



Gold
Microsoft
Partner
 Microsoft

Azure
Expert
MSP

Modernizing Legacy Data Warehouse to Synapse

WHITEPAPER

Contact Us: Microsoftsales@Hexaware.com

www.hexaware.com | © Hexaware Technologies. All rights reserved.



Contents

Sr.	Section	Page
1	Executive Summary	3
2	Introduction	4
3	On-premise Data Warehouse	5
4	Cloud Data Warehouse	6
5	Key Metrics: On-premises vs. Cloud	7
6	What benefits does Azure provide?	8
7	How do you migrate to cloud?	9
8	Key Metrics for Data Warehouse Migration to Cloud	11
9	AMAZE® for Data & AI – Value Offering	12
10	Engagement Model	14
11	AMAZE® & Hexaware Value Differentiators	15
12	Proof Point – Oracle Exadata to Azure Synapse Migration	16
13	Closing Statement	17
14	References	18
15	About Hexaware	19



AUTHOR

Dr. Ashis K. Pani

Dean and Professor Information Systems
XLRI - Xavier School of Management



CO-AUTHOR

Aditya Kaul

Product Manager
Hexaware Technologies

Executive Summary

Enterprises looking to harness the power of data transformation are moving away from traditional data warehouses that provide a siloed, proprietary and retrospective approach, to a cloud-based ecosystem that harnesses *agility, flexibility and hyper-availability* of real-time data & analytics.

However, according to leading analysts **over 50%** of the cloud migrations fail, go overbudget or run late. The root-cause of such projects points towards adoption of a *sub-optimal migration policy, incomprehensive risk assessment, and hidden costs* cropping up midway during the migration.

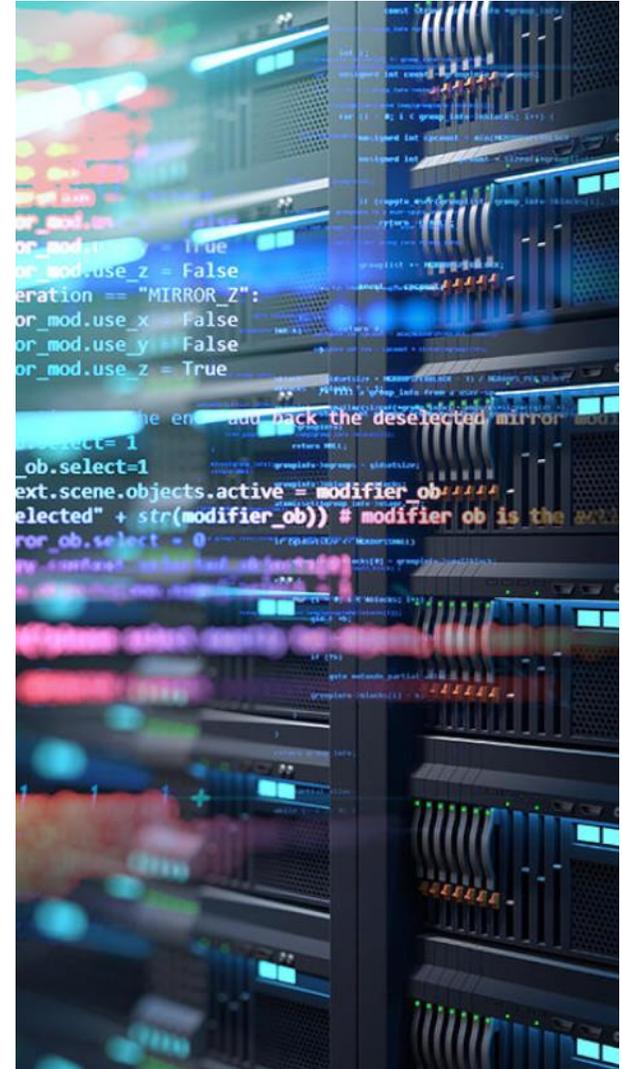
The decision of whether to migrate or not comes down to three essential parameters: **Total Cost of Ownership** (which takes care of both purchasing and operating costs), **Migration cost and timeline**, and **Process transparency** (including what data is being used at each stage).

