

Migrating from On Premises to the Cloud



By David Loshin

Sponsored by:



OCTOBER 2019

TDWI CHECKLIST REPORT

Migrating from On Premises to the Cloud

By David Loshin



555 S. Renton Village Place, Ste. 700
Renton, WA 98057-3295

T 425.277.9126
F 425.687.2842
E info@tdwi.org

tdwi.org

TABLE OF CONTENTS

2	FOREWORD
3	NUMBER ONE Articulate the advantages of migrating to the cloud
4	NUMBER TWO Understand the migration scenarios
5	NUMBER THREE Consider and mitigate migration challenges
6	NUMBER FOUR Assemble a migration strategy
7	NUMBER FIVE Consider migration approaches
8	NUMBER SIX Develop the migration plan
9	AFTERWORD
10	ABOUT OUR SPONSOR
10	ABOUT TDWI CHECKLIST REPORTS
10	ABOUT TDWI RESEARCH
10	ABOUT THE AUTHOR

© 2019 by TDWI, a division of 1105 Media, Inc. All rights reserved. Reproductions in whole or in part are prohibited except by written permission. Email requests or feedback to info@tdwi.org.

Product and company names mentioned herein may be trademarks and/or registered trademarks of their respective companies. Inclusion of a vendor, product, or service in TDWI research does not constitute an endorsement by TDWI or its management. Sponsorship of a publication should not be construed as an endorsement of the sponsor organization or validation of its claims.

FOREWORD

Cloud computing has matured, and individuals at many organizations recognize the benefits of migrating their existing on-premises reporting and analytics data warehouse applications to the cloud. Deploying a system on a cloud platform seems straightforward, as if it were as simple as replicating your on-premises system configuration on a cloud-based computing platform.

Although it is not quite that simple, cloud migration is also not as complicated as some people make it out to be. Part of the perceived complexity is perspective—many people still have concerns about deploying applications to the cloud.

Some people are concerned about security and protecting data from unauthorized access. Others are unsettled by the potential for uncontrolled data sprawl and the inability to properly govern and manage data assets. The potential challenges of managing and interoperating across a hybrid environment cause some people distress. Finally, individuals who don't completely understand cloud services pricing models worry about unbounded costs.

Although these lingering concerns still trigger apprehension, in the past few years cloud providers have not only addressed these issues, they have continued to develop services and offerings that eliminate these perceived barriers.

Cloud security mechanisms offer strong data protection, and there has been a blossoming of tools and products for data cataloging, asset management, and data lake management that reduce the specter of data sprawl. Pricing transparency coupled with streamlined methods for optimizing the use of cloud services can help with continuous cost management and control.

Aside from the broad array of computing and storage resources cloud hosts provide, there are native auxiliary services and third-party software platforms for reporting and analytics that can supersede the capabilities of traditional on-premises deployments. This suggests a need to evaluate the current data warehousing and analytics application landscape and assess its suitability for movement to the cloud as well as overall system modernization.

That being said, different applications are suited to different migration scenarios, requiring alternatives for prioritizing the systems to be moved and developing a reasonable plan to ensure a successful migration. This checklist provides guidance to the information professional who needs to justify a cloud migration initiative, educate the communities of stakeholders and data consumers about the benefits of migration, and devise a strategy and a plan for a migration.



1

ARTICULATE THE ADVANTAGES OF MIGRATING TO THE CLOUD

In recent years, many organizations have pivoted their information management strategies to take advantage of cloud computing, with increasing adoption of multiple cloud platforms for their reporting and analytics applications and workloads. If your organization has not yet made that choice, you may have an opportunity to assess the benefits of moving away from an “on-premises only” approach and advocate for a cloud migration initiative.

Aside from commonly recognized benefits (such as lowering overall IT expenses by reducing the total operational costs of running an analytics data center, making ongoing system maintenance and upgrades transparent to system managers, and reducing redundancy to ensure high availability), cloud environments provide advantages specifically relevant to reporting and analytics applications, such as:

- **SCALABLE/ELASTIC PERFORMANCE:** The capital acquisition process for on-premises systems insists that you size the configuration when the system is purchased, requiring additional effort and resources when you need to upgrade. System configurations deployed on a cloud platform can be set to automatically scale according to the computational demand and thus continuously meet performance needs.
- **FLEXIBILITY:** Cloud hosts provide a variety of options for sizing their computing and storage resources. Providers also offer a broad array of services supporting machine learning and artificial intelligence that can easily be integrated into the reporting and analytics application architecture.

