Redbooks
ibm.com/redbooks

Moving Microsoft Workloads to IBM Cloud

Matthew Price

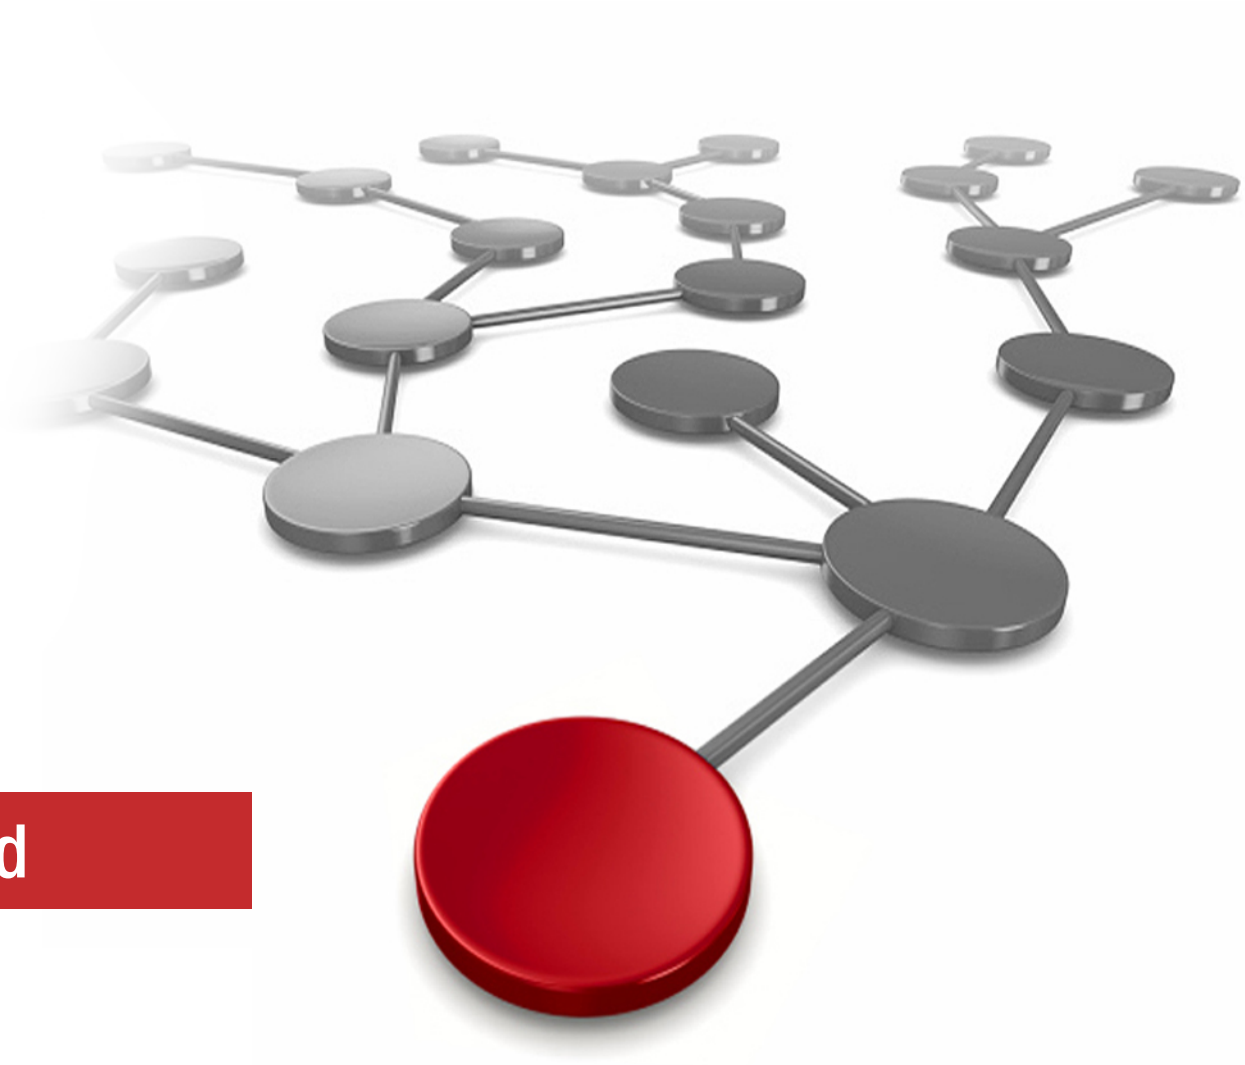

Cloud

IBM

Redpaper

Quick start guide

This IBM® Redpaper™ offers a quick, yet comprehensive, guide to moving Microsoft applications to IBM Cloud. It provides questions and considerations that help to define the scope of the cloud project, including identifying the cost structure and performing workload assessment. The paper describes the solution architecture, specifically putting together the available options and functionality to build a cloud-ready solution. The paper explores the migration options and process, and discusses how to operate the application after it is deployed.

