BI platform which keeping it available

Kubernetes deployments enable rolling upgrades to BI platform software to also be automated without disrupting usage. You can also automate the management and movement of queries, reports, dashboards and models between development, test and production environments.

Manage the movement of queries, reports, and dashboards between development, test and production

Kubernetes therefore offers many business benefits including lower operational cost (via dynamic adjustment of infrastructure), on-premises and multi-cloud portability, automated elastic scalability, high availability, workload management, and continuous integration / continuous delivery (CI/CD) of both BI software and analytical services developed on that software. It also supports automated rollout of software updates without disruption.

Separately configure, manage and scale different aspects of BI platform functionality e.g. to handle more users

Add additional services to the cluster that integrate with BI

Run different production cluster for different analytical workloads

Develop queries, reports, dashboards and models on-premises and publish to the cloud for consumption

Centrally administer users and security entitlements and push to the cloud

This plays well in the world of modern microservices based BI/ Analytics software as it allows specific BI/ analytical platform services to be isolated and separately managed in multiple different containers and container groups. For example for security or data connectivity for different groups of concurrent users can be handled separately from a BI platform analytical engine. Also multiple instances of an analytical engine could be created to scale the BI platform. This means you can separately configure resources by service type, replace service types with others if needs be (e.g. different identity providers for internal and external users) and even add new analytical services provided by other technologies e.g. to integrate BI with machine learning services for example.

Also analytical artefacts such as machine learning models can be developed and deployed as services in containers. These could then be managed, scaled separately based on concurrent usage and availability requirements and all upgraded automatically using Kubernetes. The same could be said for other types of analytical artefacts (reports, dashboards…) developed on that platform.

Kubernetes also enables a develop on-premises, deploy on the cloud approach to flourish in an analytical environment. You could set up development, test and production clusters either side of the firewall and manage resources to them all separately. Similarly, security entitlements could be defined on premises and pushed to the cloud (centralised authorisation). And as more concurrent users emerge using simpler natural language user interfaces, mobile apps, and accessing BI indirectly via operational applications, you can manage and scale it all using Kubernetes. It makes the BI platform a highly available and scalable hub at the centre of the enterprise and encourages mass consumption of insights to enable more employees to contribute to delivering business value as depicted in Figure 1.

How Is Qlik Exploiting The Power of Containers and Kubernetes?

Qlik Sense Enterprise has a microservices based architecture

Given the increasing demand for data, BI and analytics, the impact this is having on BI/analytical platforms and the new analytical operating requirements it brings, we have seen the emergence of DevOps technologies like Docker containers and Kubernetes to help provide flexibility, scalability and high availability now needed. In this section of the paper we look at how one vendor, Qlik is taking advantage of this within its analytical platform.

Different aspects of Qlick Sense Enterprise functionality can be created as different Docker containers and scaled independently

Qlik Sense Enterprise is Qlik's flagship cloud-native BI platform. It has a modern elastic microservices based architecture that allows it to group specific functionality in to separate docker containers and run them on-premises or on any cloud using Kubernetes. The architecture for this is shown in Figure 4.

Qlick Sense Enterprise can run in a multi-cloud environment

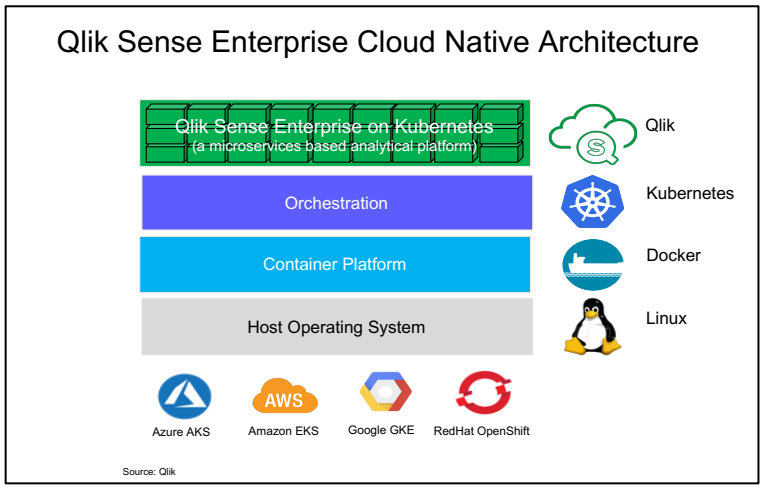

Figure 4

Qlick Sense Enterprise has different Docker containers for user and administative clients, the analytical engine, data connectivity, security and infrastructure

Within Qlik Sense Enterprise there are a number of components including:

