



White Paper

# Oracle to SQL Epic Migration

## Introduction

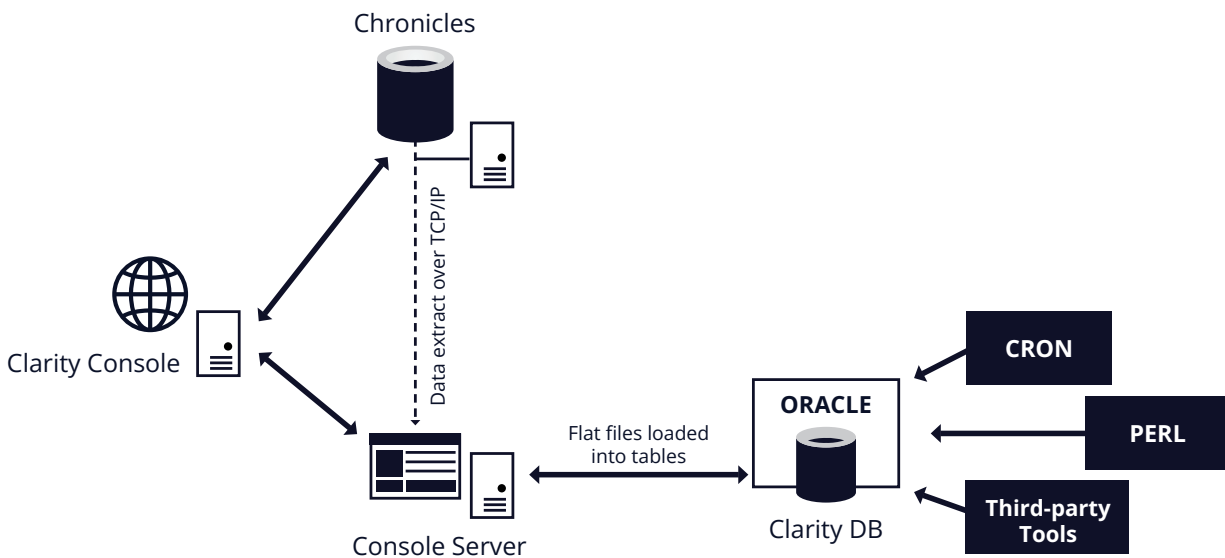
Starting in 2027, Epic will no longer support Oracle for Cogito Analytics. This puts immediate pressure on healthcare organizations leveraging Oracle for Cogito to start planning their re-platforming and migration journey to transition to Cogito on SQL Server. This may seem like a daunting and expensive task, however, we share examples further in this document that lead to a 72% reduction in licensing and support costs, while also improving performance and accelerating cloud adoption. Migrating an Epic Cogito Analytics environment from Oracle to SQL Server is often a strategic decision in an overall strategy to modernize and streamline existing workflows and processes.

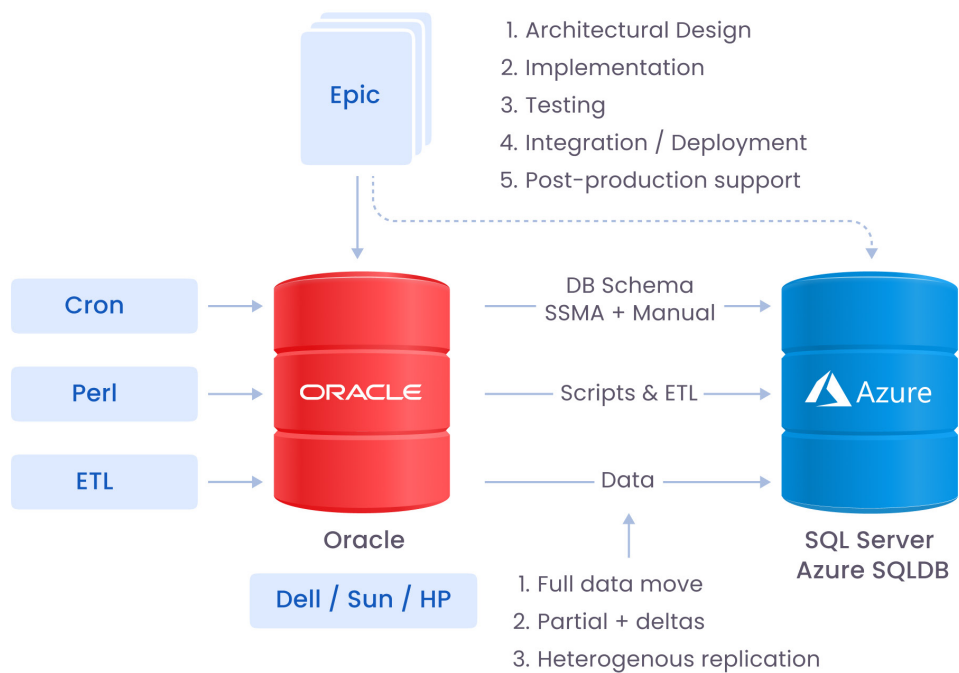
Moving Cogito from Oracle to SQL Server allows organizations to cut costs and eventually take advantage of cloud-native features that may not be compatible with Oracle. Oracle is notoriously expensive when considering licensing and support. These costs quickly increase over time as Cogito environments continue to grow and Oracle support contracts expire.