Planning

The rise of cloud computing in enterprise IT has become a pervasive paradigm for technical practitioners and business leaders. The cloud has become a means for minimizing costs, for accelerating speed of delivery with new systems and solutions, and for fueling innovation and achieving digital transformation to compete in a new business climate. But for most businesses, a glaring reality exists of hundreds or thousands of legacy systems and applications that must be considered in the journey to cloud.

To begin this process, an important aspect to understand is that businesses can leverage the power of the cloud in a number of ways. Each includes its own distinct set of pros and cons to consider. As such, recognizing that no single solution that is right for all situations is also important. To implement a successful application migration strategy, you must identify these differences early in the process.

Accordingly, one of the first questions that must be answered is: *What do you hope to gain by migrating an application to the cloud?* The answer to this seemingly simple question will influence multiple decisions that need to be made throughout the migration process.

If the only objective is to label your application *cloud-ready*, the easiest path might not require the migration of the entire application. Instead, you could stand up a front end to your application designed by, on, and for the cloud to interface with your application's back end wherever it currently resides.

Another consideration for why you should or should not migrate an application to the cloud is the need to increase the agility and flexibility of your entire organization's application portfolio. An application that is in high demand today might be rendered completely obsolete in the near future. And businesses of all sizes are now expected to respond to these market disruptions in near real time. With traditional on-premises procurement practices, the go-to solution was simply to prepare for the worst case scenario and hope for the best.

Not only is this strategy extremely cost-ineffective and slow, but it also does not guarantee business continuity in the case of infrastructure failure. With the on-demand scalability of cloud infrastructure comes the ability to quickly change direction in response to changing opportunities. Companies can focus on their core business goals and not face any slowing down caused by inflexible self-maintained infrastructure.

Assuming the application being migrated is designed to leverage the distributed nature of the cloud, business continuity and disaster recovery concerns can be addressed to the point where they are merely line items in architectural design, and data loss becomes nearly impossible. Compared to the past, when environments would strive for Six Sigma reliability, cloud practices allow for scalable solutions where, for extremely demanding applications, reliability of 16 nines and availability of 8 nines can be provided. Or customers can configure for more typical levels of reliability and availability and achieve economic savings as a result.

In the world of IT, a constant driving force is to reduce cost and increase speed. This desire to find a lower cost alternative to current environments is often what first leads to cloud consideration. An important note is that cloud might or might not be the cheapest hosting environment for your application, depending on several factors.

Here are several budgetary questions to consider:

- How are IT expenses financed within your organization?

  The reason this question is important is because you want to know whether expenses will immediately start to limit possible infrastructure as a service (IaaS) options. If your company operates on capital expenditure (CAPEX) budget, often seen in utility companies, such as power and electric businesses, then your options for creating an on-premises cloud environment to house your application might be limited.

  If your expenses are financed through an ongoing operating expense (OPEX) model, then you might have the option of leveraging an existing on-premises cloud environment, or perhaps using an off-premises public cloud infrastructure might be more affordable.

- How variable will the costs of your application be?

  If your application is business-critical and expected to run 24/7/365 at a fixed resource cost, projecting the financial costs can be easily calculated. This might be especially useful for capacity management of an on-premises cloud solution. If cost is your only consideration, and your business still plans to maintain its own data center, the cheapest hosting option will still likely be buying a budget server specified for that application and install it in your existing data center. This of course forfeits all of the other macro-level and application-level advantages that cloud hosting solutions can provide.

  Perhaps the application is based on quarterly events, such as the heavy use of online retail apps seen typically during the fourth quarter. An off-premises bare metal solution might be the best location for this app in order to leverage the stability of dedicated hardware, with minimal overhead costs due to the planned expansion and retraction of infrastructure resources around the application's peak usage. For details about an off-premises bare metal solution, go to the IBM Cloud bare metal servers web page.

  Otherwise, if the application is dependent on external triggers, which cannot be regularly predicted, you might consider placing in a virtualized cloud environment where it has near unlimited growth space based on real-time need.