- A browser based portal and client for users called Hub Client
- A browser based administrative management console client
- Client containers
- Engine containers including the associative data indexing engine, in-memory and load balancing
- Data containers to manage connectivity to data sources, reloading and scheduling
- Security containers to configure and manage authentication and entitlements
- Infrastructure containers for network external load balancing, datastore (MongoDB) and in-memory caching (Redis)
- Storage of Qlik Sense applications and content in MongoDB and Kubernetes Data Volumes

You can see right away that it is not one monolithic software application. Far from it. Multiple containers enable different elements of Qlik Sense Enterprise functionality such as security, the analytical engine, data connectivity, and Qlik Sense infrastructure to be managed in separate PODs and configured and scaled accordingly via Kubernetes to meet increasing demand of more users.

BUSINESS BENEFITS OF RUNNING QLIK IN A KUBERNETES ENVIRONMENT

Given this architecture there are several business benefits from running Qlik Sense Enterprise in a Kubernetes environment.

Kubernetes makes it possible to scale Qlick Sense Enterprise to handle more users and more applications requesting insights on-demand

Elastic, Modular Scalability To Enable Increased Usage

The first is that Kubernetes really enables businesses to put data, BI and analytics at the centre of their company and open Qlik Sense Enterprise up to many more internal and external users taking advantage of simpler user interfaces. In addition it also allows a separate high availability Kubernetes cluster to be created to make use of BI and machine learning services (queries, reports, dashboards, etc.) by integrating Qlik Sense Enterprise with 'always on' transaction processing and mobile applications. This enables operational BI and machine learning models to be embedded right into core business processes. It also enables scalability to handle more data via Kubernetes automated scale up and scale out.

Each Qlick Sense Enterprise containers can be separately scaled on a Kubernetes cluster on-premises or on any cloud

Furthermore, Qlik Sense Enterprise's microservices architecture means that different types of microservices run in different containers in different Kubernetes PODs and can be separately scaled via Kubernetes POD replicas and services that load balance requests across those containers. It can also distribute these across more nodes in a Kubernetes cluster as and when they needed. This enables specific BI platform functionality to scale separately as well as the platform as a whole to scale.

Kubernetes allows different Qlick Sense Enterprise clusters to be configured for development, tes and production all with support for automatic roll out of upgrades

In addition these Qlik Sense Enterprise containers can be brought together in separate Kubernetes deployments that are configures for development, test and one or more Qlik Sense Enterprise production clusters all with automated continuous rollout of software upgrades without disruption. All of this can be managed and adjusted declaratively via YAML. In addition, Kubernetes POD replication makes Qlik Sense Enterprise highly available (e.g. to serve up operational BI via APIs into customer facing applications on the web) and different production Kubernetes clusters could exist for different groups of users, different analytical workloads and/or different Qlik Sense applications.

Multi-Cloud Enabled For Cost And Consumption Optimisation

Develop on Qlick Sense Enterprise on premises and publish to the cloud

Qlik Sense Enterprise is platform agnostic in that it can run on Windows on-premises in the data centre, with a cloud native microservices architecture utilising Docker and Kubernetes so that you can deploy it on-premises / private cloud, in a single cloud like AWS, Azure or Google Cloud or across multiple clouds. It can also be run on Qlik Cloud Services fully managed as a service by Qlik. This flexibility allows you to create different Qlik Sense Enterprise deployments on premises or on different sizes of Kubernetes cluster for consumption in different environments. This helps optimise the cost of consumption and allows BI content authoring, security administration and data sources to be set up once on-premises and pushed to the cloud.

Automatically scale Qlick Sense Enterprise to handle more data sources both on-premises and on multiple clouds

Hybrid Data Connectivity To Meet The Demand For More Data

Also, Qlik Sense Enterprise can not only run on-premises (Qlik Sense Enterprise on Windows) and in a multi-cloud environment (Qlik Sense Enterprise on Kubernetes) but it can bridge the corporate firewall and connect to both cloud-based (e.g. Amazon Redshift, Azure SQL Data Warehouse Google Big Query, Snowflake) and on-premises data sources.

Separate But Integrated On-Premises and Cloud Environments

Easily move your Qlik Sense Enterprise environment to a different cloud with no rework

Running Qlik Sense Enterprise on Kubernetes in the cloud and on Windows on-premises inside the firewall means you can connect to on-premises data sources and push Qlik Sense applications to the cloud in as many engines and clusters as you like. This enables a develop on premises and deploy in the cloud approach. It also makes migration to the cloud easier and allows one or more analytical engines and analytical clusters in cloud production environments. You can also support multiple clouds and move off one cloud onto another with no rework.