With presence in over 60+ regions, Microsoft Azure is the fastest growing cloud platform. Its distinguished CSP network security with over 90+ certifications lends credibility to their efforts of creating a high functioning and secure cloud environment.

As cloud transformations continue to accelerate and churn out inefficiencies, **re-platforming** legacy data warehouse ecosystems to a hybrid platform seems to derive the maximum potential. Such insights are contingent upon understanding of **current DW ecosystem**, ability to generate **insights into criticality and risks** of migration, and ability to bring **automation and factory model approach** to the migration.

amaze[®]

AMAZE[®] for Data & AI (Hexaware's proprietary DW modernization tool) assesses current DW landscape by generating **cloud readiness report** and providing **fit-gap analysis** before performing a phase-wise migration.



Introduction

Businesses that have successfully migrated to cloud over the years have realized tremendous growth in their revenue streams and operations while simultaneously making their business more agile and responsive to the demands of a constantly changing environment^[1]. Driving these digital capabilities has resulted in the **cloud migration market growing at roughly 30%**^[2]. The rate of adoption has been accelerated by the pandemic where the world is looking for a more sustainable solution.

In the case of data migration specifically, **by 2022, 50% of new system deployments** will be based on a cloud data management ecosystem **and by 2023, 75% of all databases** will be on a cloud platform^[3].

However, many large organizations run the risk of an unsuccessful data migration to cloud due to a host of problems ranging from defining the right cloud architecture to operational inefficiencies in the migration process^[4]. There are also **budgetary overruns** that render the objective of migrating to cloud seem meaningless. As a result, companies face delays that translate into **unmet business potential** or, in some cases, **cloud repatriation**. Therefore, understanding the data warehouse landscape and its challenges becomes important to realize the benefits of cloud migration^[5].

Data Warehouses make use of **ETL tools** (Extract Transform and Load) which extract the data from various data sources, transforms it in the staging area and then finally, load it into the Data Warehouse system. Users can access the processed data in the Data Warehouse through Business Intelligence tools, SQL clients, and spreadsheets.

Data Warehouse on cloud differs from Data Warehouse on-premises in that it is a **platform-as-a-service (PaaS)**. A business pays for the storage space and computing power it needs at a given time and scalability is a simple matter of adding more cloud resources.





On-Premises Data Warehouse

Data Warehousing usage became prominent during the late 80s, and which saw a rapid advancement in the late 90s and early 2000s to become a core part of the Information Technology group across large enterprises.

Some of the vendors like **IBM Netezza**, **Teradata**, etc. started offering customized hardware to manage data warehouse architectures within state-of-the-art machines, to be on top of the list of priorities for CIOs and CTOs.

The benefits of having data warehouses on-premises



It allows a greater degree of control vis-à-vis usage, cost of company's hardware & software, and access grants.



Negligible network latency if any. Allows for tightly governed information security protocols.

However, since the hardware capacity is prepared for feeding the peak and ad hoc demands, utilization of resources isn't optimal, which translates to higher **CapEx and OpEx**.

Additionally, maintenance and support of both hardware and software adds to the **total cost of ownership (TCO)**. Moreover, the IT infrastructure team has full responsibility to ensure that the underlying infrastructure stays up and running efficiently, reliably, and securely.

Cloud Data Warehouse

Data Warehouses on cloud stand to gain where their on-premises counterpart fails^[6]. Capital Expenditure (or CapEx) costs are cut straightaway substituted by the pay-as-you-go model.

Cloud data warehouses also **have elasticity and scalability** to balance the data load, increase processing speeds, and serve multiple areas of the business in different geographies.

Concerns around data security are also addressed with the use of **virtual private networks (VPN)** and cloud encryption services like multi-factor authentication, making transporting of data across regions and resources incredibly secure.



Key Metrics: On-premises vs. Cloud

Whether to keep Data Warehouses on-premises or on-cloud boils down to three Key Metrics:



Total Cost of Ownership

Broadly divided into the server, storage, compute, and personnel costs, it helps you to precisely calculate **economic impact during the life cycle of an IT project(s)**.