- Is your organization already using an existing cloud environment?

  Why is this a financial question? Capacity. If your company already has cloud infrastructure (whether on-premises, or bare metal off-premises) being used by other applications, this application might be able to reside on that base infrastructure, and designed to *burst* onto its own separate infrastructure if necessary. This approach can reduce the total cost of ownership (TCO) of your applications to nearly zero if the workload is inverse to that of other applications on that environment (that is, follow-the-sun apps being paired with follow-the-moon apps).

Regardless of your organization's primary objectives for migrating your Microsoft application workloads to cloud, keep in mind, throughout the entire migration process, these important aspects:

- Some workloads are a good fit for cloud in their current state; others might require substantial work effort to operate in a cloud environment.
- Some workloads can be placed into a large number of cloud configurations based on desired benefits; other workloads have limited options for how they can be successfully deployed into the cloud to maintain full functionality.

Assessing workload fit

As mentioned in the previous section, some workloads are a good fit for cloud in their current form, and others that are not. IBM refers to this as *cloud affinity*. If the affinity is high, chances are good that migration to cloud will be successful while meeting stated business requirements and achieving desired business value. In other cases, migration will require additional up front work in order to increase the likelihood of a successful and beneficial migration to the cloud.

Examples of workloads with high cloud affinity include those with these characteristics:

- Good business/IT motivation with less pain to implement or migrate:
  - Workload elasticity or speedy deployment adds value and worth to the organization
  - Workload whose applications do not require extensive remediation
- Standard and virtualized infrastructure that already exist:
  - Workload runs on a virtualized environment that is not heavily customized
- Integration requirements that are not too significant:
  - Workload runs as a relatively stand-alone application without significant integration requirement
- Non-functional requirements that are not too stringent:
  - Workload does not have high availability or performance requirements
  - Workload does not support mission-critical transactions

Figure 1 shows workloads classified by cloud readiness.

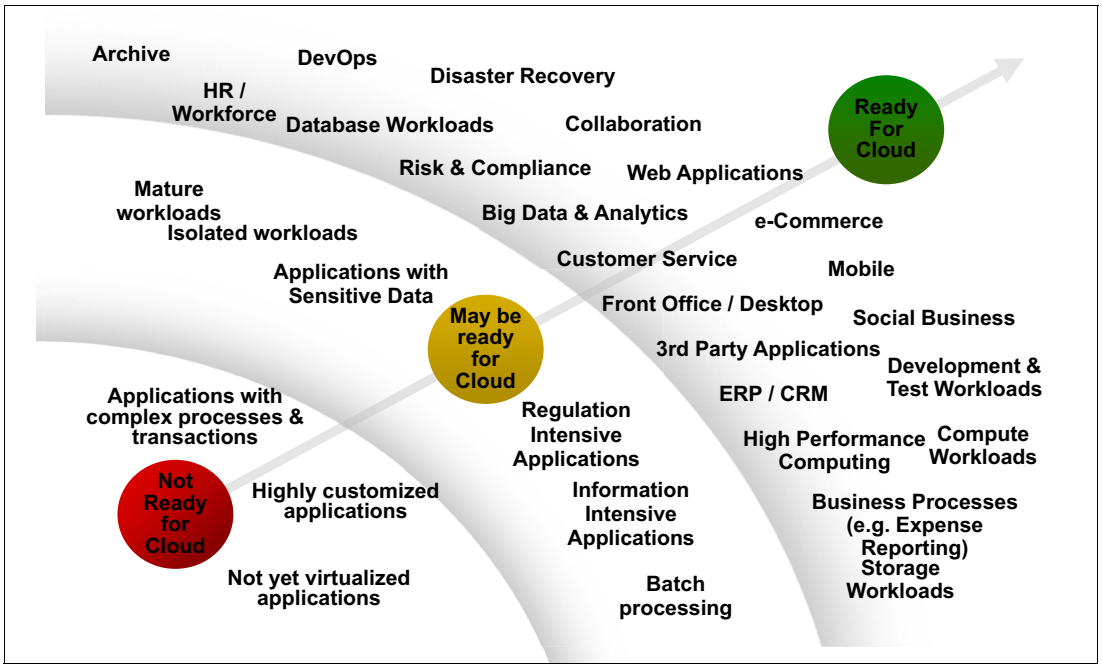

Figure 1 Workloads classified by cloud readiness

Key questions to consider in determining cloud affinity of your target workload:

- How self-contained is the workload?
- What are the scalability requirements for this workload?
- How standardized can the underlying IT infrastructure be?

