

On-Prem to Azure

A Guide to Azure Migration

Hitachi Solutions NA Headquarters
100 Spectrum Center Dr STE 350
Irvine, CA 92618

Introduction


In the dark about moving to the cloud?

It's hard to remember a time before "the cloud." In the beginning it was an unknown and emerging technology with lots of hype. But now, with changing business environments and mobile workforces, it has gained tremendous attention and momentum.

Today, the benefits of the cloud versus on-prem are proven and many—cost savings, faster deployments, better performance, on-demand capabilities, reduced complexity, scalability, and much more. And the positive impact of the cloud is almost instantaneous, with many companies reporting operational improvements within the first few months of adopting the technology.

As companies big and small "see the light," they are heading to the cloud in droves to take advantage of these exciting outcomes. The most popular reasons for cloud migration include speed and performance improvements, greater flexibility and mobile agility, and improved customer service and support.

Cloud computing continues to mature and gain popularity. Already more than 81% of all enterprises have a multi-cloud strategy laid out or in the works. And, by the end of 2020, 67% of enterprise infrastructure are expected to be cloud-based.



Are You Ready?

What is sparking the conversation

In today's competitive and dynamic business landscape, the need for cloud services has become undeniable. Gartner's research indicates that this year, for the first time, IaaS outpaced on-premises deployments, and the growth will continue well into the future. And they say market growth will be \$331.2B by 2022.

Despite statistics like this, many businesses like yours remain on the fence about moving to the cloud. We get it. While it offers exciting benefits, it can also present confusion, complexity, and – let's face it – fear. But there's usually a specific catalyst for starting the migration discussion:

Reduced operating expenses

Reduced hardware support, increased manageability, and more efficient processes, save an average of 20–30 percent on resource configuration alone.

Support scalability

When you plan for peak usage through on-premises systems, your servers are usually running at less than 20 percent utilization. The cloud releases you from this model, enabling a scale-when-you-need-it approach.

Application modernization

Two common challenges for today's businesses are over-allocated IT resources and on-premises platforms that limit the adoption of modern services.

Software end-of-support

This can be an opportunity to migrate your end-of-support workloads to Azure to strengthen your organizational security posture and ensure compliance.

- Support for SQL Server 2008 and 2008 R2 ended on July 9, 2019
- Support for Windows Server 2008 and 2008 R2 ends on January 14, 2020

Blockers to Look for:

Fear of the Unknown

Transforming your business by leveraging the cloud requires more than just technology and platform migration. It's complex and hard to even know where to begin. You don't know what to expect and you don't know what you don't know. This can be overwhelming.

Loss of Control

Some IT leaders fear losing control. Others fear the loss of accessibility to their servers and applications. Their comfort zone is when all the equipment is within four walls

Cost

Many executives fear the cost of moving to the cloud could be higher than operating the same equipment on-premise.


Lack of Trained Personal

Most in-house IT teams are already busy with service maintenance, network infrastructure, and applications. When it comes to the cloud, they don't have the time or expertise to handle the migration and ongoing management of their cloud infrastructure.

Security and Governance

Data security is on everyone's mind these days and can be a huge roadblock for many companies thinking of moving to the cloud. Is the cloud as secure as our on-prem data center? What about regulatory compliance, data loss, data breaches, and disaster recovery? All are valid concerns.

Four Steps to Migrate:



Step 1. Discovery

Where to start with your on-prem systems

Discovery here refers to surveying your entire business enterprise to discover where every instance of on-premise systems reside and the details behind each of them. Many customers who are pursuing upgrades to their infrastructure are doing so while also undergoing a true-up period, a time where they are renewing their licensing. In those cases, there may have been a lot of server deployments, some of them flying completely under the radar.

The most efficient way to scan the environment and discover existing on-premise instances and what databases are inside those instances is with the use of [Microsoft Assessment and Planning tool kit](#), or MAPs.


After scanning your domains with the MAPs tool kit, the next aspect to discovery is accomplished through interviewing your company's key stake holders, business and application owners. Be sure to document your internal customers' objectives and how they envision their desired end state. Do they want an on-premise server consolidation? Or perhaps a cloud migration? Are there any restrictions, special considerations, etc.?

Through this process, we will discover and document any unique feature usage for the servers and their databases such as Linked servers, FCI, database Mirroring, AGs, extended stored procedures, etc.