By migrating to cloud, you can add / remove/ modify resources according to your need which frees up unnecessary costs, thereby reducing overall TCO

[Click Here](#)

to estimate the cost savings you can realize by migrating your workloads to Azure



Downtime per year

The amount of time a server or application is online and accessible to the end user provides for a **greater SLA** and, therefore, lesser downtime.

Cloud users get to experience SLAs of the order of **99.99%**, compared to 99%-99.9% for on-premises. This translates into lesser time needed for support and maintenance



Time spent on data analytics

Businesses today are run on time-sensitive queries and reports, demanding companies to **reduce the time spent on non-value-adding services** such as support and maintenance.

Because legacy infrastructure is complex, we often hear that businesses continue to invest in hiring people to manage those outdated systems, even though they're not advancing data strategy or agility.

Comparing Key Metrics: On-premises vs. Cloud



To put it into perspective, we compare the key metrics for a US Enterprise looking to migrate from Teradata Vantage (on-premises Data Warehouse) to Azure Synapse (cloud Data Warehouse)

Key Metrics	On-premises	Cloud
TCO (over 5 years)	\$26.83 Mn	\$13.14 Mn
Downtime per year	8.76 Hours	52.56 Minutes
Time spent on data analytics	15%	100%

What benefits does Azure provide?



With a presence in **over 60+ regions** around the globe and 7 on their way, Azure is expanding its services far and wide to address user concerns related to availability and latency.

Policy measures pertaining to different geographies are also thereby taken care of

[Click Here](#)
to know which regions
Azure serves in



Many Windows and Linux-based software are also cloud compatible, which translates to cost savings in the form of a **single license fee** (for both on-premise and cloud) and the **amount of re-platforming required**.

Additionally, it also offers a TCO savings of the order of 40%-50%.



Having the option of a **hybrid cloud architecture** allows customers to choose which data and application they want to keep on-premises and which ones they want to migrate to cloud.

The benefit of such a structure is that they have a greater degree of control compared to a public cloud.



Microsoft Azure is the **most secure CSP network with over 90+ security certifications** that ensures only authorized users get access to the data.

Access to customer data is denied, and the least privilege is granted in case of performing tasks such as **auditing and log access**. The access-control requirements are established as per the CSP's Security Policy

[Click Here](#)
to know how Azure
Security Policy works



How do you migrate to cloud?

The journey to cloud is long and arduous, which is why it becomes important to understand the mission-critical aspects of cloud migration.

In principle, there are two major ways in which one can migrate legacy systems to data warehouses on cloud:



Rehosting

The quickest and easiest approach is to rehost (Lift and Shift) legacy data warehouse technologies to next generation hybrid, cloud, MPP, or open-source databases. However, organizations cannot simply lift and shift their data warehouses as this would be **suboptimal use of cloud economics**. It would also mean that some of the cloud-native benefits, such as continuous real-time deployment, may not be possible^[7].



Re-platforming

Re-platforming, on the other hand, is the most popular approach as it is optimized for cloud use, whether you will **rip-and-replace** the warehouse's primary platform or **augment** it with additional data platforms. Data engineers migrate the entire data schema, data pipeline and visualization tools to cloud.



Key challenges in migrating to cloud



Ability to understand the **complexity of the current state DW** environment across all the three layers of Data, Data Pipeline, and Data Visualization for effective **new data architecture** development.



Ability to **get insights** into the criticality and risks to put together a comprehensive migration plan addressing all major business impact areas.



Ability to bring in **automation** and **factory model** in the cloud native migration across all the three layers of Data Pipeline, Data and Visualization.



AMAZE® for Data & AI. is a cloud re-platforming service that takes account of the aspects of automation and necessary human expertise to deliver across all key metrics requisite for a successful cloud migration.

amaze®

Key Metrics for Data Warehouse Migration to Cloud

Migration can be done either manually by the customer or with the help of Managed Service Providers (MSPs).