- How standardized is the workload itself?
- How differentiated is the workload (is it a source of competitive advantage)?
- Is the workload available as an application or business process on the cloud?
- Does the organization have strong motivation to move the workload to run in multi-tenant environment to improve operational efficiency and for long-term cost reduction?
- What is the size of the migration or transformation effort?
- What are the data transfer requirements for the workload?
- To what degree does the workload require adherence to performance and support non-functional requirements (NFRs)?
- How large is the benefit of rapid application deployment for this workload?
- What privacy and security requirements does this workload require to meet compliance or regulatory requirements? (Are you on a data sovereignty point?)

Based on these questions, one of three labels is typically used as shorthand to identify the cloud affinity of any given application: legacy, cloud-enabled, or cloud-centric.

- Legacy applications

  Legacy applications are those typically less “ready” and not easily migrated to operate in the cloud successfully. Such application architectures are not typically designed to leverage dynamic resourcing models of the cloud, or to operate without data persistence, and can thus create dangers to critical enterprise applications.

  If an organization is moving all of its legacy workloads to the cloud, such applications would be hosted on top of cloud infrastructure, but would not be designed to inherently benefit from the dynamic scalability or agility of comparable workloads with higher cloud affinity. The alternative might be to port, refactor and redesign or rearchitect the application to more effectively function as a cloud-enabled or cloud-centric solution.

- Cloud-enabled applications

  Cloud-enabled applications are those workloads that were originally developed in a traditional data center but were repurposed for use in the cloud. The level of modification required for these applications varies based on the functionality of each app, but might be the quickest approach to migrating an existing workload to the cloud while still reaping some of the major benefits of the cloud model. A simple example of this type of application might include deployment of a cloud-native front end application with secure connections to the legacy application running behind a corporate firewall.

- Cloud-centric applications

  Cloud-centric applications are those with the highest cloud affinity. Most often these are applications designed from the ground up on and for cloud-based deployment. This includes leveraging new technologies such as, social media integration, Internet of Things (IoT) sensors, and mobile devices which may not have existed when comparable legacy applications were originally deployed. Often, these workloads could not exist in legacy non-distributed environments. As such, these applications are rarely migrated to the cloud, but might be migrated between separate clouds or between separate cloud service providers (CSPs).

Architect: Choosing the right cloud environment

After you determine the goals of your migration and the affinity of your workload type, you can begin choosing the right cloud environment and designing an architecture that fully leverages the benefits that the cloud has to offer.

Lift and shift of legacy level applications

The lowest barrier of entry for migrating an application is standing up a virtual server in a public cloud or a dedicated bare metal hybrid cloud with an existing cloud service provider, to redeploy or *lift and shift* your application to. For this type of migration, you are simply looking to re-create an application's existing environment within the cloud of your selected provider. This includes the following considerations:

- Matching system resources:
  - Memory
  - CPU
  - Storage
- Ensuring that the cloud service provider has the required certification levels.
- Deployment to a similar geography to minimize changes to latency.
- Ensure matching operating system/service packs.
- If your legacy application has dependence on services or environments that are no longer in support (such as Microsoft Server 2003), you might be forced to use a bare metal environment to install these specific dependencies.

Migration to cloud-enabled application design

To take advantage of some of the benefits of cloud, and to allow for easier control and management of applications, most applications that are migrated to the cloud undergo at least some of the following alterations:

- Virtualization and containerization to allow for automated deployment or easy portability.
- Networking reconfiguration to make best use of cloud resources, such as the IBM Cloud triple network architecture and virtual networking designs that are common to cloud-based architectures.
- Decoupling of application and data.

  Data Storage must consider inputs/outputs per second (IOPs) requirements.
- Geographic-independent application (Note that this does not apply to the data layer, which might have regulatory restrictions.)