- **SECURITY:** Cloud hosts have been diligent in securing their environments against security breaches and provide a range of capabilities to further strengthen this protection. In turn, hosted platforms (such as data warehousing providers) can layer additional controls for accessing data. Hosted platforms are also improving their data protection capabilities to support multicloud deployments.

Organizations are increasingly morphing to a multicloud strategy as a way of maintaining even more control over cost management, leveraging the best cloud provider based on user needs, and reducing risks such as outages and vendor lock-in. Assemble a business justification that considers cross-cloud capabilities for performance, portability, and protection.



2

UNDERSTAND THE MIGRATION SCENARIOS

Organizations employ many types of analytics applications, so when migrating on-premises systems to the cloud, consider how different scenarios impact the migration strategy, such as:

- **OVERLOADED DATABASES:** Some organizations rely on their OLTP and operational database environments as the resource for reporting and analytics. Database systems with these kinds of overloaded workloads are prime targets for migration. Cloud platforms allow you to segregate reporting and analytics use cases from transactional/operational ones. Offloading analytical processing to a cloud-based data warehouse eliminates the performance drag on those OLTP systems.
- **DATA WAREHOUSE APPLIANCES:** Appliances are specialty high-performance systems designed to house data warehouses for reporting and analytics that have typically incurred high up-front acquisition costs. Although organizations have benefitted from their appliance investments, these systems require ongoing upgrades to maintain pace with performance demands. Reporting and analytics applications running on an on-premises appliance are good candidates for migration to the cloud.
- **COLLECTIONS OF DATA MARTS:** In some organizations, instead of relying on a single monolithic data warehouse, the IT teams extract data to populate individual data marts that are typically aligned with siloed business functions. Although collectively these data marts support the descriptive analytics demands for the different business functions, they provide

limited holistic or systemic visibility across those functions. The cloud provides an opportunity to harmonize these data marts. Even if transaction processing systems remain on premises, their data can be extracted and moved to a unified data warehouse built for the cloud.

- **DATA LAKES:** On-premises data lakes deployed on local Hadoop environments suffer from a number of drawbacks—they are complex to set up, they are costly to build and maintain, and without imposed governance, data sets are hard to find and use once they are dumped into the lake. Cloud-based object storage provides a framework for data asset organization and cataloging, and data virtualization techniques allow structured and semistructured data sets to be integrated into a comprehensive cloud-based data warehousing solution.

In many cases the organization will deal with more than one of these scenarios at the same time. Understanding the difference will help you plan a workable and efficient migration sequence.



3

CONSIDER AND MITIGATE MIGRATION CHALLENGES

All technology projects have an inherent risk of failure, and migrating your reporting and analytics framework to the cloud is no exception. Executing a successful migration initiative demands that the technicians be aware of the risks and challenges and have ideas about how to overcome them.

Challenges include:

- **CHOOSING THE RIGHT CLOUD PARTNERS:** Migrating from an on-premises data warehouse application to the cloud involves selecting a variety of partners. As organizations consider their options for a cloud provider, multicloud hybrid environments will continue to evolve. Also important is the choice of the data warehouse software platform.
- **FINDING INDIVIDUALS WITH THE RIGHT SKILLS:** Migration requires individuals with skills in system design, data migration and integration, and configuration of cloud-based systems. You may need to find system engineers with experience in cloud development and deployment.
- **CORRECTLY ASSESSING THE LEVEL OF EFFORT:** Organizations without experience in cloud migration may not be able to assess the skills required for the migration, the skills required for long-term management of a cloud-based application, or the level of effort required for executing the migration.
- **LOSS OF DATA DURING THE TRANSFER:** In most (if not all) cases, the migrated system not only consists of a large amount of data, it also includes sensitive data that must be retained as the data volumes are moved.

- **DATA SECURITY:** Sensitive data must be protected during the migration (to prevent inadvertent exposure) and the cloud environment must provide the proper controls to satisfy internal and external data protection policies.

A preliminary scoping exercise prior to launching a migration initiative will help you identify challenges and enable you to mitigate risks. Engaging external consultants with expertise in cloud migration can help you narrow the field for selecting cloud vendors and system integrators with the right level of experience. Select and prioritize migration tasks and devise governance protocols incorporating data validation and security measures to ensure the fidelity and protection of transferred data.