In recent years, there was a push centered around incorporating artificial intelligence (AI) and machine learning (ML) in DSS systems. Cloud native features in Microsoft Azure, which are detailed later in this document, allow healthcare organizations to more quickly adopt AI and ML. Modernizing the backend, taking advantage of cloud native services, and streamlining data process flows makes data more useful for hospitals.

## Migration Process

The diagram below depicts what a traditional Oracle-based Cogito environment looks like. There are up to 2,400 tables and views in various schemas and sub systems with a large amount of stored logic. This makes migrating to SQL Server a challenging task. However, when breaking it down into different modules, it becomes manageable. The Data Mart and Reporting services modules are a good place to start the modernization journey by migrating them to SQL Server.





The typical Epic database migration includes hundreds to thousands of tables, views and reports totaling multiple terabytes of data in a custom Oracle database. Initially, an analysis of the existing database is necessary to determine the appropriate plan for the migration to minimize disruptions for your practice. It's a typical scenario to encounter customizations within the Epic reporting database. Epic allows customers to make customizations for reporting, ETL, and stored procedures. There also may be custom scripts for various purposes.

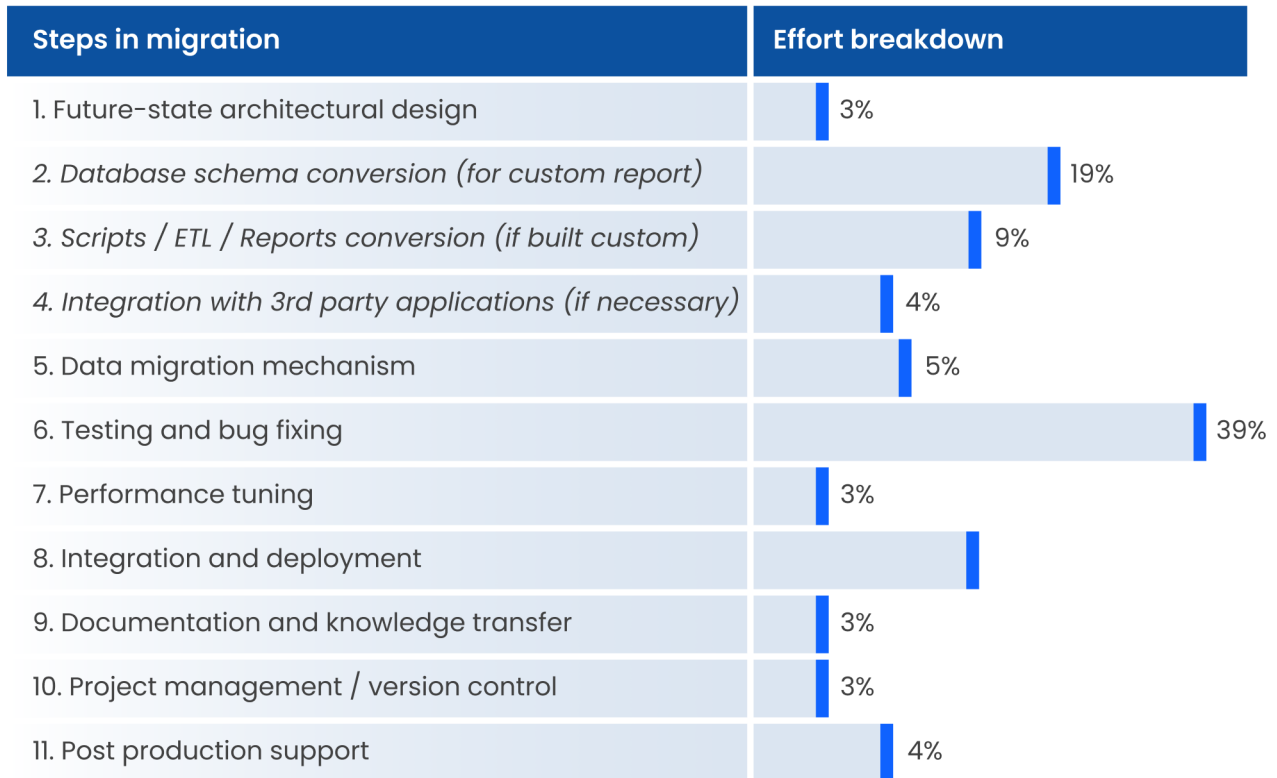
After the analysis, the next step in a migration is to convert the database schema. This includes tables, indexes, constraints, defaults, rules, stored procedures, triggers, packages, sequences, views, functions, and so on. Since the plan is to migrate within the same software vendor, the schema already exists. Only the customizations need to be moved.