Migration to cloud-centric application design

To get the most benefit from your cloud migration, incorporate as many best practices as you can from current cloud-ready application architecture designs including these:

- Microservices architecture
- Server agnostic design
- Integration of access and permissions with active directory

- Stateless application design
  - Ideally making them REST-compliant for interoperability.
  - Expectation of compute resources vanishing mid-transaction.
- Application configuration saved as a container template or server image
- Application calls back to Cloud Orchestrator for autonomous deployment of additional applications to dynamically scale
- Using existing disaster recovery processes at the IaaS layer
- Using high availability processes at the IaaS layer
- Auto-scaling rules are in place including these
  - Horizontal scaling
  - Vertical scaling
  - Floor/ceiling limits
- Integration with cloud brokerage services for cost controls
- Utilization of new cloud services such as blockchain networks

As applications increase their cloud affinity, they benefit from more and more aspects of the cloud. A cloud-centric containerized application complete with stateless design and infrastructure fault tolerance inherently has the same scalability potential that a legacy application, reconfigured to be cloud-enabled has.

Migrate: Setting up an IBM Cloud environment

This section covers the eight key steps to set up an IBM Cloud environment.

Step 1: Complete the prerequisites

Before you prepare to move your workload to IBM Cloud, perform these activities:

- Determine the types of IBM Cloud cloud configurations bare metal servers, virtual servers, or a combination.
- Determine the location of the IBM Cloud data center that fits your needs.
- Create a design plan for IBM Cloud (compute, storage and network options, security, monitoring, and management):
  - You can change your environment later. Although adding storage, CPU, and memory to instances is easy, changing the network itself is tricky. Plan your network carefully. IBM Cloud has a Triple Network Architecture (public, private, and management) network. An overview is available at The IBM Cloud network web page.
  - More information is available at the IBM Cloud Storage solutions web page.
- If you select an IBM Cloud virtual server, determine whether you want a public or dedicated virtual server. The standard offering is a public-based virtual server, which is a multi-tenant environment, suitable for various needs. If you are looking for a single-tenant environment, consider choosing a dedicated server environment. This option is ideal for applications with more stringent resource requirements. For more information about the current virtual offerings, see the Introduction to Virtual Servers topic at the IBM Cloud KnowledgeLayer® web page.

What is the difference between public and private nodes? It has to do with the concept of tenancy:

- A private node is run on a host machine that is dedicated to one customer of the cloud provider (single tenant).
- A public node is run on a host machine that is shared among multiple customers (multi-tenant).

- Determine the networking you want for your IBM Cloud environment: public, private, both.
- Ensure your virtual server is secure.
- Create an image by using the *image import* functionality that is currently limited to producing a standard image for virtual hard disks (VHDs) and ISO templates when importing an ISO.

Step 2: Log in to the Customer Portal

If you created your IBM Cloud account after July 6, 2016 follow the **IBMid Account Login** link at the bottom of the Customer Portal login page to use your IBMid to log in. Perform the following steps:

1. Log in to the Customer Portal with your credentials (see Figure 2):
   - For users in accounts that do *not* use IBMid for authentication:
     - Log in to the Customer Portal with the credentials you received through email when your account was initially created.
   - For users in accounts that use IBMid for authentication:
     - Log in to the Customer Portal with the IBMid credentials that you set up when your user ID for the account was initially created.
     - You are redirected to your Customer Portal Dashboard after your IBMid is successfully authenticated.

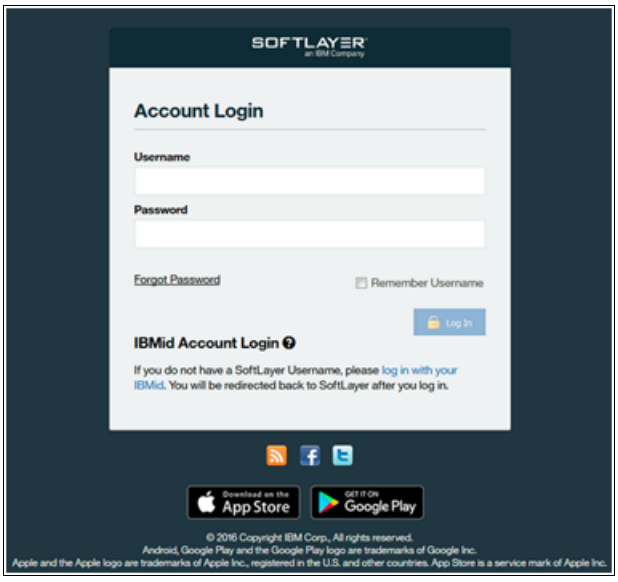

Figure 2 Account login

2. If you are a new user, you can verify your User Profile. For details, see the IBM Cloud KnowledgeLayer web page.

Step 3: Set up your IBM Cloud virtual server

Configure your cloud with virtual servers that meet your workload needs. Select your CPUs, storage, and memory (RAM). Perform the following steps:

1. You can order servers from the Order window on the Customer Portal home page. Click **the Devices link**.

   The Order Devices window opens. It shows the types of servers available.

