



CLOUD MIGRATION AND INTER - CLOUD MOBILITY ASPECTS

Mikhail Gloukhovtsev

Sr. Cloud Solutions Architect

Orange Business Services

Mikhail.Gloukhovtsev@orange.com

Table of Contents

1. Introduction	4
2. Initial Application Migration to Cloud	5
2.1 A Variety of Factors Determining Cloud Migration Strategy	5
2.1.1 • Cloud Service Type – SaaS, PaaS, or IaaS and the Cloud Migration Strategy	7
2.1.2 • Cloud Service Deployment Model – Public, Private, Hybrid Clouds and the Cloud Migration Strategy	8
2.1.3 • Cloud Migration Strategy: Migration Types	9
2.1.4 • Network Solutions for Cloud Migration Process: AWS Direct Connect, Microsoft ExpressRoute, Google Carrier Interconnect, and Orange Business VPN Galerie	10
2.2 Migrating Unstructured Data to the Cloud	15
2.2.1 • V2V Cloud Migration	18
2.2.1.1 • Offline Migration Using vSphere Replication	18
2.2.1.2 • Migration to VMware vCloud Air Using vCloud Connector	20
2.2.1.3 • Cloud Migrations Using Long-Distance vMotion	21
2.2.2 • Data Migration to Cloud: VPLEX	21
2.2.3 • Migrating Unstructured Data to the Cloud Using Provider's Export/Import Services	22
2.3 Migrating Structured Data to the Cloud	23
2.3.1 General Considerations for Migrating DBs to the Cloud	23
2.3.2 Migrate to Cloud-Based DB Service or to Run DBs on Cloud VMs?	24
2.3.3 Migrating DBs by Restoration from Backup and Synchronization	25
2.3.4 Migrating DBs from MS SQL Server to Azure SQL Databases	26
2.3.4.1 Migrating DBs to Azure SQL Databases by Using SQL Server Management Studio	28
2.3.4.2 Migrating DBs to SQL Database by SQL Database Migration Wizard	29
2.3.4.3 Migrating DBs by Updating the Source DB and Then Deploying It to Azure SQL DB	30
2.3.4.4 Migrating DBs by Using Transactional Replication	31
2.3.5 Migrating SQL DBs to Azure VMs	33

2.3.6 Cloud Provider Services for Migrating DBs	35
2.3.7 Migrating Oracle DBs to Cloud	36
3. Inter-Cloud Migration and Workload Portability	36
3.1 Cloud Interconnectability	37
3.2 Data and Workload Portability in Cloud	38
4. Conclusion	40
5. References	41

Disclaimer: The views, processes or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes or methodologies.

1. Introduction

When I work with my customers to help them define their cloud computing strategy, one of the main questions they raise is about the best ways to move the application workloads to the cloud. The customers understand that the migration to the cloud-based IT services involves various transformations – business, organizational, and architectural changes. They want to know how much effort is required, how to minimize the transformation impact on the IT services provided to the business, and what changes in the application architecture are needed.

A lot of publications describe how to make “Journey to the Cloud” successful.¹⁻⁵ They review various aspects including transformation of business processes along with changes in the IT service models and organizational structures. For example, the decision on the right sourcing for application workloads (private, hybrid, or public cloud) is based on financial analysis (OPEX vs. CAPEX, licensing and usage models, time-to-market requirements), functional and service level requirements, and regulation (data security) compliance. Applications are profiled to create a tiered application lexicon for right-sourcing and migrating to the cloud. Various levels of the application data flow redesign may be required.

The implementation of cloud computing is multifaceted. The goal of this article is to consider challenges and solutions associated with cloud migrations. Specifically, we will focus on the technology and methodology of the application and related data migration solutions from the practitioner’s point of view. For reviewing other aspects of the cloud migration (cloud migration frameworks (e.g. CMOTION and CloudMIG),¹ general holistic methodologies such as model driven architecture (MDA) transformation, business case development, migration planning, application assessment, proof of concept, etc.), readers are referred to the numerous publications.¹⁻⁵

We will consider:

- *Moving data to/from the cloud* (the initial migration to the cloud and data recovery in case of service termination)
- *Migrating data between different cloud environments*

In relation to data migrations, I will also briefly discuss the application data portability enabled by using various technologies such as containers and cloud storage gateways.

Challenges of cloud migration, its complexity, and risks may discourage some companies from undertaking their journey to cloud. As a result, the migration project may be postponed despite understanding that delays in getting benefits of cloud-based IT services impact company competitiveness. The author hopes that this article addressing the practical aspects of cloud migrations will help improve the effectiveness of migrating to cloud.

2. Initial Application Migration to Cloud

2.1 A Variety of Factors Determining Cloud Migration Strategy

A strategy for migrating applications to cloud depends on many various factors, among them:

1. Cloud service type – SaaS, PaaS, or IaaS;
2. Cloud deployment type – private, public, or hybrid;
3. Application architecture – single- or multi-tier, usage pattern, application interdependency;
4. Application criticality – acceptable downtime window;
5. Amount of the data to be migrated;
6. Data security requirements;
7. Whether physical-to-virtual (P2V) server conversion is part of the migration to cloud.

In some cases, a complete application redesign may be required to move the application to the cloud. Alternatively, an incremental hybrid cloud-based approach can be implemented by moving cloud-ready application components to cloud while keeping the legacy part in the data center with the goal of transforming the cloud-unfriendly components to make it possible to move them into the private cloud.

What makes a migration to the cloud a success? First of all, it is about meeting objectives of IT service delivery to the business units:

- Low cost for the migration
- Online migration or a short downtime window
- Minimal impact on existing IT services
- Opportunity for P2V conversion
- Data security during the migration process

Here are just a few of the many challenges for cloud migration:

- Significant downtime can be required to migrate critical applications
- Procurement of migration software tools and hardware may be required
- Purchasing (short term leasing) additional network bandwidth may be needed for the migration
- Configuration changes to meet cloud IT standards
- Cutover planning; while the goal is to minimize application downtime, can you afford to keep your existing environment up and synchronize the data with the new cloud-based environment until the new environment is tested and validated so that you can redirect all the traffic to it?
- Is a staging environment to reconfigure and validate the migrated applications and/or virtual machines before moving them into the production cloud environment needed?

Based on the application taxonomy, application move groups should be identified by the application teams to allow for optimal sequencing during migration. These move groups bundle applications and their associated server and storage infrastructure that need to move together during the migration.

For each application move group, the following will be defined:

- Application migration strategy tier;
- Data in-flight encryption requirements;
- Migration downtime;
- Downtime risk range;
- Critical dependencies;
- Server and data migration schedules;
- Data synchronization solutions.

To achieve the goals of the cloud migration, best practice is to set up what can be called a Migration Factory including development of application and/or environment migration solutions and related procedures, creating or selecting tools, and establishing the vision for the end-state environment based on using the cloud services. It is important to get agreement from all stakeholders about the criteria of migration success and establish governance and processes for moving from the legacy infrastructure to the cloud.

As all these factors mentioned above are important for successful migration to cloud, we will keep them in mind while we discuss the main topic of this article - the technology aspects of the cloud migration solutions.

2.1.1 • Cloud Service Type – SaaS, PaaS, or IaaS and the Cloud Migration Strategy

Gartner⁶ suggests five general methodologies for migrating legacy systems to the cloud:

1. Rehost on infrastructure as a service (IaaS)
2. Refactor for platform as a service
3. Revise for IaaS or PaaS
4. Rebuild on PaaS
5. Replace with software as a service

Other reviews of cloud migration strategies categorize the migration based on the target cloud - migration to IaaS, migration to PaaS, and migration to SaaS.⁷⁻¹⁰

Software as a Service. In the case of moving to SaaS, application migration per se is not required as the existing application is replaced with a SaaS delivery option. While this migration strategy through replacing legacy system by SaaS does not involve reengineering of the legacy system, the existing application data may need to be imported into the new platform. This may result in a need for data conversion to a new data format. If only some application functionality is outsourced to the cloud, the business process should be revised to integrate the cloud and traditional services.

Platform as a Service. Migration to PaaS is usually chosen for migrating business applications that are based on standard application platforms such as JavaEE or .NET platforms. In the PaaS model, the application owner keeps control over the applications and the service provider offers an environment for application development including the full lifecycle up to the application deployment (for example, Cloud Foundry, OpenShift, BlueMix, etc.). If the selected PaaS does not support some application features, the application migration to the PaaS requires some level of the application redesign. The legacy system will be migrated to the PaaS by system refactoring or rebuilding on PaaS, according to Gartner.⁶

At the PaaS level, storage unit for application data is a database (relational or non-relational (aka NoSQL) database) with the related storage platform. Therefore, the application data migration involves data export and import using a format required by the PaaS DB.

Infrastructure as a Service. In contrast to SaaS and PaaS, in the case of IaaS the migration scope is broad with many various options. Migration to IaaS results in deploying the application on the cloud service provider's servers. While IaaS is the best choice for moving applications to the cloud when there is no time to reengineer the applications for a cloud, compatibility of the server OS and hardware features used by the existing applications (for example, some CPU features may be important for a given application) should be reviewed and application porting may be a necessary step in the migration to IaaS.

Data migration to IaaS may involve data format conversion between different storage platforms, for example, moving block-data to object-based storage. Depending on the application criticality and data amount, the data migration can be done keeping the application online, or using offline data transfer with initial data copy (or a so-called pre-copy process) with following delta synchronization later. If the data volume is large, a "swing storage system" can be used to copy data at the current data center, ship the storage system to the cloud provider data center, and synchronize the delta.

2.1.2 • Cloud Service Deployment Model – Public, Private, Hybrid Clouds and the Cloud Migration Strategy

There are a few key considerations that determine your decision on which cloud deployment model – Public, Private, or Hybrid Cloud – to use:

- *Security and regulatory compliance requirements* make private cloud the preferred model as it provides better security controls, customized security policies, and operational procedures. Depending on the application architecture, you may consider using private cloud storage in hybrid cloud deployment, for example, private cloud storage solution offered by NetApp.¹¹ You can choose using a public or hybrid cloud if the provider's multitenancy security policies meet your requirements.
- *WAN traffic:* If an application generates intensive data traffic that in some cases needs to be encrypted, WAN bandwidth costs and potential performance risks may exclude a public cloud from your option list. You may also consider low-network latency solutions offered by some public cloud providers if the cost of such solutions is outweighed by the benefits of the provided public cloud services; see Section 2.1.4 for more details.

- *Migration to the cloud:* Private cloud solution, particularly if it is developed and managed by your IT team, allows you to minimize changes in your IT standards while migrating to your private cloud. If your private cloud architecture can effectively deliver cloud services while using similar server, storage, and/or data protection platforms in the existing and end-state environments, it makes the migration to the cloud much easier. Less changes in the application architecture and minimal configuration changes mean less risks of application migration to the cloud. On the contrary, when you migrate to a public cloud, the use of your own image may be not supported or you have to adhere to the image compliance policy of the cloud provider.
- *Legacy application integration:* Applications dependent on the legacy applications running on platforms such as a mainframe and AS400 are better candidates for migration to the private cloud.