4

ASSEMBLE A MIGRATION STRATEGY

A successful migration must ensure alignment across the collection of stakeholders and participants, and therefore a strategy must encompass people, process, and technology. More to the point, a migration strategy guides the organization's approaches to requirements analysis, project scoping, skills assessment, and overall planning for launching a series of system migrations.

This leverages a sequence of steps:

- **ESTABLISH ORGANIZATIONAL READINESS:** Corporate stakeholders must be aligned with the objectives of a move to a hybrid environment that may span multiple cloud platforms. In addition to preparing a business justification articulating the benefits of migration, assess the level of readiness to launch a cloud migration initiative by sharing the advantages of the cloud and demonstrating to key stakeholders that their success factors will be addressed as part of the initiative.
- **SOLICIT AND DOCUMENT REQUIREMENTS:** Ultimate success means that the needs of the end users and data consumers (such as data scientists, data analysts, and users of executive dashboards) will be met.

Engage the user communities to understand their current expectations, current roadblocks, and existing challenges. Synthesize a set of requirements that address current as well as future needs for a reporting, business intelligence, and analytics environment.

- **ESTABLISH LEVELS OF SERVICE:** Specify expectations for observing defined service levels for accessibility, usability, performance, security, and uptime/availability.
- **DEVISE THE TECHNOLOGY ROAD MAP:** Synthesizing the requirements and the anticipated levels of service provides a framework for a technology road map that identifies candidate cloud providers and hosted applications, guidelines for choosing one cloud host over another for particular services, and a series of steps for data preparation, migration, and system instantiation.

These steps allow you to evaluate the analytics application landscape and identify candidate applications for migration. Characterize the applications in terms of their on-premises operational context and profile the associated workloads.

Use that acquired information to map the applications' execution and storage performance requirements to corresponding cloud resources and consider the effort to migrate each application. Prioritize based on identifying those applications that are easiest or fastest to migrate that provide the greatest economic benefit. At that point you can consider the different migration approaches.



5

CONSIDER MIGRATION APPROACHES

As organizations become more mature in their use of multicloud hybrid information environments, their approaches to data warehouse migration will also become more refined. The spectrum of migration approaches can be quite broad, but here are three approaches to migrating reporting and analytics systems to the cloud:

- **LIFT AND SHIFT:** In this approach, which is often the obvious default first attempt at migration, the cloud environment is a replica of the on-premises environment. The same underlying database management system is implemented on the cloud, and the data is moved from the on-premises version to the version on the cloud.

The benefits of this approach include a fast time to value, little need for requirements analysis, few modifications to the system, and a cloud-based version that operates the same way the legacy system did. Among the drawbacks is the fact that the new system does not provide any new functionality or address emerging analytics requirements.

- **MINIMAL RENOVATION:** In this approach, improvement opportunities are identified by leveraging the scalability and the performance characteristics of the cloud and the hosted platform. For example, instead of replicating the same vendor's relational database management system as a cloud implementation, you might select a vendor whose data warehouse system was specifically designed to work in a cloud environment, allowing you to exploit integrated optimizations.

Improvements include faster data access and faster analytics because distributed execution and parallel computation reduce resource contention. In addition, cloud services that can benefit or improve performance can streamline data pipelines and accelerate processing. The most significant drawback remains: the migrated system is still only providing the same functionality as the legacy on-premises implementation.

- **MODERNIZATION:** In this approach, the business processes supported by the data warehouse are reviewed to identify new business requirements and evaluate ways that a rearchitected, cloud-based data warehouse solution can satisfy current and new requirements. Subsequently, a rearchitected, cloud-based data warehouse can become the centerpiece of an analytics environment that consolidates the siloed on-premises BI and reporting applications.