2. Click the server type you want, and select hardware, software, and services options.

   For virtual server options, you choose which data center to provision in, and some of the basic server specifications such as number of virtual CPUs (vCPUs), RAM, and whether the storage will be through SAN or local storage. With bare metal offerings, additional customization options are available that must be considered, for example number and type of internal drives, and hard drives, such as Serial Advanced Technology Attachment (SATA), serial-attached SCSI (SAS), or solid-state drives (SSD). All provisioning options will also provide you an opportunity to provision the operating system of your choice.

   For details about selecting storage, see the IBM Cloud Storage solution web page.

   Click the **Continue Your Order** link.

3. You can also specify any of the following items:

   – Public and private VLAN (if your account has more than one provisioned).
   – Provisioning scripts that run automatically after the server is provisioned.
   – Secure Shell (SSH) keys, allowing for a more secure login.
   – User metadata, which is server-specific data that can be passed to provisioning scripts.
   – The host and domain name for your server.

4. Read the IBM Cloud Service Agreement. To continue, accept the agreement.

5. In Devices tab, click **Device List**, and watch for your server to become active.

Note: When the first server is provisioned in a pod or data center, the account will receive a new public and private VLAN. If you select a private VLAN-only deployment, your server will be deployed without a public VLAN assignment.

With your server now available you can do the following tasks:

- Find your server in the Device List.

  All of your devices are listed in the Device List, where you can manage devices, upgrade devices, or generate bandwidth usage charts.

- Choose to manage your device.

  You can interact with servers in the Customer Portal in both the Snapshot view (a summary of your device) and on the Device Details page (a fully detailed list):

  – To view and interact with your server in the Snapshot View, click the arrow next to the Device Name to expand the view.

  – To view and interact with your server from the Device Details page, click the device name of the server.

- Record the IP address and credentials for the server in a secure place that allows you to quickly access details without logging in to the Customer Portal.

- Update credentials for operating system and software.

  All software loaded onto your device during the provisioning process was assigned temporary credentials by our systems. These credentials are viewed and managed on the Passwords page of each device in the Customer Portal. Use these temporary credentials to access your software for the first time. As a preferred practice, change the password to your software after accessing it for the first time. Use a strong password to contain a combination of letters, numbers, and symbols.

  Optionally, password updates might be stored on the Passwords page for each device. However, understand that when storing passwords within the portal, any person with access to the account and appropriate permissions can view passwords stored on the Passwords page.

  For detailed instructions about viewing and managing your software credentials, see the Add, Delete and Update Software Users and Passwords topic in IBM KnowledgeLayer.
- Set up monitoring to evaluate server availability and potentially when to scale.
- Secure your system by implementing a hardware firewall.
- Schedule backups.

Step 3a: Select your IBM Cloud storage option

IBM Cloud virtual servers can be deployed with primary storage based on local disk or SAN, and with portable storage volumes as the secondary storage. If your application requires faster disk I/O, resiliency, or long-term flexibility, you can match your virtual server's storage to its application.

Portable storage volumes are secondary storage solutions that are exclusively available on IBM Cloud virtual servers. They can be connected to one virtual server at a time and are an ideal solution when looking to transfer data between virtual servers in any data center on the IBM Cloud network. Portable storage volumes are useful for database applications that access raw, unformatted block-level storage, and for moving large data sets between virtual servers. You can also use iSCSI (block) storage, which is common for Microsoft SQL databases and Windows clusters in general.

For more details, see the IBM Cloud Storage solutions web page.

Step 3b: Understand your IBM Cloud networking option

IBM Cloud uses a triple network architecture. For the rack, you use the IBM Cloud public and private networks, and also use the services that are built into the IBM Cloud management network. Each rack within an IBM Cloud data center has at least 10 gigabits per second (Gbps) of redundant (20 Gbps total usage) bandwidth for both the private network and public network.

This approach allows for a throughput of 10 Gbps throughout the IBM Cloud network down to each rack. The management network, although it is a separate network, uses the private network bandwidth for access to the systems.

For more information, and to chat with a network expert, see the following web pages:

- The IBM Cloud network page describes IBM Cloud, a network of networks.
- Networking services in IBM Cloud page describes networking solutions.