The following key metrics help understand which option of the two is better:



Cost of Migration

Most third-party cloud data warehouse migration tools available in the market come with a **high license cost**, don't offer customization and none of the platforms cover the migration of all the three layers of Data Pipeline, Data and Visualization. Additionally, the **level of automation** provided by such tools further adds to the total cost



Duration or Time to Market

The time it takes for a customer to migrate their Data Warehouses from on-premises to cloud. Most solutions in the market automate schema but offer a **manual approach for ETL code conversion and report migration**. End-to-end automation reduces both time and probability of an error occurring during migration^[8].

Also, one of the biggest pain points for customers is the disruption caused in the day-to-day operations during migration. A **well-planned phase-wise migration** with AMAZE® brings down the overall downtime for migration.



AMAZE[®] for Data & AI – Value Offering



Customers, while manually migrating data warehouses to cloud, typically face a **higher cost of migration and a schedule overrun**, among other perils (e.g.: data security). With AMAZE[®], however, they can realize obvious benefits within those metrics.

Putting into the context of the same US Enterprise:

	Cost of Migration	Time to Market
Manual Migration	USD 980K	6 Months
AMAZE [®]	USD 390K	3.5 Months
Improvement	60% Reduction	40% Faster

While it is reasonable to assume that roughly **50%-80%** of the work can be automated, some key processes such as the configuration of data and analytics for cloud, creating new environment dependencies, testing, and validating data need human intervention.

Data 		Data Pipelines 			Data Visualization 		
On-Premise Technology	Synapse	On-Premise Technology	Lift & Shift	Re-platform ADF	On-Premise Technology	Lift & Shift	Re-platform ADF
Microsoft SQL Server	<input checked="" type="checkbox"/>	Microsoft SQL Server Integration Services	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Microsoft SQL Server Reporting Services	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
ORACLE	<input checked="" type="checkbox"/>	ORACLE DATA INTEGRATOR	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	SAP BusinessObjects	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
teradata.	<input checked="" type="checkbox"/>	Informatica	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	COGNOS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
NETEZZA	<input checked="" type="checkbox"/>	sas	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	MicroStrategy	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
ORACLE EXADATA	<input checked="" type="checkbox"/>	talend	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	+ a b l e a u	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
amazon REDSHIFT	<input checked="" type="checkbox"/>	IBM Datastage	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	Qlik Q	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
SYBASE	<input checked="" type="checkbox"/>	SAP BODS	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	ORACLE BUSINESS INTELLIGENCE	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
snowflake	<input checked="" type="checkbox"/>						
TCO Savings  ~40-50%				TCO Savings  ~10-15%	~25-30%		
Automation  80%				Automation  70%			
		TCO Savings  ~4-6%			~50-60%		
		Automation  60%					

Automated Analysis Automated Analysis & Conversion

Re-platforming (Cloud Migration) Process

The entire re-platforming (cloud migration) process is divided into three phases:



Data Capture & Assessment

In the Data Capture & Assessment phase, we analyze all the three layers of the customer's data warehouse and analytics landscape that covers **Schema, Data ETL scripts and Visualization reports**. The outcome of this stage is a **cloud-readiness report** that lays out the complexity of the environment and measures migration readiness while highlighting the keywords/syntax from on-premises code that will be migrated in an automated way as well as need manual intervention.

AMAZE® uses **Metadata extraction** while assessing the landscape. It ensures that no customer data is exposed (or manipulated), thereby improving the security of the process. This phase is almost fully automated (**95% and above**) owing to AMAZE®'s automation-based discovery tool.



Readiness and Build

Next, in the Readiness & Build phase, based on the assessment results and the services required by the customer, the best cloud architecture (hybrid/single/multi) is laid out. Then with the help of a cloud-readiness report from the assessment phase, AMAZE® **maps the objects** of data warehouse, ETL pipelines and reports from source to target. This object mapping also helps to **identify the gaps** which would require manual intervention while migrating.

In most of the scenarios, simple objects are easily mapped to the target which are converted automatically, while medium and complex objects are partially converted manually. We also **provision the cloud resources** laid out in the target state architecture.