2.1.3 • Cloud Migration Strategy: Migration Types

Based on the five general cloud migration methodologies mentioned in Section 2.1.1, the following migration types can be identified:^{10,12,13}

Type I: Replacing some of application architectural components with cloud offerings. This is the least invasive type of migration exemplified by migrations to SaaS and PaaS (Section 2.1.1). While the application data and/or business logic migrations are still required, the migration is less complex compared with other migration types. Using the Azure SQL Database service in place of a local SQL server database is an example discussed in Section 2.3.4. It is important that the application vendor validates and supports the SaaS cloud platform you move your application to.

Type II: Migration of some of the application functionality to the cloud. In this migration type, one or more application layers, or a subset of architectural components from one or more layers implementing a particular functionality are moved to the PaaS cloud. Using a combination of Amazon SimpleDB and EC2 instances to host the data and business logic is an example of such migration.

Type III: Migration of the whole application stack of the application to the cloud. This migration type is used to move to an IaaS cloud. For example, VM clones containing encapsulated

application can be imported to the cloud-based environment (see Section 2.2). Then changes of the application configuration to make it work in the new cloud environment, e.g. changing the storage access, IP addresses, DNS, backup settings, etc., should be done. The main challenge is that the reconfiguration is application-specific and usually involves some parameters that should be determined a priori by testing.

Type IIIa: A green-field deployment with the application data import into the cloud. As some cloud providers may not allow the use of the customer VM images, an alternate way is to create an AWS or Azure VM and install the application on it (see Section 2.3.2). In this case a new VM is provisioned using the cloud's existing image catalog and then the application is installed starting from scratch. The existing data should be secure copied to the cloud-based application. While this may be seen as the most straightforward solution, creating the application environment that requires the application data migration can be a time consuming approach.

Type IV: "Cloudifying" the application by completely replacing the entire application functionality with a composition of services running in the cloud (SaaS).

2.1.4 • Network Solutions for Cloud Migration Process: AWS Direct Connect, Microsoft ExpressRoute, Google Carrier Interconnect, and Orange Business VPN Galerie

What network connectivity do you need to move your applications and application data to a cloud? When you consider the cloud-readiness of your applications, such aspects as their network bandwidth, application tolerance to network latency, and security requirements for the traffic to/from the cloud are key factors in decision making. This article will focus on the network requirements specifically for the cloud migration process.

While the Internet is a simple, flexible, and cost effective way to migrate to cloud services and use them after migration is completed, several questions should be answered:

- How long will it take to migrate data with the available bandwidth?
- Should additional bandwidth be ordered only for the migration to meet the application downtime windows and migration schedule? Can it be purchased for a short period of time – just for the migration?
- How secure is the data traffic? Are the security requirements for data in-transit met?
- Is redundant connectivity needed to mitigate network reliability issues and complete the migration on time?

A secure virtual private network (VPN) is one way to manage the security of data during its migration to a cloud environment. A VPN can operate securely over the Internet and still provide high levels of security through encryption.

Do you need to purchase additional bandwidth? To address the potential network bandwidth issues during the cloud migration, use of WAN optimization appliances is recommended. Potential performance gains from using WAN accelerators can be assessed by using WAN-emulation appliances such as HPE/Shunra Virtual Enterprise appliance combining appliance-based WAN emulation and application performance analysis.¹⁴

For facilitating data migration and in case of bandwidth-heavy applications which are sensitive to network latency, cloud providers offer dedicated network connections. For example, AWS Direct Connect lets customers establish a dedicated network connection between their network and one of the AWS Direct Connect locations. Each AWS Direct Connect connection can be configured with one or more virtual interfaces. Virtual interfaces may be configured to access AWS services such as Amazon EC2 and Amazon S3 using public IP space, or resources in an Amazon Virtual Private Cloud (VPC) using private IP space. Partitioning this dedicated connection into multiple virtual interfaces enables maintaining network separation between the public and private environments. Bandwidth of 50 Mbps, 100 Mbps, 200 Mbps, 300 Mbps, 400 Mbps, and 500 Mbps can be ordered from any AWS partners supporting the AWS Direct Connect.^{15,16}

Similarly Microsoft offers the ExpressRoute to access the Microsoft cloud over a dedicated private connection facilitated by a connectivity provider. Using the ExpressRoute, connections can be established to Microsoft cloud services, such as Microsoft Azure, Office 365, and CRM Online. Connectivity can be from any-to-any (IP VPN) network, a point-to-point Ethernet network, or a virtual cross-connection through a connectivity provider at a co-location facility. As ExpressRoute connections do not go over the public Internet, they are more reliable, with higher speeds, lower latency, and better security than typical connections over the Internet.

Google Cloud Interconnect presents Cloud Platform customers with several connectivity options to connect their network to Google. One of them is the Carrier Interconnect, an option enabling customers to connect their network to a service provider with a direct connection to Google. This connection helps to provide higher availability, higher security, and lower latency since the

traffic travels from the customer's data center to one of the Google's partners – Carrier Interconnect service providers – which then passes it directly to Google via network peering.

VMware vCloud Air offers the Direct Connect that is a private, high-throughput, dedicated connectivity option for connecting on-premises or co-located environment to vCloud Air. The standard vCloud Air connection is routed over the public Internet and provides up to 1 Gbps or up to 300 Mbps bandwidth for the Dedicated Cloud or Virtual Private Cloud, respectively. Using the Direct Connect, VMware provides port connectivity and connection from vCloud Air to meet-me room: up to 10 Gbps for Dedicated Cloud and up to 1 Gbps for Virtual Private Cloud / Disaster Recovery Cloud. Private line is provided by a partner and is a separate obligation and billing relationship from vCloud Air subscription.¹⁷

Business VPN service providers offer integration services allowing customers to extend their on-premises environments into a cloud or to a few different cloud providers. For example, Orange Business Services (OBS, EMC partner as Cloud Service Provider (Gold)) provides Business VPN Galerie that brings cloud services inside enterprises' VPNs.¹⁸ Business VPN Galerie was launched in 2011 in anticipation of the growing demand from multinational corporations in their digital transformation for high performance, secure access to Internet-based cloud and web services. It extends the corporate VPN to the cloud via fully-secured gateways with end-to-end high performance and reliability. As shown in Figure 1, cloud providers such as AWS, Azure, and Google Cloud Platform connect to the Orange Business Services network via Business VPN Galerie to make their services seamlessly available to any Business VPN customer. Business VPN customers wishing to access cloud providers' services simply subscribe to a Business VPN Galerie option in addition to Business VPN. The cloud providers' services are then seen as belonging to customer's private network.^{18,19}

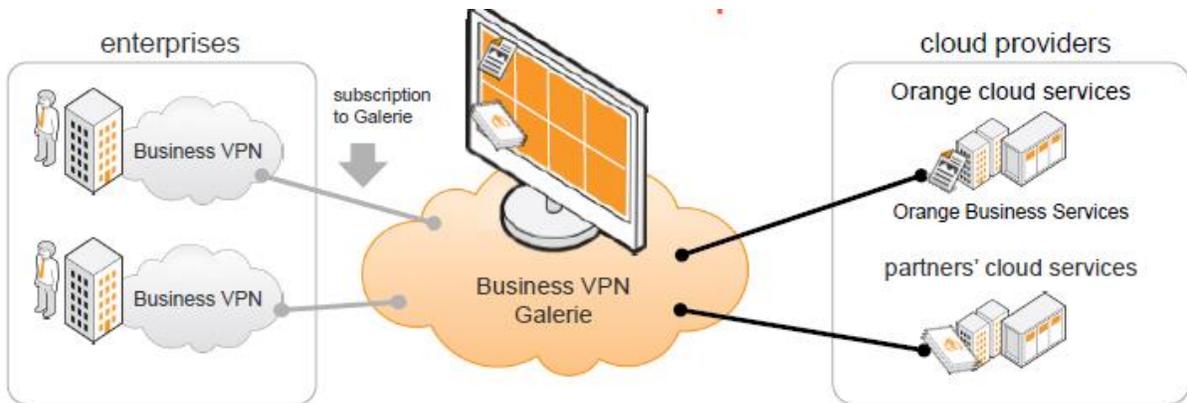


Figure 1: Orange Business VPN Galerie (Ref.18).

The OBS Business VPN Galerie is designed for optimal cloud services delivery (Fig. 2).

- simplicity** ▶ Orange manages connection to cloud providers making them virtually part of the WAN
- reliability** ▶ 100% availability SLA (connection to provider)
▶ end-to-end connection management and support is handled by Orange
- performance** ▶ cloud traffic goes on-net
 - Business VPN performance SLAs apply
 - performance is predictable
- security** ▶ cloud traffic goes on-net
 - Business VPN traffic isolation applies
 - traffic is not exposed to the Internet

Figure 2: Benefits of the OBS Business VPN Galerie

The Business VPN Galerie provides direct access via a secure, high-performance and managed network with a single and trusted partner to Google Cloud Platform datacenters, to AWS via the Direct Connect (Fig. 3), to the Orange Flexible Cloud, and to the all cloud-based services accessible via Microsoft ExpressRoute: Azure – the Microsoft IaaS/PaaS solution and Microsoft Services – SaaS solutions from Microsoft (Fig. 4).^{18,19}