Once all of your on-premise instances and databases are known, we work through discovery of who the Application Owners and SMEs are internally and identify what applications are associated with\connect to a given database. In cooperation with application owners, SMEs, and in-house DBAs, determine which databases are intended for upgrading\migration. Next, we determine whether the application itself requires upgrading to support the newer version of SQL Server or its migration to the cloud (Azure SQL Database, Data Warehouse, etc.).

For databases created by third-party ISV applications, this might get a little tricky. Use your support agreements and ISV sales reps to find out what upgrade or cloud restrictions there might be. This activity is often the reason a database is still running on older versions of an RDMS. If you still need the functionality of an application but your vendor goes out of business or your party ways, companies like Hitachi Solutions can help you re-write it.

It is important to identify all the applications that are associated with, or connect to, a given database. Processes as simple as SP_WHO2 can provide us with details such as login names, host names, and program names. Of course, this is not fullproof as users and applications may connect to a database in an infrequent manner. To get a clearer picture, we often resort to extended events or server-side tracing to capture connection information from those occasionally connected processes.



Step 2. Planning

The **what** and **how** of your migration

Planning is the next stage of the process. Simply put, now that we have the chess board setup and we know where all the pieces are and who the players are, we can begin to strategize on the what and how of your particular migration project and how it should proceed. The more effort we put in at this stage usually translates to less difficulty later on.

Since upgrade efforts often entail multiple servers and hundreds of databases, we like to have a system of classification and then build upgrade\migration processes around those classifications, not around the individual database.

A typical approach is to classify databases into one of three levels based on how important the data is to your business. The ranking may also reflect the enterprise-level features currently being used, or that you desire to start using the newer version of SQL Server. For instance, today your SLA's for high availability as they relate to RPO and RTO are being met, but going forward, the application owners want tighter standards. This may translate to the new platform will need to support rolling patching of the OS. In that case, databases that are happy to be part of log shipping may now need to be part of an availability group—an enterprise edition level feature for the most part.

Application owners and SMEs are best equipped to make these judgment calls in cooperation with your DBAs. These classifications might imply different kinds of servers and configurations as well as the amount of pre-upgrade evaluation and testing effort expended. Top-level systems (business or mission critical), for example, may want high availability technologies like Availability Groups with very small RTO and RPO standards.

Often, customers want these “tier 1” databases to undergo deeper testing in a lab environment to ensure no performance problems after the upgrade.

Additionally, these “tier 1” databases may demand a more complex set of migration and rollback plans that minimize service outages. Establish up front which databases must meet a stringent migration window of opportunity. It is likely these will be the same databases that will need advanced performance tuning prior to migration and require advanced HA\DR strategies with very small RPOs and RTOs. The planning for migration should include examining all the possibilities; from highspeed Express Route connectivity to the cloud to third-party replication products like DoubleTake. It could also be as simple as using the Data Migration Assistant (DMA) tool. Straight backup\restore is widely used, as is log shipping—especially when the databases are very large. The main point here is whatever defines a tier 1 versus a tier 2 database for you will also help inform you as to what are the appropriate migration processes for each tier.

Databases designated as tier 2 with a mid-level priority could have looser standards for all sorts of things, from the form a migration might take to HA standards. As such, for on-premise deployments businesses might use log shipping for migration as well as creating warm standby servers for their HA\DR requirements.

Lowest-level databases (aka “tier 3”) might require only nightly backups to satisfy any HA\DR requirements. In fact, the checklist for upgrading the lowest level of priority databases from SQL Server are usually much simpler.

For on-premise consolidation efforts where databases come from many instances, evaluate grouping compatible databases. This advice also works well with Azure's SQL Managed Instances if moving to the cloud is your intention.

- Similar workload characteristics (DSS\OLAP\Reporting versus higher volume OLTP)
- (DSS\OLAP\Reporting workloads like higher settings for max degree of parallelism and optimize for ad hoc workloads)

- System collation conflicts need to be avoided, MAP tool kit can help with discovering those for your existing instances
- Identify and group databases that share inter-dependencies
- Security requirements (TDE, CCC, C2, xp_cmdshell, Ad Hoc Distributed Queries, Ole Automation Procedures, etc.)

Understand feature differences then decide which databases will go where to get the most for your money and your best competitive advantage

- On-premise: Versions (2016, 2017, 2019)
- On-premise: Editions (Enterprise, Standard)
- Cloud: (SQL Database, SQL Data Warehouse)
- Cloud: (Basic, Standard, Premier, Elastic Pools, Managed instances, Reserved Instances, etc.)