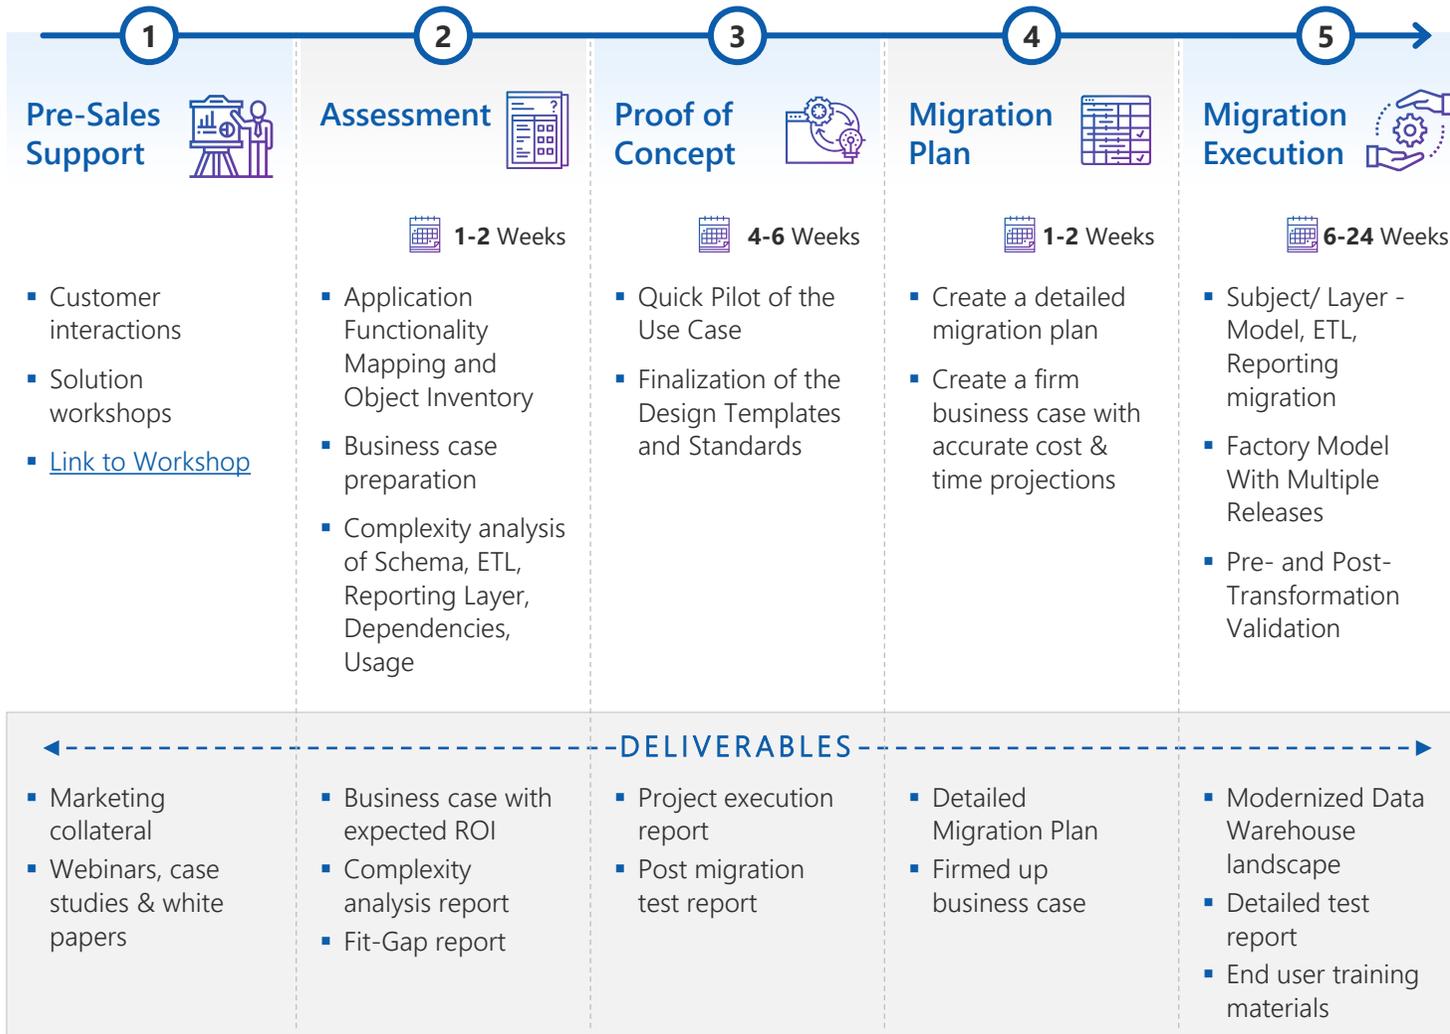
Migration to Azure phase

Lastly, the Migration to Azure phase sees an **incremental phase-wise implementation** of data ecosystem migration (data schemas, objects, and pipelines) and conversion (reports and dashboards). We avoid a big-bang migration approach which can potentially prove risky and lead to a significant amount of downtime. In each of these phase-wise migrations, **key metrics, benchmarks, guidelines, and best practices** are agreed upon mutually.

The biggest value add of this approach is that the entire migration process happens with significantly lesser downtime and since no externally licensed product is used for migration, **50%-60% cost savings** is realized in the customer's data ecosystem.

Engagement Model

We follow a field tested & process driven execution model to guarantee repeatable outcomes with exceptional & consistent quality while going the extra mile to deliver great customer experience



Call to Action

An offer to get started



- Assessment Workshop to generate sizing report & Data Warehouse re-platforming possibilities
- We can access ECIF funding for the workshop / POCs
- Create business case with defined benefits and expected ROI from the engagement

Joint Proof of Concept



- POC for migrating specific workloads identified as part of the assessment
- Create a detailed modernization roadmap with target recommendations & timelines

AMAZE[®] & Hexaware Value Differentiators

amaze[®]



Time to Market

- Automated data, pipelines & visualization assessment, and accelerated migration
- Re-platform on-premises data landscape to cloud within 5 to 7 months



Proven Migration Framework

- Encrypted metadata driven repeatable data migration framework
- Delivers consistent quality and savings
- Follows zero data leakage policy



Cost Effective Migration

- No external licensed products used for migration
- Optimization on Azure translates into significant TCO savings



Faster
40-60% faster migration



Better
Up to 80% reduced migration issues



Cheaper