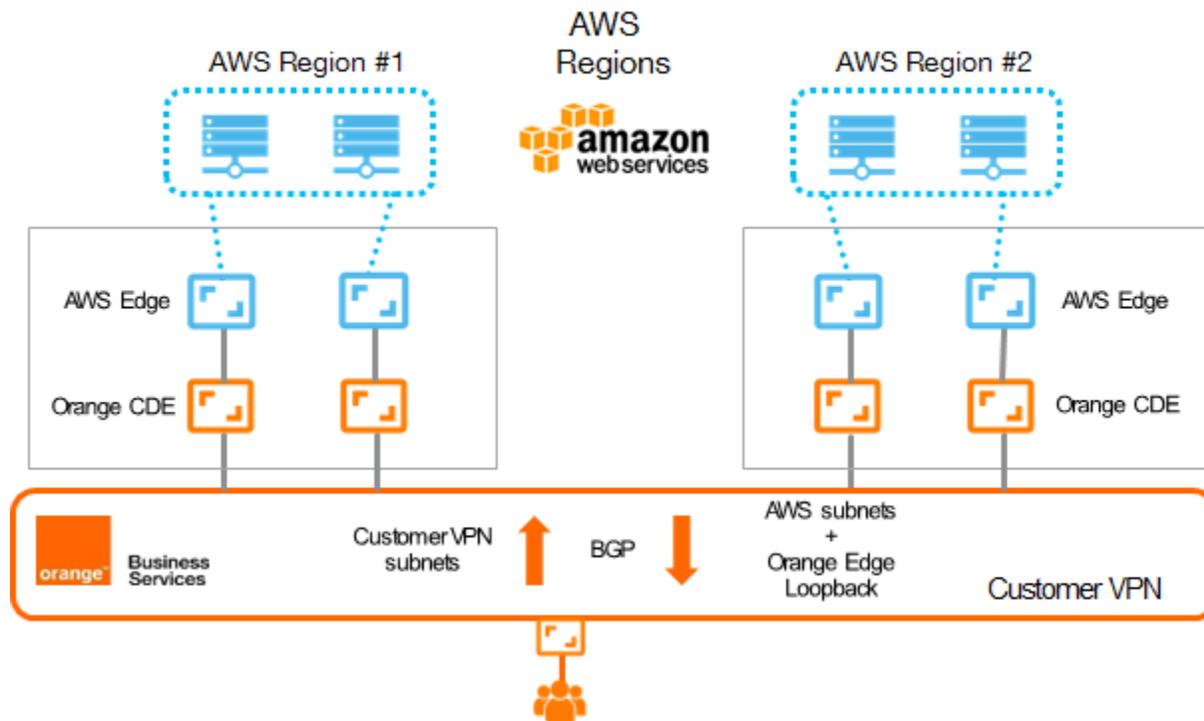


Figure 3: Business VPN Galerie Connections to AWS

As shown in Figure 3, each AWS Region is connected directly to the BVPN Galerie network and BVPN Galerie is connected to AWS through the CDE (Orange side) and the Direct Connect (AWS side) connections. In case of multiple AWS VPCs, each VPC will be connected independently to the customer VPN.¹⁹

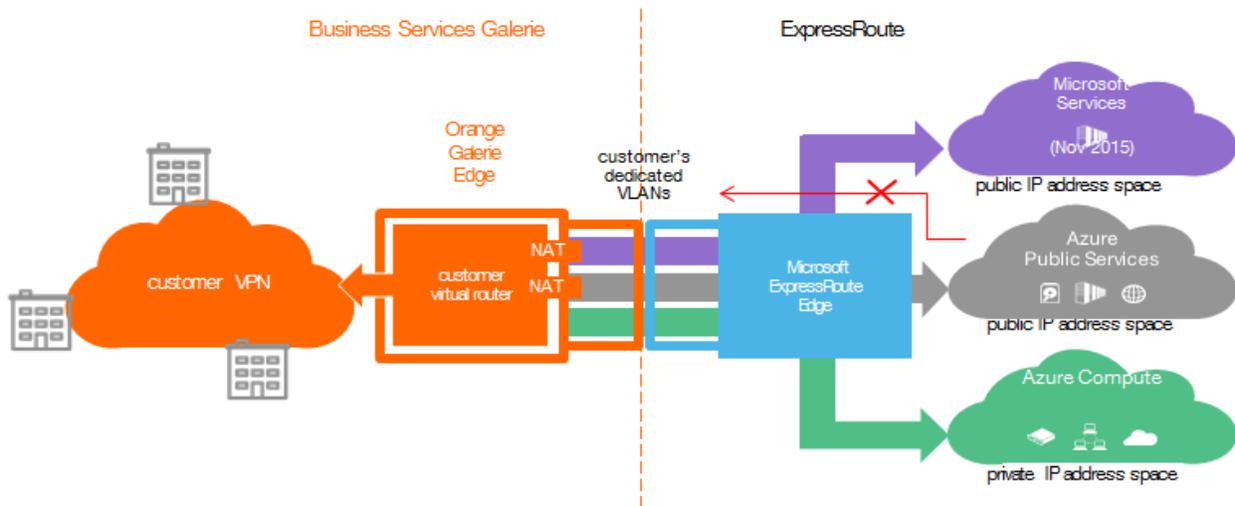


Figure 4: Business VPN Galerie Connections to the Services Accessible via ExpressRoute

The Orange Business VPN Galerie provides a foundation for the Managed Multi-Cloud Orchestration Platform that offers customers choices of public and private cloud services and access to various cloud services with the same user experience control. As a result, it presents a secure way to migrate applications and application data to cloud provider services. The performance and reliability guaranteed at the Business VPN Galerie SLA levels are critically important for planning the migration to cloud or application migration between different cloud service providers.

2.2 Migrating Unstructured Data to the Cloud

This category of migration solutions includes both migrating unstructured application data and migrating VM images encapsulating applications with their data. As many of these migration methods require application downtime, the data is synchronized with the source instance of the application after creating the new application environment in the cloud. As already mentioned (Section 2.1.3), there are several main approaches for unstructured data migration:

- Migrate a server image (image – a file which contains a complete operating system and, in most cases, server management tools and middleware software). As the guest OS and device drivers used in a cloud are unlikely to be an exact match of those in the legacy environment, the server should be reconfigured in order to meet the cloud environment standards. The existing server management tools need to be replaced by tools specific to a given managed cloud. The next steps are to make required changes in the application configuration and then synchronize the data with the source instance of

the application to reflect the data changes taking place during the migration when the source instance is kept online. Server migration may include conversions like Physical-to-Virtual (P2V), Physical-to-Image (P2I), or Virtual-to-Virtual (V2V). P2V and P2I conversions act similarly: P2V instantiates the image into a target hypervisor or cloud whereas P2I captures the image as a file and stores it on disk.

- Create a VM using the cloud image catalog, install the application, and copy the application data from the source system. Alternatively, the application data can be recreated by restoring them from the backup that is done using the cloud backup services (see Section 2.3.3). An advantage of using an image from the cloud catalog is that it already has the OS and management tools properly configured. However, the new VM's kernel, software packages, and library versions may differ from those in the legacy environment. As a result, some modifications may be required to run the application properly in the new VM.
- Online VM migration (Sections 2.2.1.3 and 2.2.4)

VM cloning can be hot/online – that is, done on a running system or cold/offline when a clone is created for a VM taken offline. The time of creating an online clone depends on the system activity level. Offline cloning requires less time and has greater chance for creating a consistent image.

There are several commercial tools to enable migration of workloads to virtualized environments. These tools make a complete copy of the existing server software by means of capturing the contents of disks and, optionally, of server hardware by documenting source HW component details. For example, Novell PlateSpin Migrate²⁰ operates at the image level for its “move workload” and “copy workload” features. PlateSpin Migrate provides a “server synch” feature to synchronize deltas between the original and target “workloads”. Vision Double Take Move²¹ gives users an option of migrating either an entire server image or just data to the virtualized environment. ZConverter Cloud Migration supports OpenStack Cloud Migration and CloudStack Cloud Migration by using ZConverter imaging migration technology.²² On-premises physical or virtual servers can be migrated to OpenStack or CloudStack hypervisors (KVM, Xen, and VMware) without reinstalling OS and applications.

Sometimes migration to cloud requires VM conversion, for example, from Hyper-V VM to VMware VM, and it can be done using tools like VMware vCenter Converter.²³

Table 1 provides a summary of the server migration methodologies.

	Scenario 1	Scenario 2	Scenario 2a	Scenario 3	Scenario 3a
Existing server type	Virtual	Virtual	Virtual/ Physical	Physical	Physical
Target server type	Virtual	Virtual	Virtual/ Physical	Virtual	Virtual
Data Amount	Small/Medium, <1 TB	Large, > 1 TB	Large, > 1 TB	Small/ Medium, <1 TB	Large, > 1 TB
Server Function	Application servers	VMs with large data volumes	File, Database, Mail servers	Application/ Middleware	Physical servers with large data volumes
Migration solution	Create a clone as OVA file and securely copy it to the cloud environment. Resync data by replication over WAN or a dedicated link (see 2.1.4) by using a replication tool e.g. vSphere Replication	Clone VM as OVA file, encrypt and move it to a portable storage system, ship the storage system to the target site and use it for seeding the replication. Start replication to	Migration Type IIIa (see 2.1.3). Set-up a new server using the cloud catalog, install app and import the data transferred using a portable storage system. Use storage	P2V conversion, copy the OVA file to the cloud and import to vSphere cluster. Resync data with replication over WAN or a dedicated link (see 2.1.4) using a replication tool e.g. vSphere	Use tools to convert the server without app data to virtual image (P2V), copy over WAN and import the image to the cloud platform. Move the app data separately using a portable

		synchronize the VMs.	“swing” array for large data volumes (>10 TB)	Replication	storage or replication at the app level. Use storage “swing” array for large data volumes (>10 TB)
--	--	----------------------	---	-------------	--

Table 1: Methodologies of Server Migrations to Cloud

In relation to Scenario 2a, please note that cloud service offerings are not limited to virtual servers – for example, Internap offers bare-metal servers dedicated servers to meet requirements of high-performance workloads.²⁴

If you keep your own licenses, you need to consider whether the vendor provides any options for transferring the licenses to the cloud. For example, Red Hat offers the Red Hat Cloud Access to transfer subscriptions for Red Hat software to a Red Hat Certified Cloud Provider.

2.2.1 • V2V Cloud Migration

2.2.1.1 • Offline Migration Using vSphere Replication

Let’s consider V2V cloud migrations (Scenarios 1 and 2) in more detail. There are many solutions for V2V inter-site, data center-to-cloud migrations. As the scope of this article does not allow us to have a deep-dive review of them, we will consider just an example of migration of VMware ESXi VMs. However, the general steps are applicable to other hypervisors (Hyper-V, Xen, and KVM) and other migration tools.

After the pre-migration steps overviewed in Section 2.1 are complete, move groups are identified. You have to decide how many VMs can be migrated simultaneously; ideally an entire move group. While migrating many VMs concurrently allows you to use the network more efficiently, the migration time of each individual VM increases, resulting in high probability of migration failures due to network issues since network conditions, such as available bandwidth and congestion, fluctuate over time.

Once a source VM is cold-cloned in ova format, you have to decide how the image will be moved to the cloud. The image size, available network bandwidth and migration time targets are the factors to take into account. Small images can be copied by using, for example, an FTP

server such as FileZilla, BrickFTP, Wing FTP Server, and Cerberus FTP Server. You may choose to use managed file transfer (MFT) services with industry standard encryption and secure protocols for transferring your sensitive data. MFT services can provide detailed reports that can be used for PCI DSS, HIPAA, SOX, and GLBA compliance, if needed. Using MFT for your migration to cloud may be a cost-effective option if you build a hybrid cloud and expect an intensive data transfer between your on-premises environment and the cloud. You may also consider WAN optimization solutions like Riverbed SteelHead, Cisco Wide Area Application Services (WAAS and Virtual WAAS), and Silver Peak physical (NX) and virtual (VX) WAN optimization appliances. Network solutions for cloud migration are reviewed in Section 2.1.4.