Sizing instances for either on-premise bare metal to virtualization or to Azure's cloud offerings for IaaS/PaaS

- Infrastructure as a Service (IaaS): {VM's, Reserved Instances}
- Platform as a Service (PaaS): {SQL Database, Elastic Pools SQL Databases, SQL Managed Instances, SQL Data Warehouse}
- Depends on having existing server specs and resource usage statistics particularly for the three pillars of compute resources; memory, disk IOPS and sizes, and CPU.
 - Keep in mind that 100% CPU consumption on a low-end and older processor is not the same as if the metrics are coming from a processor with large L3 cache and plus 3GHz processors (just remember there is context with every comparison)
- Projecting what it might cost to operate

Establish performance baselines

- Use tools such as SQL Server Profiler\Extended Events, sys.dm_io_virtual_file_stats, and PerfMon to gather data indicating typical performance measurements
- Ensure that the broadest possible sets of commands are captured in traces
 - For later assessment of the performance impact to upgrading your tier 1 most important databases, this aspect is crucial. If you are not familiar with the workload characteristics for these databases, then be careful to not assume

all traces are equal. Some workloads happen only on Mondays, or the last day of the month. This awareness may lead you to capture traces for different times of the day, week, or month.

Finding upgrade blockers/issues is extremely important

- As much as possible we want to avoid hitting production because
 - Some examinations (i.e., DBCC CHECKDB) have overheads that are resource intensive and could cause performance problems if done in production
 - Code remediation work should not be done in production
- Setting up upgrade lab environment allows for discovering processing that might fail or change behavior
 - Considerations are lab servers with adequate disk space, RAM, and CPU to handle the multiple restoration of databases for examination and replay testing
 - Good place for installing and launching various tools for conducting assessments: MAPS, DMA, DEA, etc.
- Inspect the restored databases for any
 - Compatibility issues with SQL Server 2017 or cloud versions of Azure SQL Database or Data Warehouse
 - User data or meta data corruption

Determine whether applications require any of the following

- Changes to connectivity settings
- Changes to authentication mode
- Measures to prevent SQL Injection

The more thought and effort we put into the planning stage for an upgrade will always aid you in achieving the best experience possible. What is listed here is generally applicable to any enterprise-wide upgrade effort, however, every environment has its own unique requirements. Your job is to make sure those requirements have been documented and are included in your planning activities.

Step 3. Migration

Your migration checklist

Identify pre-upgrade tasks

- Identify tasks that might be accomplished before the upgrade and without downtime, for instance:
 - Determine where it is possible to pre-install or enable whichever version of .NET Framework that is applicable to your end state version of SQL Server. For instance, for SQL Server 2019, .NET Framework 4.6.2 is the standard.
 - Pre-install your end state SQL Server version and be sure to use “Slipstream”\”Product update” technique to avoid all known install and upgrade bugs. It also aids in faster setup of SQL Server and makes the process simpler and repeatable.

Estimate required downtime

- Perform test upgrade and record the time to completion, then use those timings in refining your upgrade plans
- Allow enough downtime so the upgrade process and acceptance testing can be completed successfully
- Allow time for rollback if unexpected issues arise

Develop upgrade checklists


- Decide now if you want to “prepare” the databases just before the upgrade for things like:
 - Detail the steps required for taking the systems offline for a period of time and bringing them back online
 - Detail the steps to take during the upgrade processes

Identify backup and restore operations

- Plan to back up the targeted legacy databases
- Assume nothing and verify the backups (checksum for both the database and the backups then taken)

Determine acceptance and rollback steps

- Identify how the organization will accept the upgrade, and how it will make the “go/no-go” decision:
 - Verify tests to ensure applications using the upgraded database servers will run as expected and required
 - If available, enlist the support of the quality assurance (QA) team to develop appropriate acceptance tests



- Determine exactly when and how a rollback to the legacy Server might be required.
- Test the rollback plan.