Step 4: Create Windows operating system images to be deployed in IBM Cloud

To create a `cloudbase-init` image template in IBM Cloud IaaS, complete the following steps:

1. Deploy single virtual server from a public image:
   a. Select a public image, for example Windows Server 2012 Standard Edition (64 bit).
   b. Type: Small.
   c. No key.
   d. No user/password.
   e. No volume attaches.
2. Log on to the virtual machine:
   a. The virtual machine password is generated at first. So once the virtual machine is reported as `ACTIVE`, open the IBM Cloud portal and navigate to your devices list.
   b. Expand the virtual machine that you provisioned and click the **show password** box to reveal the administrator password.
   c. Use this password to log on to the virtual machine by using RDP with the IP address provided.
   d. Depending on the load on the IBM Cloud data center, you are using, it might take up to 20 minutes before the login information becomes available.
3. Install `cloudbase-init` on the virtual machine:
   a. Download the installer from this web address:

      `https://www.cloudbase.it/downloads/CloudbaseInitSetup_Beta_x64.msi`
   b. Run the installer.
   c. Enter the correct administrator user name for your version of Windows. For example, `Administrator` for the English version. Be sure that you select the **use metadata password** option.
   d. Click **Next** and wait until the installation completes. Do *not* select to run **`sysprep`** or to shut down the virtual machine.
   a. Click **Finish** to close the installer.
4. Copy the IBM Cloud metadata service to the virtual machine:
   a. The built-in version of `cloudbase-init` does not support loading metadata from IBM Cloud IaaS. Therefore, the `cloudbase-init` installation on your virtual machine must be extended with a small file that implements an IBM Cloud metadata service.

      Download the file from the appropriate `ico _server`:

      `http://<ico_server>:9797/downloads/scripts/SoftLayer/windows/cloudbase-init/`
   b. Copy the `slservice.py` file to the services folder of your `cloudbase-init` installation. This is the default folder:

      ```
      C:\Program Files (x86)\Cloudbase
      Solutions\Cloudbase-Init\Python27\Lib\site-packages\cloudbaseinit\metadata\services
      ```
   c. Adjust the configuration settings of `cloudbase-init`:
      i. Open the `cloudbase-init.conf` file in an editor. This is the default folder:

         ```
         C:\Program Files (x86)\Cloudbase
         Solutions\Cloudbase-Init\conf\cloudbase-init.conf
         ```
      ii. Make sure that it contains these lines:

         ```
         metadata_services=cloudbaseinit.metadata.services.slservice.SLServiceplugins=clou
         dbaseinit.plugins.windows.setuserpassword.SetUserPasswordPlugin
         ```

d. Create a private image from the virtual machine:

   i. Leave the virtual machine (do not shut it down, do not run **`sysprep`** or anything).

   ii. Return to the IBM Cloud portal and click on the virtual machine to open its details.

   iii. In the action menu, select the **create image template** action and provide an image name.

   iv. You can now use this private image for provisioning with IBM Cloud Orchestrator.

e. (Optional) Delete the original virtual machine; consider this important information:

   - After the new image is created, you can safely delete the original virtual machine by using the IBM Cloud Orchestrator UI.
   - Be aware that creating the private image template might take 20 minutes depending on the size of the virtual machine.
   - Until the image creation transaction is completed, the original virtual machine cannot be deleted.

Notes:

- The password that you enter during provisioning is visible to a anyone who can log on to the provisioned virtual machine. So the best approach is to change the password as soon as possible after the first login.
- Setting the password during provisioning works only if the chosen password complies with the password policy of the Windows operating system on the image. If the password you chose at provisioning time does not comply with the password policy, the password is not set. Then, you are able to access the virtual machine by using the password that is originally generated by IBM Cloud that you can reveal by using the IBM Cloud portal.

These instructions are available in the Creating Windows operating systems images topic at IBM Knowledge Center. That topic refers to SoftLayer, which is now named IBM Cloud.

Step 5: Access server instances

When accessing server instances consider these various methods:

- To access a server running Windows from a Windows computer, you can use Windows Remote Desktop program.
- To access the command prompt on the server from your workstation, you can also use any of the Secure Shell tools available, such as OpenSSH in Linux, or System X or PuTTY on Windows.
- If you generate SSH keys and transfer them to your server when ordering, you can use them to log in securely.
- You can also set up a VNC server on your server, and install, run, and access Linux GUI environments from clients on Windows.
- The initial root/administrator password for your server is available in the Passwords page in Device Details. This password is not automatically kept in sync with password updates on the device.