It is advisable to use a staging zone where migrated VMs are reconfigured to meet the standards of your cloud environment before moving them into production. The staging zone includes a vCenter server and datastores. Various solutions can be used for data synchronization between the source and target VMs. For example, the vSphere Replication can copy changed data from a live VM in the on-premises datacenter to a staging zone datastore in the cloud (Fig. 5). The vSphere Replication does an initial full synchronization of the source virtual machine and its replica copy. For large VMs and low inter-site bandwidth, a seed copy of data can be placed at the destination to reduce the amount of time and bandwidth required for the initial replication. A seed copy of a virtual machine consists of a virtual machine disk file that can be placed at the target location through almost any mechanism (Table 1).

A portable storage system, for example, Synology NAS DiskStation 1815+, can be used for transfer of a high data volume. VM images are encrypted and copied to the NAS system that is shipped in a locked container to the cloud provider to copy the images to a datastore in the cloud environment.

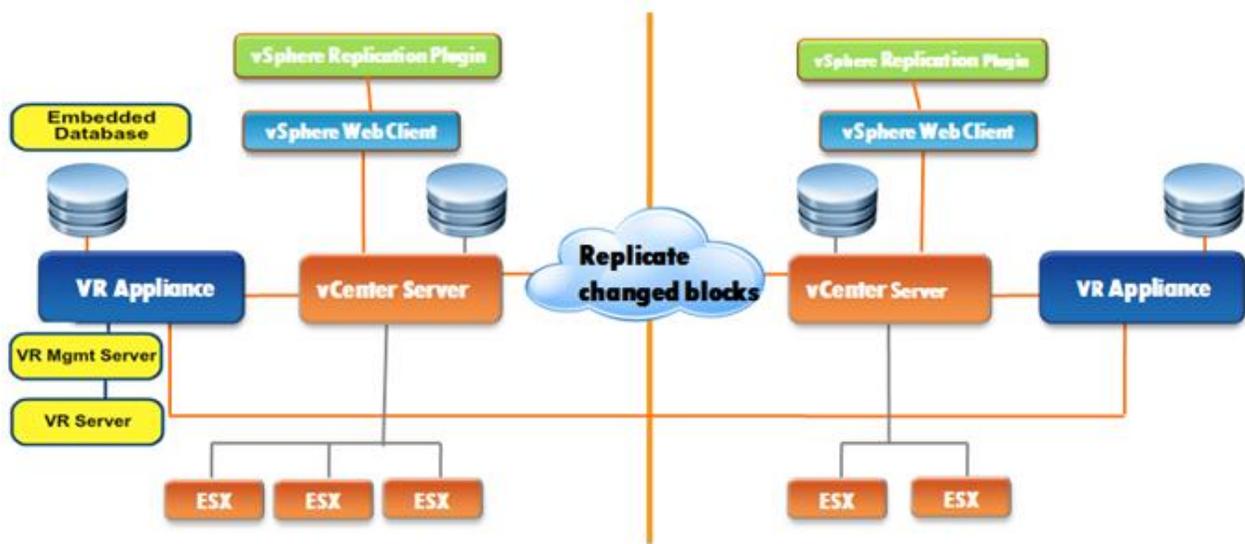


Figure 5: Using vSphere Replication for Migration to Cloud (Ref.25)

Once the synchronization is complete, the VM is imported into the staging zone vCenter Server, powered up, and reconfigured as needed. Once all the VMs in a given move group are validated, they are shut down, the staging zone datastores containing their images are unmounted, rezoned, and mounted to the production cluster in the cloud environment. Alternatively, data between the datastores can be copied at the storage array level. After importing the VM into the production vCenter Server, the data is synchronized using vSphere Replication, the replication stopped, and finally the application functionality is validated before redirecting the users to the new server in cloud.

2.2.1.2 •Migration to VMware vCloud Air Using vCloud Connector

If you use VMware vCloud Hybrid Service, VMware vCloud Connector (vCC) allows you to connect your local datacenter to a cloud provider datacenter and initiate transfer of vApps or VMs and templates to and from your datacenter and between clouds that have been added to your vCloud Connector instance. Objects can be copied between two vSphere clouds, between two vCloud Director-based clouds, and between vSphere and vCloud Director-based clouds.

vCC has a feature – Datacenter Extension – that can be used to extend a private datacenter to a public vCloud so that you can move workloads (virtual machines and vApps) from your private datacenter to a public vCloud. While moving a workload, vCC stretches its private network boundary to the public vCloud, so that it continues to get all its networking properties from the private network. As a result, the VMs retain their original network settings and continue to have

the same IP addresses and MAC addresses in the public vCloud as if they were still in the source private datacenter.

The vCloud Connector Offline Data Transfer feature enables you to migrate large numbers of VMs with terabytes of data, vApps, and templates between private and public clouds. Using vCC, data is exported to a portable pre-configured NAS appliance that is shipped to the cloud provider that imports the data into the vCloud Air Service. The vCC encrypts the data before writing it to the NAS device, ensuring a secure transfer. The vCloud Air Service provider uses vCC to import the data to a vCenter server datastore associated with the vCloud Air service. The vCC uploads the data, decrypts it, and moves it into the vCloud Air service datacenter. The vCC then applies the network, power, and deployment settings that are specified during export.

In a similar fashion, the AWS Connector for vCenter can be used to migrate a customer's virtual machines to Amazon EC2.

2.2.1.3 •Cloud Migrations Using Long-Distance vMotion

VMware vSphere 6.0 introduces Long Distance vMotion enabling VM migration over long distances. As vMotion supports RTT (round-trip time) latency of 100 milliseconds or less between hosts (previously only 10ms), it allows you to move VMs from your datacenter into a cloud datacenter. Long distance vMotion is supported with both no shared storage infrastructure and with shared storage solutions such as EMC VPLEX[®] Geo which enables shared data access across long distances. Long Distance vMotion requires L2 network connectivity because it keeps the IP address of the VM. The vMotion network can be secured either by being dedicated or encrypted (VM memory is copied across this network).

2.2.2 •Data Migration to Cloud: VPLEX

EMC VPLEX extends VMware functionality by enabling vMotion to be used across datacenters to move and relocate VMs, applications, and data over distance.²⁶

Online application migration between data centers

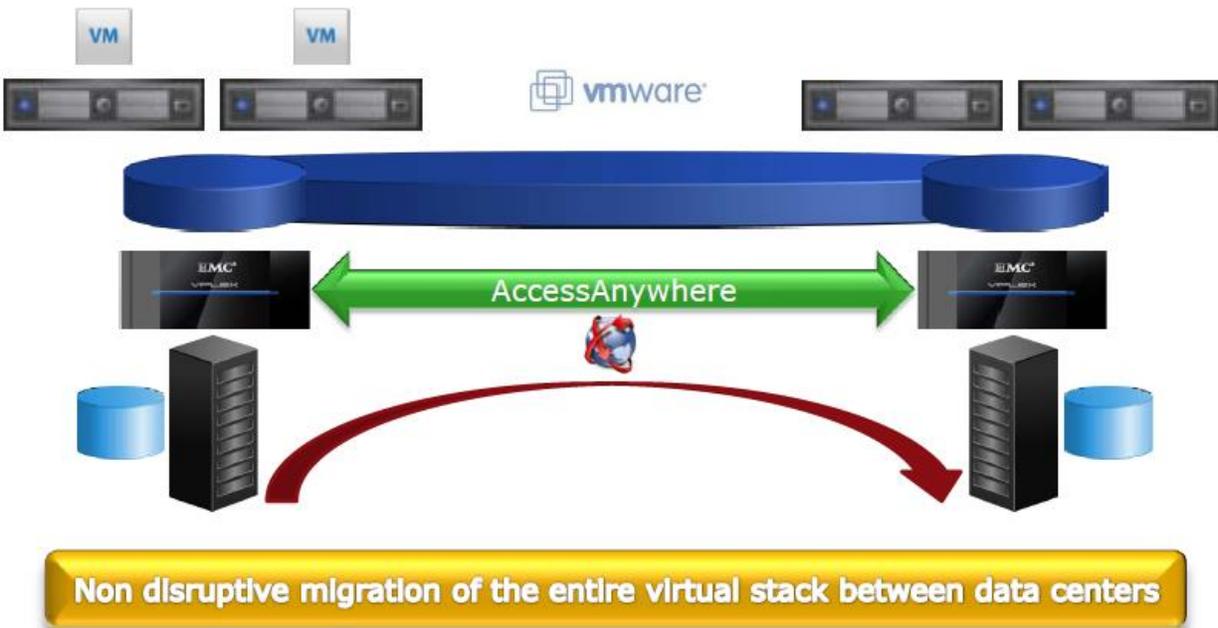


Figure 6: VM Mobility with VPLEX Metro (Ref.27)

VPLEX Metro makes it possible to migrate online the entire virtual stack between data centers. EMC VPLEX Geo provides the asynchronous replication over distance and enables information in the public, private, or hybrid cloud to be shared, accessed, and moved within, across, and between data centers. However, as the potential value of VPLEX Geo will be diminished by the upcoming Active-Passive replication technologies with VMware VVOLs and vMotion over long distances (see Section 2.2.1.3), EMC announced End-of-Life (EOL) for VPLEX Geo in July 2015.

2.2.3 • Migrating Unstructured Data to the Cloud Using Provider's Export/Import Services

When you need to move large amounts of data to the cloud, transferring data across the Internet can be cost and/or time prohibitive. The major cloud service providers offer using import/export service as a solution for large data transfers.

With *the AWS Import/Export Service*, users load their data on USB 2.0 or eSATA storage devices and ship them via a carrier to AWS. AWS then uploads the data into the designated buckets in Amazon S3.

In a similar way, when you use *the Microsoft Azure Import/Export* service to transfer a large set of file data into Blob storage, you can copy data to hard drives, encrypt the data, and send the disks to an Azure data center, where your data will be uploaded to your storage account.²⁸

The Google Cloud Storage Transfer Service can be used to quickly import online data into Google Cloud Storage. It can be used to move data from other storage providers as the data source can be an Amazon Simple Storage Service (Amazon S3) bucket or an HTTP/HTTPS location. The Google Cloud Platform also suggests using *the Offline Media Import/Export* service that is a third party (Iron Mountain in NA, Prime Focus Technologies in EMEA and APAC) solution to load data into Google Cloud Storage by sending customer physical media, such as hard disk drives (HDDs), tapes, and USB flash drives, to a third party service provider who uploads data on the customer's behalf.

2.3 Migrating Structured Data to the Cloud

2.3.1 General Considerations for Migrating DBs to the Cloud

Your selection of migration method depends upon the application/DB tolerance for downtime, the size and complexity of the DB, and the bandwidth of the connection to the cloud. Network bandwidth challenges for cloud migrations have been reviewed in Section 2.1.4. Cloud Export/Import services offered by cloud service providers (see Section 2.2.3) can be used to move large DBs, for example, a 100 TB Oracle DB with 50 GB of daily changes. A full backup of the DB is done to tapes, a portable storage system, or to a “swing” storage array. The encrypted data is then sent to the cloud provider who restores the DB from the full backup. Following that, the cloud DB instance is synchronized with the source DB using native DB tools.