Post-upgrade tasks

- Remaining tasks might include the following depending on how those database's are used:
 - Correct the rows, used pages, reserved pages, leaf pages and data page counts for each partition in a table or index
 - Rebuild cubes, reconfigure log shipping, and reconfigure availability groups
 - Test a failover cluster
 - Verify that SQL Server Agent jobs run correctly

After upgrade monitor for

- CPU usage and response times and compare to previous baseline
- Problem caused by a change in SQL Server 2016, it seems that SQL Server got a little stricter about doing subtraction of integers from time types
- Work out which databases need to be available for performing post-upgrade checks and migrate those first
- Change app connection strings as necessary
- Watch for apps that it have a connection string embedded in their code in which you may not have been aware of earlier

Step 4. Testing

Creating a Lab Environment

Testing is integral to ensuring system health before final cutover. Many migration tools allow you to mimic the production environment in the cloud and enable you to fully test the system without affecting the on-premises or cloud production versions. Once replication is complete, simply start your application or workloads using the isolated environment option, and take some time to test your startup script or runbook for errors. When you're fully satisfied both function as expected, it's time to perform the final cutover.

For your most important databases where performance must not be compromised for even a single stored procedure or statement, using the Database Experimentation Assistant (DEA) tool is your answer.

The DEA tool is used in a lab environment that can scale up to handle tremendous workloads by enlisting many other servers that together throw trace captures against both a baseline servers and your intended upgrade version.

What is important for a performance assessment is to capture a functional-coverage workload, ensuring the broadest array of commands as possible. This includes special processing like beginning and end of week, month, quarter, and end-of-year operations. Include nights and weekends or when special workloads are running.

Get to Know Azure

Azure is the leader at the junction of infrastructure, PaaS, and software.

In the traditional data center model, your organization owns the entire stack from the physical hardware to the individual applications that your organization uses to run its business.

Many companies incorporate a virtualization layer to help manage these physical assets, which helps tremendously but still does not remove you completely from managing the physical aspect of the data center.

With Azure Virtual Machines, you are able to abstract away the hardware layer, the virtualization layer, and the management of these two layers into the virtualization of a cloud network that makes up infrastructure as a service.

And mobile services take this abstraction layer many steps forward with each workload becoming more abstract and specialized.

The farther to the right you go, the less control and work is needed to manage the underlying infrastructure.

54
Azure Region

4500+
Peering Locations

83%
of the company workload will
be stored on the cloud by 2020

81%
of enterprises have a
multi-cloud strategy



Developer Services

-  Visual Studio Team Services
-  Azure DevTest Labs*
-  VS Application Insights*
-  HockeyApp
-  Developer Tools

Management & Security

-  Azure Portal
-  Scheduler
-  Automation
-  Log Analytics
-  Key Vault
-  Security Center*

Compute

-  Virtual Machines
-  Virtual Machine Scale Sets
-  Cloud Services
-  Batch
-  RemoteApp
-  Service Fabric
-  Azure Container Service

Web & Mobile

-  Web Apps
-  Mobile Apps
-  Logic Apps*
-  API Apps
-  API Management
-  Notification Hubs
-  Mobile Engagement
-  Functions*

Data & Storage

-  SQL Database
-  DocumentDB
-  Redis Cache
-  Storage: Blobs, Tables, Queues, Files and Disks
-  StorSimple
-  Search
-  SQL Data Warehouse*
-  SQL Server Stretch Database*

Analytics

-  Data Lake Analytics*
-  Data Lake Store*
-  HDInsight
-  Machine Learning
-  Stream Analytics
-  Data Factory
-  Data Catalog
-  Power BI Embedded*

Internet of Things & Intelligence

-  Azure IoT Suite
-  Azure IoT Hub
-  Event Hubs
-  Cortana Intelligence Suite
-  Cognitive Services*

Media & CDN

-  Media Services
-  Content Delivery Network

Identity & Access Management

-  Azure Active Directory
-  B2C*
-  Domain Services*
-  Multi-Factor Authentication

Hybrid Integration

-  BizTalk Services
-  Service Bus
-  Backup
-  Site Recovery

Networking

-  Virtual Network
-  ExpressRoute
-  Traffic Manager
-  Load Balancer
-  Azure DNS*
-  VPN Gateway
-  Application Gateway

Conclusion

Azure the only consistent, comprehensive hybrid cloud

In conclusion, the process for migrating can run the gamut from very simple and straight forward to extremely sensitive and challenging. Perhaps the Pareto Principle of 80/20 is a good way of thinking about how one can approach upgrading databases and other on-premise systems in mass. In which case there are two primary challenges, how to upgrade large numbers of databases that are fairly simplistic in nature, and how to handle those that are either complex or have very high performance expectations. These challenges have often led many of our customers to keep putting off their upgrading efforts.

Hitachi Solutions has been involved with cloud and database migration projects for over two decades and across a multitude of industries. We can help your company make the transition a successful and painless affair.