40-50% Reduction of TCO

Hexaware Differentiators



IP CO-sell Ready Offering

MS sellers can leverage our IP validated offerings and retire their quotas faster



Azure Expert MSP & AMP

Best in class, leveraging ECIF funding available for assessments, POCs and Migrations



Microsoft CAF Enabled

Consistent delivery of quality through repeatable platforms validated by MSFT

Proof Point – Oracle Exadata to Azure Synapse Migration



Customer Profile

US headquartered
Biopharmaceutical
Enterprise

\$2.5 Bn
Revenue



20,000+
Employees
worldwide



Challenges

- **Slow time to market:** Real time data ingestion wasn't possible from current DW and analytics landscape required considerable manual interference & ops overheads.
- **Absence of secure migration framework:** Customer wanted accelerated and secured migration to cloud with minimal downtime and cost overruns
- **High migration cost:** In absence of a planned migration framework, manual migration was both risky and expensive
- **High TCO:** Storage & processing of IoT data generated from clinical trials & research was proving to be extremely expensive on Exadata

Solution

- Automated analysis and conversion coupled with our subject matter expertise helped optimize the process, which in turn **shortened the time to market**
- We delivered consistent quality & cost savings with our repeatable migration framework and executed **secured data migration, with zero data loss**
- Zero license cost & high automation level across data, pipeline & visualization **reduced migration cost** significantly.
- Using Oracle Exadata to Azure Synapse analytics, we were able to **optimize the TCO** of Data Warehouse & Analytics landscape

Benefits

**60%
Faster**

Time to Market
Improvement



**50-80%
Better**

Automation while migrating
data, data pipeline & data
visualization



**45%
Cheaper**

Total Cost of
Ownership
Reduction



Closing Statement

Additionally, the analytical capabilities of the customer can be enhanced by making use of cloud native development of applications and open-source projects.