There is one more important aspect of migrating to cloud, namely, what computer tier and cloud storage tier you select for your migration. Using higher pricing tiers for migrating large DBs will allow you to reduce the migration time – it means less application downtime and network usage charges. The pricing tier determines how much you are charged for the DB you create in cloud. For cost optimization, you can use higher pricing tiers to achieve shorter deployment times and then reducing the pricing tiers after the deployment has been completed. If your cloud provider charges you per hour, a shorter migration time may offset the additional cost of using the higher pricing tier.

2.3.2 Migrate to Cloud-Based DB Service or to Run DBs on Cloud VMs?

When you consider moving your DBs to the cloud, you have two main options: either migrate your data to the DB service your cloud provider offers – such as Azure SQL Database, Oracle Database Cloud Service, and Amazon Relational Database Service – or move your DB to a VM in the cloud. Your decision depends on whether your DB is comparable with the DB services offered by the provider, the functionality of the cloud-based DB services and, of course, cost efficiency.

Essentially it is a question of what solution you choose for hosting your DBs in the cloud:

- A database native to the cloud, also known as a platform as a service (PaaS) database or a database as a service (DBaaS) that is optimized for software-as-a-service (SaaS) application development. It offers compatibility with the majority of the features of conventional SQL DBs.
- DB on a VM: a DB is installed and hosted in the cloud on VMs – it is known as an infrastructure as a service (IaaS) (Section 2.1.1).

It is easy to create an AWS or Azure VM and then installation of a DB is almost the same process as installing it in your datacenter. As the range of the VM OSes supported by cloud providers is broader compared to the DB services, this option is more flexible. For example, you are not limited to using a Microsoft OS as in deploying your DB to the Azure SQL DB service – you can use a Linux flavor supported on an Azure VM and install a MySQL instance. However, running DBs on the cloud-based VMs means you take the responsibility for managing them (backups, HA, security, DB patching, etc.) whereas DB services offered by cloud providers typically include a lot of the features you need as out-of-the box.

Of course, the service cost is a key component in your decision making. Usually cloud-based DB services do not include the license cost explicitly as it is built into the monthly recurring charge billed to you. On the other hand, if you install the DB onto cloud-provider VMs, then you need to license the DB for each virtual machine instance. Furthermore, your cloud provider should be on the DB vendor list of the approved providers of “Authorized Cloud Environment”. In other words, the cloud platform should be supported by the DB vendor. If you need to run your DB on cloud-based VMs for a short period of time, the question is whether the DB vendor offers short-term licenses. While such licenses are less expensive – for example, a one-year term

license can be purchased from Oracle for as little as 20 percent of the product list price – using cloud-provider DB services may be more cost-effective in this case.

You may also be able to migrate your existing licenses to some cloud service providers. For example, you can move your MS SQL server license to Azure for running SQL Server in an Azure VM by using License Mobility through Software Assurance on Azure. In this case, you only pay for Azure compute and storage costs associated with the VM.

2.3.3 Migrating DBs by Restoration from Backup and Synchronization

Alternatively, the source database can be backed up to the cloud and then the backup data can be restored to the destination DB in the cloud. Oracle offers the Oracle Secure Backup (OSB) Cloud Module (a separately licensed product) enabling an Oracle Database to send its backups to Amazon S3. It is compatible with Oracle Database versions 9i Release 2 and above. The Oracle Secure Backup Cloud Module uses the Oracle Recovery Manager (RMAN) SBT interface allowing for integration of external backup libraries with RMAN. Therefore, RMAN compression and encryption capabilities can be used for faster and secure data transfer. Recently, Oracle introduced its own Oracle Cloud Database Backup Service (ODBS), enabling backup storage in the Oracle Cloud.

While both cloud database backup services have similar features, there are some differences in the default options for security. Oracle Database Cloud Backup Module (ODCBM) enforces mandatory RMAN encryption of the backup data – the backup data is encrypted at the source, in transit, and at rest in the Cloud (RMAN encryptions to backup to ODBS do not require licensing Advanced Security Option). In contrast to ODBS, the default option in the AWS solution is not securing the data at rest or in flight. Of course, encryption for both data at rest and in-flight can be enabled (and it should be enabled to mitigate risk of theft or unauthorized access). The Amazon S3 Server Side Encryption can be used for the data backup in S3.

To move an existing Microsoft SQL Server deployment to Amazon Relational Database Service (RDS), Azure SQL database, or other public cloud service, some solutions available on the market can be used.²⁹⁻³² For example, CloudBerry Backup for MS SQL Server uses native SQL Server backup routines to perform database backup and then move it directly to one of the public cloud storage services including: Amazon S3, Amazon Glacier, Windows Azure, HP Cloud, Rackspace, OpenStack, and others.³³

MS SQL Server supports storing backups to the Microsoft Azure Blob storage service by specifying URL as the backup destination. This feature is available in SQL Server 2012 SP1 CU2 or later and is also referred to as SQL Server Backup to URL. The backup file which is stored in the Microsoft Azure Blob storage service can be made directly available to either an on-premises SQL Server or another SQL Server running in a Microsoft Azure VM, without the need for database attach/detach or downloading and attaching the VHD.

Once the data is restored from backup to an Azure SQL database, the data needs to be synchronized with the original on-premises database. This can be done using SQL Data Sync³⁴ which synchronizes data across multiple SQL Server and Windows Azure SQL Database instances. SQL Data Sync uses a Hub and Spoke architecture. The target Azure SQL DB can be assigned as the Hub. DBs synchronizing their data form a so-called sync group which is used to specify the databases, tables, and columns to synchronize as well as the synchronization schedule. When a database is added to a sync group, the sync direction is defined. By default, the sync is bi-directional. Both sync to the Hub and sync from the Hub are possible options.

In order to synchronize an on-premises SQL Server Database to a Windows Azure SQL Database or vice versa, SQL Data Sync Agent needs to be installed and configured on the local machine. Once the Sync Agent is installed, a new Windows service (Microsoft SQL Data Sync) is available. The required Sync Agent configuration includes information on where the on-premises data source can be found, the credentials used to access the data source, and any other information that the synchronization requires. To connect the SQL Azure Sync Agent with the one on the premises, a key needs to be shared between them.

2.3.4 Migrating DBs from MS SQL Server to Azure SQL Databases

To illustrate database migration solutions to the cloud, let us consider migrating MS SQL DBs to the Azure SQL Databases as an example.

Migration Solution	Solution Description	Use Cases
Using SQL Server Management Studio (SSMS) to migrate a compatible database	The database is deployed to Azure SQL Database using SSMS. The database can be deployed directly or exported to a BACPAC file (see 2.3.4.1) which is then imported to create a new Azure SQL database.	Use when the database schema requires upgrade and the changes can be handled by the SQL Azure Migration Wizard.
Migrate a near-compatible database using SQL Azure Migration Wizard	The SQL Azure Migration Wizard generates a migration script containing schema or schema plus data in BCP format (see 2.3.4.2).	Use when the database schema requires upgrade and the changes can be handled by the Wizard.
Update database schema off-line using Visual Studio and SAMW and deploy with SSMS	The source DB is imported into a Visual Studio DB project for processing offline. After that, SQL Azure Migration Wizard is run across all the scripts in the project to apply a series of transformations and corrections. The project is targeted at SQL DB V12 (preview) and built and any remaining errors are reported. These errors are then resolved manually using the SQL Server tool in Visual Studio. Once the project builds successfully, it is published back to a copy of the source DB. This updated DB is then deployed to Azure SQL DB by using SSMS. If schema-only migration is required, the schema can be published directly from Visual Studio directly to Azure	Use when the database schema requires more changes than can be handled by SAMW alone.

	SQL DB.	
--	---------	--

Table 2: Migrating to Azure SQL DBs (Ref. 35)

2.3.4.1 Migrating DBs to Azure SQL Databases by Using SQL Server Management Studio

In cases where a database schema is already compatible with the latest Azure SQL Database Update, the migration to Azure is relatively simple as only copying the database to Azure is required. Using SQL Server Management Studio (SSMS) makes it possible to migrate a DB in a single-step process by ‘deploying’ the database to Azure SQL Database. The “Deploy Database to Microsoft Azure Database” wizard in SSMS migrates a compatible SQL Server database directly into the Azure SQL Database server (Figure 7).

An alternative two-step process uses the SQL Database Import Export Service by first exporting a data-tier application (DAC) file (known as a BACPAC file for the file extension, .bacpac) and importing it to an Azure SQL Database server to create a new database. Unlike the database copy requiring source and destination database servers, the BACPAC file is standalone and can be moved to wherever it is needed. Creating a BACPAC file includes copying the object definitions from the database following a bulk copy of the data from the user tables. The resulting BACPAC file contains sufficient information to completely replicate a database. After the BACPAC file is copied into the Azure Blob storage service, DAC import can be used to create a new database containing all of the objects and data.

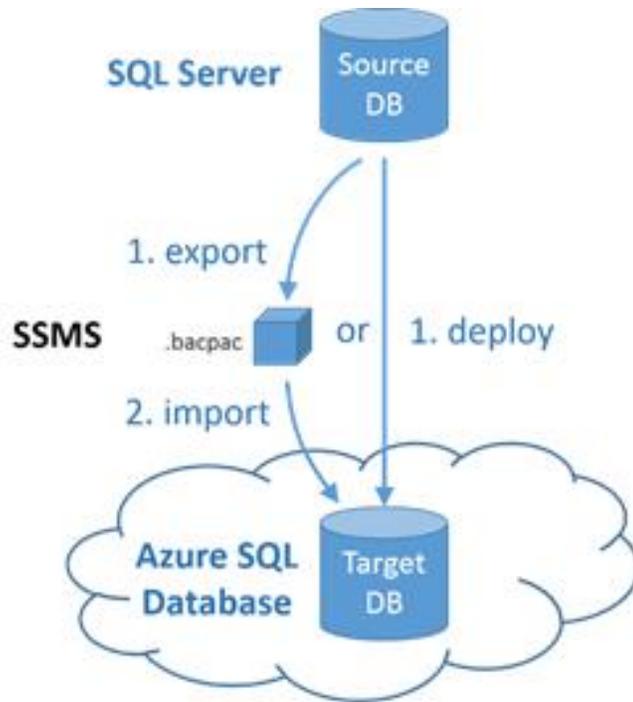


Figure 7: Migrating a Compatible DB Using SSMS (Ref.35)

Deploying directly from SSMS will always deploy the schema and data, while the two-step export/import process always deploys the schema and provides an option to deploy data from all or a subset of the tables.