Individual And Group Level Security

By separating Qlick Sense Enterprise security microservices into its own container, Kubernetes allows you to support different identity providers for internal and exteral users

When running in a multi-cloud environment we know already that providing access to BI by internal and external users on the cloud or on mobile apps elevates the importance of data privacy and access security. Qlik Sense Enterprise accommodates this by supporting identity provider and edge authentication containers to manage different identity providers authenticating via the OpenID standard. Identity providers allow for a common identity to be used in all Qlik Sense Enterprise on-premises and cloud deployments so that common security rules and policies can be defined once and used everywhere. It also supports user and user groups with personal and shared collaborative areas called *spaces* that allow for single user and collaborative development on Qlik Sense Enterprise. In this case the user and user group attributes are provided by the Identity Provider. Once Qlik Sense applications have been developed they can publish into managed production spaces (e.g. in the cloud) for consumption by authorised users either within or outside the enterprise.

Centralised Administration

Qlick Sense Enterprise allows you to keep single user and user group developed content separate using spaces to maintain privacy before punlishing to production spaces for easy consumption

Administration can be centrally managed to ensure consistency across all Qlik Sense Enterprise deployments on-premises and in multiple clouds

Administration can be managed centrally with the Qlik Sense Management Console. This is a multi-cloud setup console that allows you to manage multiple Qlik Sense Enterprise environments all from the same console. For example you can manage Qlik spaces, users, identity providers and web integrations as well as distribute Qlik Sense applications, metadata and connectivity credentials across a multi-cloud target Qlik Sense Enterprise environment. You can also implement common licensing across multiple Qlik Sense Enterprise deployments.

CONCLUSIONS

Qlik Sense Enterprise is built to enable specific aspects of its functionality to be separately scaled to meet different needs such as more users and more data

Kubernetes allows multiple Qlick Sense Enterprise clusters to be deployed for different user communites and different workloads

All of it can be configured, managed and continuously updated in a multi-cloud environment

The Qlik Sense Enterprise modern microservices architecture enables key components of platform functionality to be isolated, packaged up into containers and run on Kubernetes. This purpose of this is to enable different aspects of the platform functionality to scale to handle more users, more data, more data sources, and integration of BI with more on-premises or cloud applications in a modern data driven enterprise. Kubernetes provides flexible deployment of Qlik Sense Enterprise on premises and on one or more clouds.

It also makes it possible to define and configure different independent clusters for development, test and different production workloads so that BI content, policies and rules can be centrally developed in a low-cost environment and pushed to workload and cost optimised production Kubernetes clusters running in a multi-cloud hybrid computing environment. All of this can be centrally managed and administered with a common approach to access security, user management, licensing and more. Furthermore, it can all be done with high availability where needed, automated rolling software upgrades with no disruption to service with continuous integration and continuous delivery (CI/CD).

Given the pressure to lower total cost of ownership and move data, BI and analytics to the centre of the enterprise to drive value right across the business and beyond, Qlik Sense Enterprise is well placed to fulfil this role.

About Intelligent Business Strategies

Intelligent Business Strategies is an independent research, education and consulting company whose goal is to help companies understand and exploit new developments in business intelligence, machine learning and advanced analytics, data management, big data and enterprise business integration. Together, these technologies help an organisation become an *intelligent business*.

Author

Mike Ferguson is Managing Director of Intelligent Business Strategies Limited. As an independent IT analyst and consultant he specialises in data management and analytics. With over 38 years of IT experience, Mike has consulted for dozens of companies on data strategy, big data, machine learning and advanced analytics, data governance, master data management, data warehousing and enterprise architecture. He has spoken at events all over the world, written numerous articles and is Chairman of Big Data LDN, one of the fastest growing and largest data and analytics conferences in Europe. Formerly he was a principal and co-founder of Codd and Date Europe Limited – the inventors of the Relational Model, a Chief Architect at Teradata on the Teradata DBMS and European Managing Director of Database Associates, an independent analyst organisation. He teaches popular master classes in Data Warehouse Modernisation, Designing, Managing and Operating a Multi-Purpose Data Lake, Machine Learning and Advanced Analytics, Big Data and Analytics Fundamentals, Enterprise Data Governance and Master Data Management, Modern Data Architecture, Real-Time Analytics and Operational BI.

8 Paddock Chase, Poynton
Cheshire, SK12 1XR
England
Telephone: (+44)1625 520700
Internet URL: www.intelligentbusiness.biz
E-mail: info@intelligentbusiness.biz