Step 6: Create your software bundle

If you did not yet create your software bundle, learn more about creating them in the Managing software bundles topic at IBM Knowledge Center.

Step 7: Add users to your account

Adding a user to your account is easy, but remember that giving a user access to the web interface essentially gives that user access to your server room and everything within. Use permissions to limit user access. Create users based on the roles they will perform. You can apply any of the following three default permissions templates. You likely need to create customer users and apply the templates as a base for those users.

- View Only User: Has the privilege to view tickets and basic statistics.
- Basic User: Can manage almost everything concerning the services and devices but cannot cancel a device, If a basic user orders a service that incurs a charge, someone who can approve charges to the account must approve it before the service is processed.
- Super User: Has the same rights as the master account. Limit the number of these users if you have them at all.

In IBM Cloud IaaS, you need to set the permissions for the user and decide which instances, if any, that user is able to access. The permissions that you set can be even more fine-grained if you use the quick filters. By using quick filters, you can grant permission automatically to the user for any future instances of the same type.

Step 8: Select your backup and recovery option

IBM Cloud offers various options for backup and recovery.

If you want to use EVault Backup

EVault Backup is an automated agent-based backup system that is managed through a centralized web administration console called WebCentralControl (WebCC). It allows you to back up data between servers in one or more data centers on the IBM Cloud Network. To use EVault, complete the following steps:

1. Order an EVault backup volume for the server you want to back up.
2. Install the EVault Client (also known as Agent) on your system. Instructions for agent installation are available at the IBM Cloud backup services web page.
3. Create a backup job using EVault WebCC.

If you want to use R1Soft

Another backup solution available from IBM Cloud is R1Soft Server Backup. R1Soft Server Backup provides high-performance disk-to-disk server backup, featuring a central management and data repository. It protects data at block level, and unique disk blocks on the server are stored only once across all recovery points, increasing storage efficiency.

R1Soft is ordered as an add-on for a bare metal server. IBM Cloud provides the license based on the amount of backup agents that the customer needs. For more information about R1Soft, see the R1Soft Server Backup Manager web page.

If you want another option for backup

In IBM Cloud IaaS, you may install any middleware on your virtual servers and bare metal servers. IBM Cloud clients can implement their own backup solution on their bare metal servers by using software such as Veeam, IBM Spectrum Protect™, Veritas, NetBackup, EMC Networker.

If you want to use snapshot and replication for Endurance storage

Snapshot technology can be used as a first line of defense in your backup and recovery plan. You can easily combine Endurance snapshots with any of external backup solutions. As an extension to the Snapshot functionality, a replication functionality is available for Endurance storage.

Replication uses one of your snapshot schedules to automatically copy snapshots to a destination volume in a remote IBM Cloud data center. During a disaster, you can fail over to your replica site and mount the replicated volume to any server in the failover IBM Cloud data center.

For more information about snapshots and replication, see the following topics at the KnowledgeLayer web pages (to access these topics, log in with your IBM Cloud credentials):

- Take a snapshot
- Set up replication

Conclusion

With your new infrastructure installed, you can now move your application to its new hosting environment. By moving your Microsoft application to the cloud, you should now be able to use at least some of the benefits that businesses across all industries have been taking advantage of to set themselves apart in the market place. Whether your application is a small mobile app designed for direct customer communication, or a back-office app used for routine batch processing, by staging it in a cloud environment, you have now increased its cloud affinity and can begin experiencing, firsthand, the value that cloud can bring to your business.

For additional help, chat with IBM at the MSSQL and Windows servers web page:

Authors

This paper was produced by a specialist working with the International Technical Support Organization, Austin Center.

Matthew Price is an IBM Global Cloud Advisor in the US. He has 8 years of experience in cloud computing. He holds a degree in Computer Science from Northern Illinois University. His areas of expertise include Cloud Architecture and IT Project Management. He has written blogs on cloud computing related topics.

Thanks to the following people for their contributions to this project:

Vasfi Gucer, LindaMay Patterson, Diane Sherman
International Technical Support Organization, Austin Center

Janet Whaley Sifers, John F Scott
IBM

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks
- Follow us on Twitter:

  http://twitter.com/ibmredbooks
- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806
- Explore new IBM Redbooks® publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm
- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo)  ®	IBM Spectrum™	Redbooks®
IBM®	IBM Spectrum Protect™	Redpaper™

The following terms are trademarks of other companies:

KnowledgeLayer, and The Planet are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

REDP-5428-00

ISBN 0738456624

Printed in U.S.A.