2.3.4.2 Migrating DBs to SQL Database by SQL Database Migration Wizard

The SQL Azure Migration Wizard can be used to generate a T-SQL (Transact-SQL) script from a source database which is then transformed by the Wizard to make it compatible with the latest SQL Database Update and then connect to Azure SQL Database (an empty DB to be created) to execute the script against an empty Azure SQL database. The script can be generated with schema only or can include data in the BCP (Bulk Copy Program) format (the bcp utility (Bcp.exe) is a command-line tool that uses the BCP API). The SQL Database Migration Wizard allows you to select or exclude specific SQL objects and then it creates SQL scripts suitable for Azure SQL Database and makes it possible to migrate data between on-premises SQL Servers and Azure SQL Database servers.

2.3.4.3 Migrating DBs by Updating the Source DB and Then Deploying It to Azure SQL DB

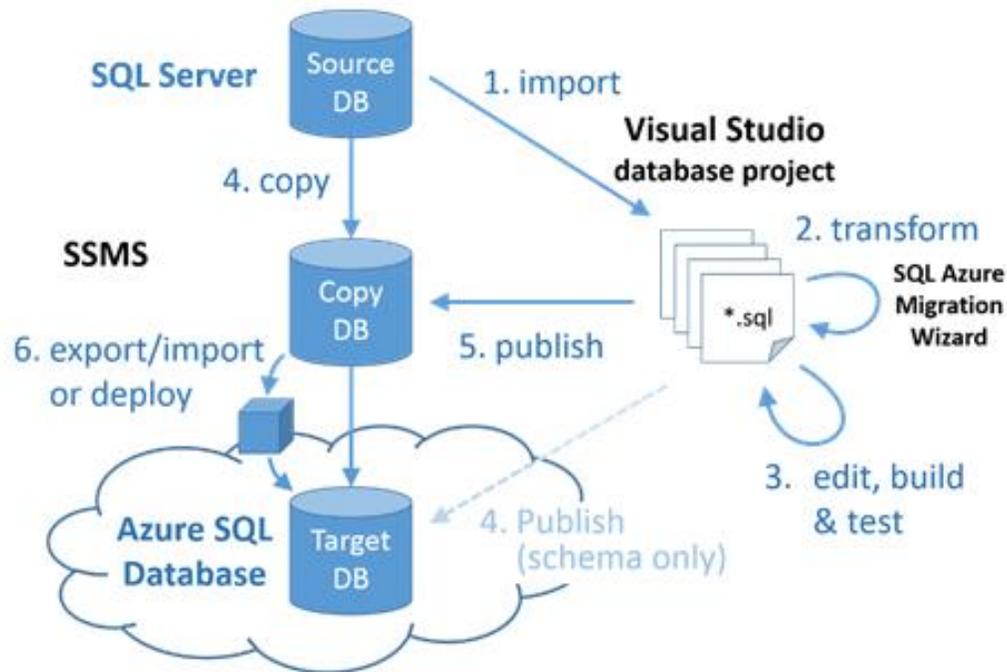


Figure 8: Update DB Schema Off-line Using VS and SAMW and Deploy with SSMS (Ref.35).

The solution presented in Figure 8 can be used for migrating a database to the preview of the latest Azure SQL Database Update V12 when schema changes that cannot be addressed using the SQL Azure Migration Wizard (SAMW) are required. Changes may be needed if the database uses SQL Server features that are not or not yet supported in Azure SQL Database. The solution includes creation of a database project from the source database by using Visual Studio. The project's target platform is then set to Azure SQL Database V12 and the project is built to identify all compatibility issues. Many but not all compatibility issues can be fixed by SAMW, so it is used to process all the scripts in the projects as a first pass. Using SAMW is optional but highly recommended as SAMW will detect compatibility issues within the body of functions, stored procedures, and triggers which will not otherwise be detected until deployment.

Remaining issues will be identified by building the project after processing the script files with SAMW. They must then be addressed manually using the T-SQL editing tools in Visual Studio. After the project building is completed successfully, the schema is published back to a copy (recommended) of the source database to update its schema and data in-situ. The updated database is then deployed to Azure, either directly or by exporting and importing a BACPAC file, using the techniques described in section 2.3.4.1.

As the schema of the database is updated in-situ before deploying to Azure, it is strongly recommended to perform this update on a copy of the database. The Visual Studio Schema Compare tool can be used to review the full set of changes that will be applied to the database before publishing the project.

2.3.4.4 Migrating DBs by Using Transactional Replication

If you need to keep your source SQL DB in production during the migration, you can use SQL Server transactional replication to migrate it to the Azure SQL DB.³⁶

Transactional replication is based on using the SQL Server Snapshot Agent, Log Reader Agent, and Distribution Agent. The Snapshot Agent prepares snapshot files containing schema and data of published tables and database objects, stores the files in the snapshot folder, and records synchronization jobs in the distribution database on the Distributor.

The transaction log is monitored by the Log Reader Agent of each database. The Log Reader Agent copies the transactions marked for replication from the transaction log into the distribution database. The copies of the initial snapshot files from the snapshot folder and the transactions held in the distribution database tables are copied to the Subscriber (Azure SQL DB) by the Distribution Agent. With the SQL Server 2016 preview, it is now possible to configure your Azure SQL Database as a transactional replication subscriber to the on-premises publisher (see Figure 9).

When using transactional replication for the migration to Azure SQL DB, all changes in the data or schema that take place during the migration are replicated to the target Azure SQL DB. After completing the migration, the applications will be pointed to the target Azure SQL DB instead of pointing them to the source on-premises DB. Once transactional replication copies all the changes from the on-premises database and all the applications point to Azure DB, the replication can be stopped and the Azure SQL DB becomes the production system.

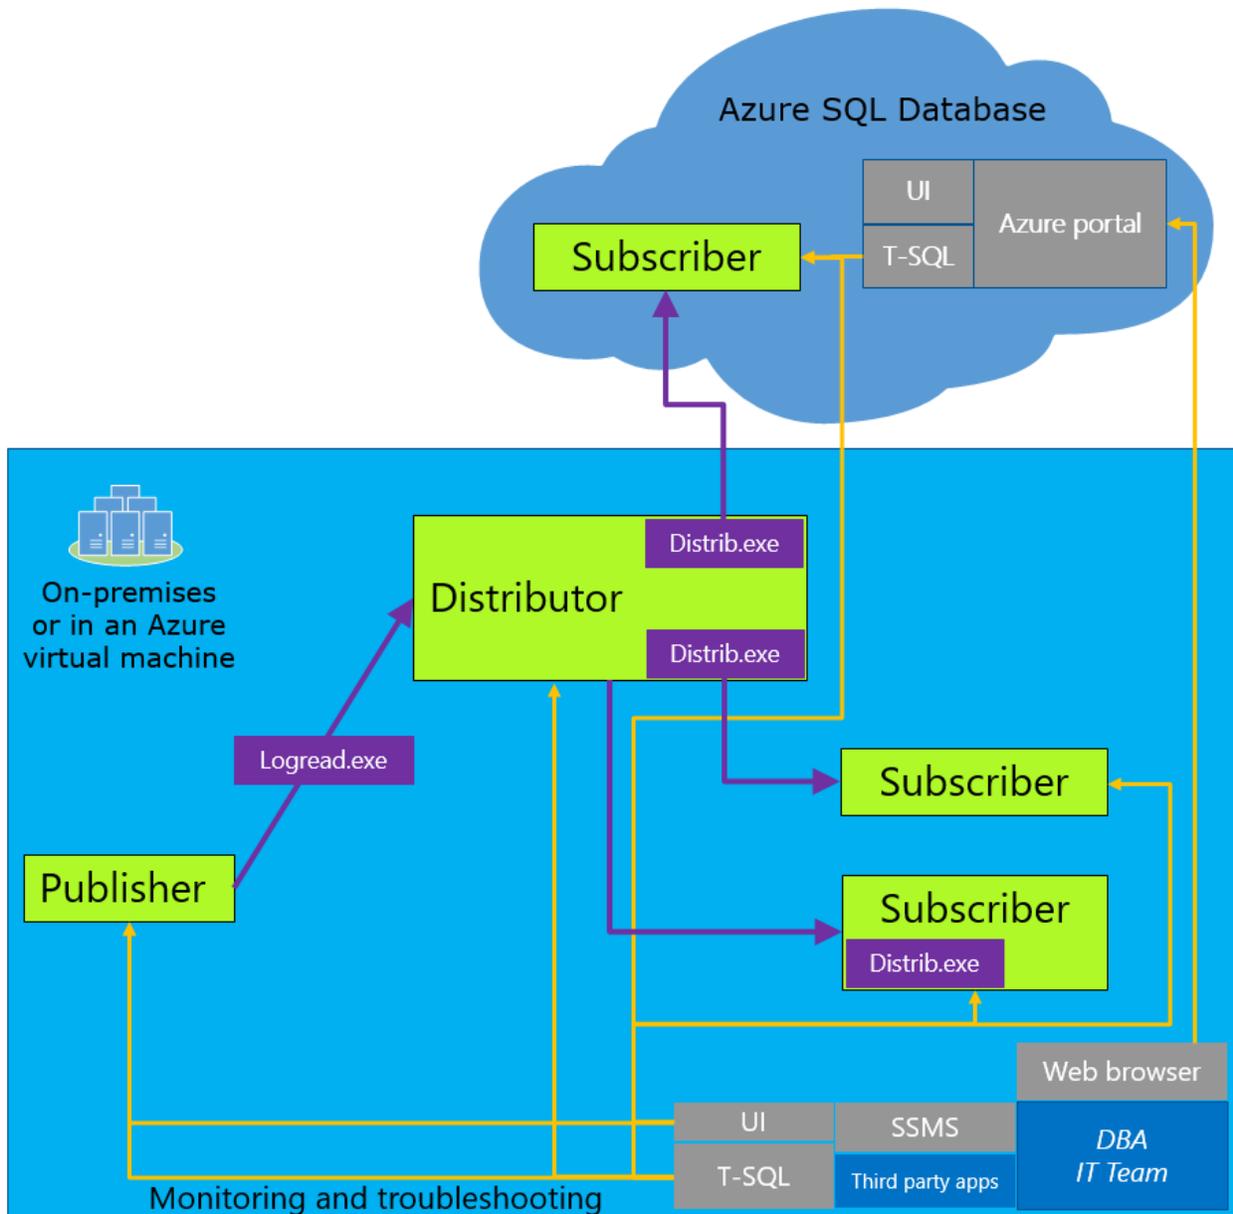


Figure 9: Migrating a Compatible DB Using Transactional Replication (Ref.36)

Before a new transactional replication Subscriber can receive incremental changes from a Publisher, the Subscriber must contain tables with the same schema and data as the tables at the Publisher. The initial dataset can be copied to the target DB either by using a snapshot that is created by the Snapshot Agent and distributed and applied by the Distribution Agent or restoring from the source DB backup (see the section: Migrating DBs by restoration from backup and synchronization).