Doing business using cloud data warehouses is no longer a thing of the future as every organization looks towards adopting cloud for far-reaching benefits. It therefore makes sense to invest in technologies such as automation and real-time analytics to achieve a successful migration of legacy Data Warehouses or systems at low cost & risk and high speed & return.



References

1. TechRepublic.com. 2020.
Report: 95% Of Businesses Have Migrated Critical Applications to The Cloud. [online]
Available at: <https://www.techrepublic.com/article/report-95-of-businesses-have-migrated-critical-applications-to-the-cloud/>
2. Mordorintelligence.com. 2020.
Cloud Migration Market | Growth, Trends, Forecasts (2020 - 2025).
[online]
Available at: <https://www.mordorintelligence.com/industry-reports/cloud-migration-services-market>
3. Thanaraj, R., 2020.
Cloud Data Warehouse: Are You Shifting Your Problems To The Cloud Or Solving Them?
Gartner Research.
4. Friedman, T., 2009.
Risks and Challenges in Data Migrations and Conversions.
Gartner Research.
5. Friedman, T., 2009.
Best Practices Mitigate Data Migration Risks and Challenges.
Gartner Research.
6. Smith, D., 2017.
Cloud Strategy Leadership | Gartner Insights on How and Why Leaders Must Implement Cloud Computing.
Gartner Research.
7. Business Reporter Production. 2020.
Migrating Legacy Applications to The Cloud.
Forbes
8. Transforming Data with Intelligence. 2019.
Plan Carefully When Migrating to A Cloud Data Warehouse. [online]
Available at: <https://tdwi.org/articles/2019/11/07/ta-all-plan-carefully-when-migrating-to-cloud-data-warehouse.aspx>

About Hexaware

Hexaware has a global presence across the Americas, EMEA & APAC to ensure we cover every time zone, language zone & regulatory zones to effectively serve our customers. When it comes to our Microsoft partnership, we have the necessary expertise with over 2000+ Microsoft certified consultants and many of those consultants holding multiple certifications.

We are an Azure expert MSP (Managed service partner) with over 410% ACR growth in 2020. Our offerings are Microsoft CAF (Cloud adoption framework) enabled which ensures consistent delivery quality through repeatable platforms which are validated by Microsoft.

Learn more about Hexaware at <https://www.hexaware.com>

Contact us: Microsoftsales@hexaware.com



HEXAWARE



NA Headquarters

Metro 101, Suit 600,101
Wood Avenue South, Iselin,
New Jersey – 08830
Tel: +001-609-409-6950
Fax: +001-609-409-6910



India Headquarters

152, Sector – 3
Millennium Business Park
'A' Block, TTC Industrial Area
Mahape, Navi Mumbai – 400 710
Tel : +91-22-67919595
Fax : +91-22-67919500



EU Headquarters

Level 19, 40 Bank Street,
Canary Wharf,
London - E14 5NR
Tel: +44-020-77154100
Fax: +44-020-77154101



APAC Headquarters

180 Cecil Street,
#11-02, Bangkok Bank
Building, Singapore - 069546
Tel : +65-63253020
Fax : +65-6222728



Australia Headquarters

Level 26, #44, Market
Street Sydney,
NSW 2000, Australia
Tel : +61 2 9089 8959
Fax : +61 2 9089 8989



Safe Harbor Statement

Certain statements in this press release concerning our future growth prospects are forward-looking statements, which involve a number of risks, and uncertainties that could cause actual results to differ materially from those in such forward-looking statements. The risks and uncertainties relating to these statements include, but are not limited to, risks and uncertainties regarding fluctuations in earnings, our ability to manage growth, intense competition in IT services including those factors that may affect our cost advantage, wage increases in India, our ability to attract and retain highly skilled professionals, time and cost overruns on fixed-price, fixed-timeframe contracts, client concentration, restrictions on immigration, our ability to manage our international operations, reduced demand for technology in our key focus areas, disruptions in telecommunication networks, our ability to successfully complete and integrate potential acquisitions, liability for damages on our service contracts, the success of the companies in which Hexaware has made strategic investments, withdrawal of governmental fiscal incentives, political instability, legal restrictions on raising capital or acquiring companies outside India, and unauthorized use of our intellectual property and general economic conditions affecting our industry.