6

DEVELOP THE MIGRATION PLAN

After assembling a strategy for migrating data warehouses for reporting and analytics, the next pragmatic task is to develop a migration plan consisting of a sequence of practical steps. An example migration plan encompasses these tasks:

- **DETAILED ASSESSMENT:** Collect intelligence about the existing systems that have been identified as candidates for migration and provide a more complete characterization of their architecture, functionality, consumer communities, performance expectations, and suitability to a cloud environment. Consider the different migration scenarios.
- **PRIORITIZATION:** Evaluate systems based on predefined criteria to characterize the benefits and drawbacks of moving to the cloud, prioritize based on those measures, and select one or more systems for migration.
- **APPROACH:** For each system or application to be migrated, specify an approach based on the migration scenario and the anticipated benefits. Choose a cloud host as the target for the migrated application.
- **DESIGN:** Design an application architecture for the system to be migrated to the selected cloud environment. If necessary, rearchitect the underlying application design, especially if opting for an alternative system platform.
- **CONFIGURATION:** Identify the cloud resources to be used, determine the sizing, and configure the resources. If the cloud host's services accommodate auto-scaling, specify the criteria for automated scalability.
- **SYSTEM DEVELOPMENT:** Develop modifications to the on-premises implementation that need customization for the cloud. Design and develop any new code for the migrated application.
- **PROCESS DEVELOPMENT:** Develop the migration sequence and develop continuous update utilities to allow for dual operation until the migrated system's acceptance criteria have been met.
- **DATA MIGRATION:** Bulk transfer system data along with any necessary data reconfiguration.
- **INITIAL TEST:** Launch and test the cloud environment and the system configuration with a sample of the migrated data.
- **FULL MIGRATION:** Migrate the full data set into the target application environment on the cloud.
- **DUAL OPERATION:** While the legacy system remains in operation, operate in parallel and incorporate continuous updates. Implement quality controls to demonstrate that the new systems operate correctly.
- **ACCEPTANCE AND TRANSITION:** Shut down the legacy system and transfer all operations to the migrated system.

Of course, this is a rough outline of a migration plan, but it covers the key steps in the transition.

AFTERWORD

This checklist has focused on fundamental steps for preparing your organization to migrate reporting and analytics applications to the cloud by examining three different stages of the process: understanding and sharing the value of migration; assessing scenarios, potential challenges, and roadblocks; and assembling a strategy and plan for migration. An underlying theme has been that migration may be more complex than you expect.

These complexities can be overcome. Even if the array of offerings from each cloud provider is confusing, systems engineers can help determine the right sets of resources and services to meet your applications' needs. Platform providers such as those who have designed their products to be deployed on the cloud have already considered many of the issues that their customers have faced as they migrate their on-premises systems to one or more cloud environments.

Taking a measured and thorough approach to assessing the reporting and analytics landscape and prioritizing data warehouse workloads that can easily be moved to a data warehouse built for the cloud will help to ease the process and provide faster time to value.



ABOUT OUR SPONSOR



Snowflake is the only data warehouse built for the cloud, enabling the data-driven enterprise with instant elasticity, secure data sharing, and per-second pricing across multiple clouds. Snowflake combines the power of data warehousing, the flexibility of big data platforms, and the elasticity of the cloud at a fraction of the cost of traditional solutions. Snowflake: Your data, no limits. Find out more at [snowflake.com](https://www.snowflake.com).

ABOUT TDWI CHECKLIST REPORTS

TDWI Checklist Reports provide an overview of success factors for a specific project in business intelligence, data warehousing, analytics, or a related data management discipline. Companies may use this overview to get organized before beginning a project or to identify goals and areas of improvement for current projects.

ABOUT TDWI RESEARCH

TDWI Research provides research and advice for BI professionals worldwide. TDWI Research focuses exclusively on analytics and data management issues and teams up with industry practitioners to deliver both broad and deep understanding of the business and technical issues surrounding the deployment of business intelligence and data management solutions. TDWI Research offers reports, commentary, and inquiry services via a worldwide membership program and provides custom research, benchmarking, and strategic planning services to user and vendor organizations.

ABOUT THE AUTHOR



David Loshin, president of Knowledge Integrity, Inc. (www.knowledge-integrity.com), is a recognized thought leader, TDWI affiliate analyst, and expert consultant in the areas of data management and business intelligence. David is a prolific author on topics related to business intelligence best practices. He has written numerous books and papers on data management, including *Big Data Analytics: From Strategic Planning to Enterprise Integration with Tools, Techniques, NoSQL, and Graph* and *The Practitioner's Guide to Data Quality Improvement*, with additional content provided at www.dataqualitybook.com. David is a frequent invited speaker at conferences, online seminars, and sponsored websites and channels including [TechTarget](#) and [The Bloor Group](#). His best-selling book, *Master Data Management*, has been endorsed by many data management industry leaders. David is also the program director for the Master of Information Management (MIM) program at University of Maryland's College of Information Studies.

David can be reached at loshin@knowledge-integrity.com.