2.3.5 Migrating SQL DBs to Azure VMs

Table 3 lists the typical methods for migrating SQL DBs to Azure VMs.

Migration Solution	Source database version	Destination database version	Source database backup size constraint	Notes
Use the Deploy a SQL Server Database to a Microsoft Azure VM Wizard	SQL Server 2005 or later	SQL Server 2014 or later	Up to 1 TB	Fastest and simplest method to be used whenever possible to migrate to a new or existing SQL Server instance in an Azure VM.
Use the Add Azure Replica Wizard	SQL Server 2012 or later	SQL Server 2012 or later	Storage limit for Azure VM	Minimizes downtime, use when you have an AlwaysOn on-premises deployment.
Use SQL Server transactional replication (see 2.3.4.4)	SQL Server 2005 or later	SQL Server 2005 or later	Storage limit for Azure VM	Use when you need to minimize downtime and do not have an AlwaysOn on-premises deployment.

Perform on-premises backup using compression and manually copy the backup file into the Azure VM (see 2.3.3)	SQL Server 2005 or later	SQL Server 2005 or later	Storage limit for Azure VM	Use only when you cannot use the wizard, such as when the destination database version is less than SQL Server 2012 SP1 CU2 or the database backup size is larger than 1 TB (12.8 TB with SQL Server 2016).
Perform a backup to URL and restore into the Azure virtual machine from the URL (see 2.3.3)	SQL Server 2012 SP1 CU2 or later	SQL Server 2012 SP1 CU2 or later	> 1 TB (for SQL Server 2016, < 12.8 TB)	Generally using backup to URL is equivalent in performance to using the wizard and not quite as easy.
Detach and then copy the data and log files to Azure Blob storage and then attach to SQL Server in Azure virtual machine from URL	SQL Server 2005 or later	SQL Server 2005 or later	Storage limit for Azure VM	Use this method when you plan to store these files using the Azure Blob storage service and attach them to SQL Server running in an Azure VM, particularly with very large databases.
Convert on-premises machine to Hyper-V VHDs, upload to Azure	SQL Server 2005 or later	SQL Server 2005 or later	Storage limit for Azure VM	Use when bringing your own SQL Server license, when migrating a

Blob storage, and then deploy a new VM using the uploaded VHD				database that you will run on an older version of SQL Server, or when migrating system and user databases together as part of the migration of database dependent on other user databases and/or system databases.
Ship hard drive(s) using Windows Import/Export Service (see 2.3.6)	SQL Server 2005 or later	SQL Server 2005 or later	Storage limit for Azure VM	Use the Windows Import/Export Service when manual copy method is too slow, such as with very large databases.

Table 3: Main Methods for Migrating SQL DBs to Azure VMs (Ref. 37)

2.3.6 Cloud Provider Services for Migrating DBs

Amazon offers the AWS Database Migration Service to migrate data to and from all widely used commercial and open-source databases. Both same-platform migrations such as Oracle to Oracle, as well as migrations between different database platforms, such as Oracle to Amazon Aurora or Microsoft SQL Server to MySQL are supported. As the AWS Database Migration Service ensures that all data changes to the source database that occur during the migration are continuously replicated to the target, the source database can be fully operational during the migration process. After the initial database migration is complete, the target database will remain synchronized with the source for as long as you choose, giving you the flexibility to switch your applications from the source to the target database at a convenient time, reducing risk and minimizing application downtime. The change data capture feature of the AWS Database Migration Service continuously captures and applies all data changes from the source

database to the target instance after the start of the migration process. This minimizes downtime as the source database can be kept operational during the migration process.

Similar DB migration services are offered by Oracle Database Cloud Services, Microsoft Azure, Google Cloud Platform, IBM BlueMix, and other providers.

2.3.7 Migrating Oracle DBs to Cloud

The general aspects of the DB migration to cloud that we have considered in Sections 2.3.1-2.3.3 and 2.3.6 apply to Oracle DBs migration. The size limitation of this article does not allow us for a more detailed consideration and readers are referred to reviews on this topic.³⁸

Regarding Azure, note that there is no DBaaS offering for Oracle, so if you want to use Azure, you would be responsible for management of the operating systems and DB in the virtual machines. Microsoft provides ready-built images of Oracle database on Windows operating systems, but there are currently no pre-built images for Oracle products on Oracle Linux. However, Microsoft supports Oracle Linux running inside an Azure VM.

3. Inter-Cloud Migration and Workload Portability

While private cloud is the primary type of cloud infrastructure currently in use by many companies, there is a substantial growth of the multi-cloud approach resulting in implementations of hybrid cloud. There are various benefits offered by hybrid cloud as exemplified by the EMC Enterprise Hybrid Cloud (EHC).³⁹ The EHC empowers IT to be a broker of cloud services – providing the control and visibility that IT organizations need, and the on-demand self-service that developers and application users expect.

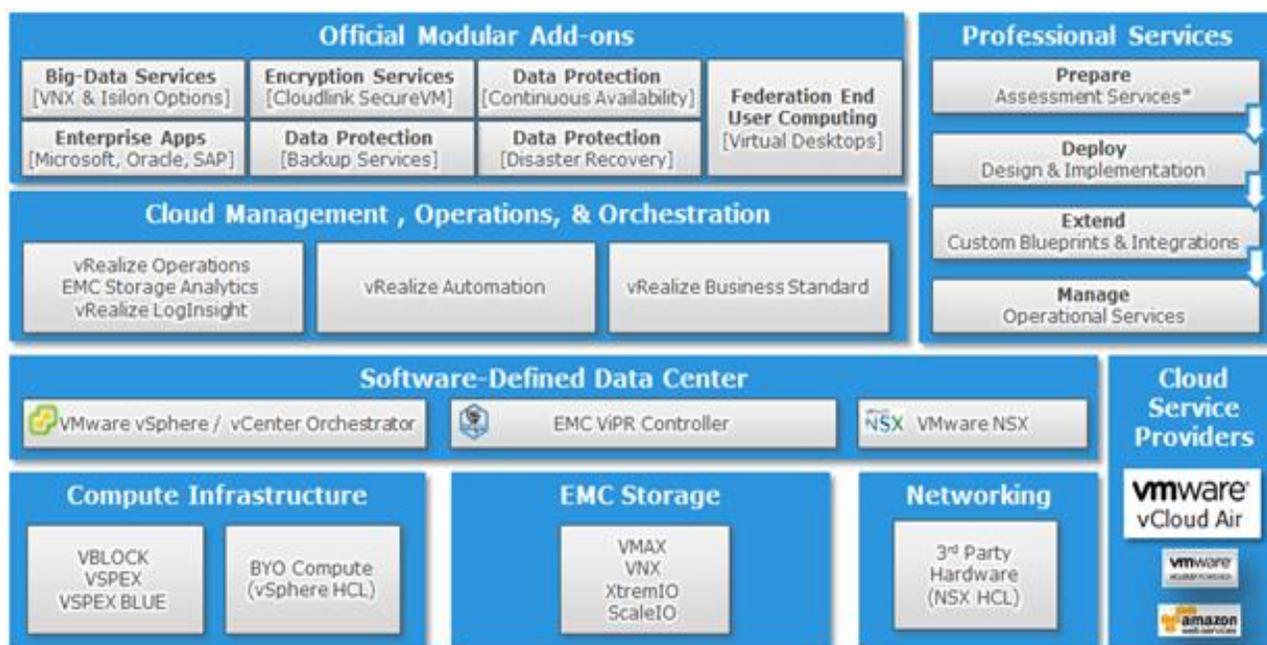


Figure 10: EMC Federation Enterprise Hybrid Cloud (Ref.39)

As a hybrid cloud is a mix of private and public cloud services, customers are concerned with the need for interoperable cloud environments and management solutions that incorporate standard mechanisms to create consistencies across all VMs. The growing demand for hybrid cloud is driving the need for interoperability and openness across on-premises and public cloud environments. The Open Data Center Alliance (ODCA)⁴⁰ considers the interoperability from two perspectives; interconnectability and portability:

- Interconnectability – the parallel process in which two coexisting environments communicate and interact.
- Portability – the serial process of moving a system from one cloud environment to another.

Let us take a brief look at both of them.

3.1 Cloud Interconnectability

Equinix offers a broad choice of cloud service providers, such as VMware vCloud Air, AWS, Google Cloud Platform, and Microsoft Azure, with direct connections to these cloud services via Equinix Cloud Exchange™, cross connects, or Ethernet services.⁴¹

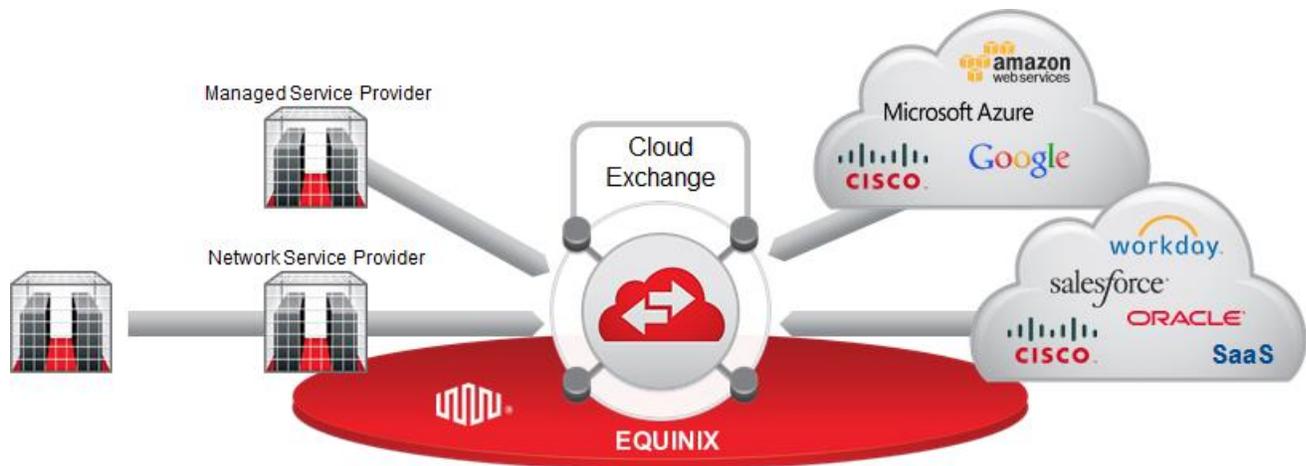


Figure 11: Equinix Cloud Exchange (Ref.41)

Equinix Cloud Exchange provides virtualized, private direct connections that bypass the Internet to provide better security and performance with a range of bandwidth options. This enables companies to build hybrid cloud solutions meeting their business needs. Furthermore, cloud services providers can benefit from connectivity to Orange Business VPN Galerie service on the Equinix Cloud Exchange (see Section 2.1.4).

3.2 Data and Workload Portability in Cloud

Data and workload portability requires the ability to move not only VMs and application data between clouds but the environment metadata as well. Environment metadata is often very cloud provider-specific. Account structures, users, permissions, policies, load balancers, and the like vary from cloud to cloud.

A recently initiated standardization effort from OASIS – the Topology and Orchestration Specification for Cloud Applications (TOSCA)⁴² – has the goal of improving the portability and manageability of applications by composing a service once and playing it on any cloud. TOSCA provides a framework for development of orchestration for the entire application lifecycle by modeling your topology once, and then deploying it to your infrastructure of choice, all while managing, monitoring, scaling, and healing everything in the same place. For example, using TOSCA Cloudify provides out-of-the-box integrations with leading clouds, infrastructure, and tooling (from OpenStack to VMware to Docker using Chef, Puppet, or even SSH with Fabric).⁴³

There are promising developments in using containers (Docker, CoreOS, etc.) for migration between different clouds. Using containers means that programmers may not need to rewrite the code for each new operating system and cloud platform. However, there are some

challenges as well: for example, not all applications are “container-friendly”, containers have security-related limitations, etc.⁴⁴

Another aspect of data portability is the integration and mobility of cloud- and on-premises storage. This portability can be achieved by using cloud storage arrays and gateways. For example, EMC CloudArray[®] pushes data to public or private clouds of the customer’s choice (Figure 12).⁴⁵

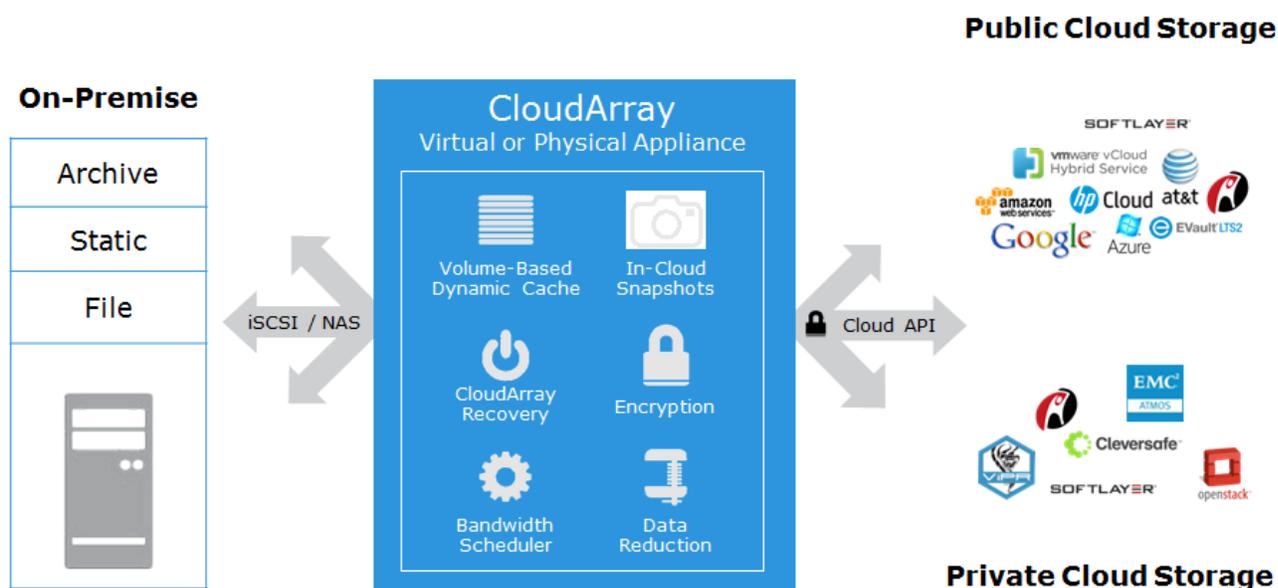


Figure 12: EMC Cloud Array (Ref.45)

A need in cloud storage gateways stems from the fact that cloud service providers like Amazon offer a cloud-based object store with interfaces such as REST or SOAP, but traditional applications expect storage resources with block-based iSCSI and Fibre Channel interfaces or file-based interfaces, such as NFS and CIFS. A cloud storage gateway (major vendors include EMC, Avere, Amazon, CTERA Networks, Emulex, F5 Networks, HP, Microsoft, and NetApp) converts file level/block level requests into cloud storage service provider API’s. This eases migration between cloud service providers without having to rewrite numerous API’s because those API’s are built-in to the cloud storage gateway. Cloud gateways can make cloud storage appear to be a NAS filer, a block-access storage array, or a backup target.⁴⁶

Size limitations of this article do not allow us to discuss these key topics in more detail. However, several discussions of them can be found in references.⁴²⁻⁴⁶

4. Conclusion

Migration to the cloud raises many questions. It is not possible to cover them all in this article whose goal is to consider the practical aspects of cloud migrations. I hope that this overview will help improve the effectiveness of the migration to cloud. Frequently, my customers ask which migration type is most suitable for them. This article provides a review of various cloud migration types and shows that the selection of the best solution depends on the individual needs of each organization.

5. References

1. S. Frey, W. Hasselbring. Model-based migration of legacy software systems to scalable and resource-efficient cloud-based applications: the CloudMIG approach. In Proceedings of the 1st International Conference on Cloud Computing, GRIDs, and Virtualization, p. 155, 2010.
2. M. Gloukhovtsev. Does the Advent of Cloud Storage Mean “Creation by Destruction” of Traditional Storage? EMC Proven Professional Knowledge Sharing, 2013.
3. D. C. Marinescu. Cloud Computing: Theory and Practice. Elsevier, Waltham, 2013.
4. M. Kavis. Architecting the Cloud: Design Decisions for Cloud Computing Service Models. Wiley, Hoboken, 2014.
5. V. Tran, J. Keung, A. Liu and A. Fekete, Application Migration to Cloud: A Taxonomy of Critical Factors, in Proceedings of the 2nd International Workshop on Software Engineering for Cloud Computing, 2011.
6. <http://www.gartner.com/newsroom/id/1684114>
7. A. Khajeh-Hosseini, D. Greenwood, J. W. Smith and I. Sommerville. The Cloud Adoption Toolkit: supporting cloud adoption decisions in the enterprise. Software – Practice and Experience, vol. 42, p. 447, 2012.
8. V. Andrikopoulos, T. Binz, F. Leymann and S. Strauch. How to Adapt Applications for the Cloud Environment: Challenges and Solutions in Migrating Applications to the Cloud. Computing, vol. 95, no. 6, p. 493, 2013.
9. R. Rai, G. Sahoo and S. Mehruz. Exploring the factors influencing the cloud computing adoption: a systematic study on cloud migration. SpringerPlus, vol. 4:197, 2015.
10. C. Tang, B. C. Tak, LongWang, H. Huang, S. Baset. On the Challenges and Solutions for Migrating Legacy Distributed Applications into Cloud. IBM Research Report, 2014.
11. <http://www.netapp.com/us/solutions/cloud/private-storage-cloud>
12. K. Sabiri, F. Benabbou. Methods Migration from On-premise to Cloud. IOSR Journal of Computer Engineering. vol .17, p. 58, 2015.
13. V. Andrikopoulos, T. Binz, F. Leymann, S. Strauch. How to adapt applications for the cloud environment . Computing, vol. 95, no. 6, p. 1, 2013.
14. <http://www.techvalidate.com/portals/hp-software-application-performance-testing-shunra>
15. J. Varia. Migrating your Existing Applications to the AWS Cloud. Amazon, 2010.
16. S. Morad. Amazon Virtual Private Cloud Connectivity Options. Amazon, 2014.
17. <http://vcloud.vmware.com/service-offering/direct-connect>

18. D. Loi. Aligning Applications and Connectivity to Enable Fast And Safe Cloud Computing. Orange Business Services, 2015.
19. <http://www.orange-business.com/en/products/business-vpn-galerie>
20. <http://www.novell.com/products/migrate>
21. <http://www.visionsolutions.com/products/dt-move.aspx>
22. <http://www.zconverter.com/index.php/solution/CloudMigration>
23. <http://www.vmware.com/products/converter>
24. <http://www.internap.com/bare-metal>
25. <http://www.vmware.com/products/vsphere/features/replication>
26. <http://www.emc.com/storage/vplex>
27. A. Palekar and P. Danahy. VPLEX: Continuous Operations for VMware Environments. EMC, 2013.
28. B. Johnson. Windows Azure Data Storage. Wrox Press, Indianapolis, 2014.
29. Oracle Database Backup To Cloud: Amazon Simple Storage Service (S3). Oracle Technical White Paper. Oracle, 2015.
30. Oracle Database Backup Service. Oracle Technical White Paper. Oracle, 2015.
31. Amazon.com leverages the AWS Cloud for Database Backups. Amazon, 2012.
32. C. Zechmeister, A. Tomic, R. Ravirala, and J. Nunn. Enterprise Backup and Recovery On-Premises to AWS. Amazon, 2014.
33. Introducing Cloud Backup for MS SQL Server. The Cloudberry Lab Whitepaper. Cloudberry, 2015.
34. H. Roggero. Windows Azure SQL Data Sync. Idera White Paper. 2015.
35. K. Lu. Azure SQL Database Migration Cookbook. Microsoft, 2015.
36. C. Rabeler. Migrate SQL Server database to SQL Database using transactional replication. <https://azure.microsoft.com/en-us/documentation/articles/sql-database-cloud-migrate-compatible-using-transactional-replication>.
37. C. Rabeler. Migrate a SQL Server database to SQL Server in an Azure VM. <https://azure.microsoft.com/en-us/documentation/articles/virtual-machines-migrate-onpremises-database>.
38. A.S. Sait. Strategies for Migrating Oracle Database to AWS. Amazon, 2014.
39. <http://www.emc.com/en-us/cloud/hybrid-cloud-computing>
40. Virtual Machine (VM) Interoperability in a Hybrid Cloud Environment Rev. 1.2. Open Data Center Alliance, 2013.
41. <http://www.equinix.com/services/interconnection-connectivity/cloud-exchange>

42. S.Tummalapalli, R. Kanth .P.Yuvaraj, and S.Velagapudi. TOSCA Enabling Cloud Portability.. International Journal of Advanced Research in Computer Engineering & Technology. vol. 2, p.974, 2013.
43. <http://getcloudify.org>
44. D. Linthicum. Understanding the Challenges of Containers.
<http://www.cloudtp.com/2015/12/04/understanding-challenges-of-containers>
45. <http://www.emc.com/storage/cloudarray>
46. B. Panchanathan. What, Why and How of Cloud Storage Gateway. EMC Proven Professional Knowledge Sharing, 2013.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.