Next, any custom code written for Oracle is converted for SQL Server. During its legacy, our migration partner created the SQL Server Migration Assistant (SSMA) to help their customers migrate more efficiently. They have an intimate working knowledge of this tool. SSMA was designed to automate aspects of database migrations to SQL Server from various databases. While this is an excellent tool, it doesn't convert 100% of the code automatically, so be prepared for some manual efforts.

Custom reports are commonly found in Epic and many other off-the-shelf systems. These custom reports need to be converted for the new database. This is another area where several approaches are possible. Our migration partner reduced the number of reports to convert by 85% by establishing lineage between reports and determining which reports were derivatives of others. They then built a tool to extract the query from a report, convert it via SSMA, then feed it back into the report.

The most difficult step is usually the data migration itself. Whether it's an initial full move of the data or a partial data move while capturing changes while migrating the remaining data, there are diverse options to complete the data migration. Depending on the downtime requirements and a few other factors, this influences the method and tools used to move the data.

Our migration partner, ABCloudz is a full-service technical consulting firm with over 20 practices. Their legacy includes the creation of SQL Server Migration Assistant (SSMA) and Schema Conversion Tool (SCT). ABCloudz provides services around platform, application, and database **Migration/Modernization**, application and analytic systems **Development**, and **Managed Services**.



*(The steps in the migration process as described above apply to Epic to SQL Server or Azure SQL Database. For simplicity, the destination is referred to as SQL Server. These steps may vary slightly for other COTS applications.)*

## Cost Analysis

The scenario below highlights the cost savings that can be achieved by migrating from Oracle to SQL Server.

In this example, the Production Clarity DB running on Oracle is sized with 128 vCPU, which, in turn, requires 64 Oracle processor licenses due to the Intel 2:1 core factor. We also need to consider the annual support costs, which is traditionally 23% of the licensing costs. At list Oracle pricing, taken from Oracle's Technology Guide Price List, this costs \$47,500 per processor per year plus any licensable add-ons like Advanced Security Option, Partitioning, In Memory, etc. In this example, it is assumed that Oracle Partitioning is in scope as a licensable add-on as it is common in data warehousing environments. Oracle Partitioning costs \$11,500 per processor per year at list price. The first-year costs are as follows:

<b>Oracle Processor Licensing 64 x \$47,500</b>	<b>\$3,040,000</b>
<b>Oracle Partitioning Add-on 64 x \$11,500</b>	<b>\$736,000</b>
<b>Oracle Annual Support</b>	<b>\$699,200</b>
<b>Total</b>	<b>\$4,475,200</b>

If we run the same exercise under the assumption that the customer receives an Enterprise Level 65% discount, we see the following costs for the first year:

<b>Oracle Processor Licensing 64 x \$16,625</b>	<b>\$1,064,000</b>
<b>Oracle Partitioning Add-on 64 x \$4,025</b>	<b>\$257,600</b>
<b>Oracle Annual Support</b>	<b>\$244,720</b>
<b>Total</b>	<b>\$1,566,320</b>

Now let's compare that to first year SQL Server licensing and support costs calculated below, that are taken from the Azure Pricing Calculator:

<b>MS SQL M128 (128 vCPU)</b>	<b>\$420,480</b>
<b>MS SQL Professional Direct Support</b>	<b>\$12,000</b>
<b>Total</b>	<b>\$432,480</b>

Even with a 65% Oracle discount, the same configuration for SQL Server is 72% more cost effective than the Oracle configuration.

Taking this a step further, organizations often need to overprovision vCPU in the cloud to get a higher level of storage performance that is required for the same level of storage performance that they were achieving on-prem. Using the same example as above, let's assume that we need move to the next largest VM size in order for the database to perform better than, or equivalent to, it was running on-prem.

<b>MS SQL M208 (208 vCPU)</b>	<b>\$683,280</b>
<b>MS SQL Professional Direct Support</b>	<b>\$12,000</b>
<b>Total</b>	<b>\$695,280</b>

Although it is still much cheaper than the Oracle solution on-prem, or in the cloud, there is a 38% increase in licensing costs due to the need to overprovision.

Silk provides a virtual storage platform that is built on native cloud resources for databases that need ultra-high storage throughput to connect into and leverage the data network rather than the storage network. Because the Silk platform utilizes the data network, the performance of the database is not bound by the limits of smaller VM shapes. This allows Silk to help organizations reduce overprovisioning across the estate. Silk allows organizations to lift and shift databases to the cloud achieving the same, or better, performance as on-prem, while utilizing the same amount of vCPU. This helps organizations control database licensing costs.

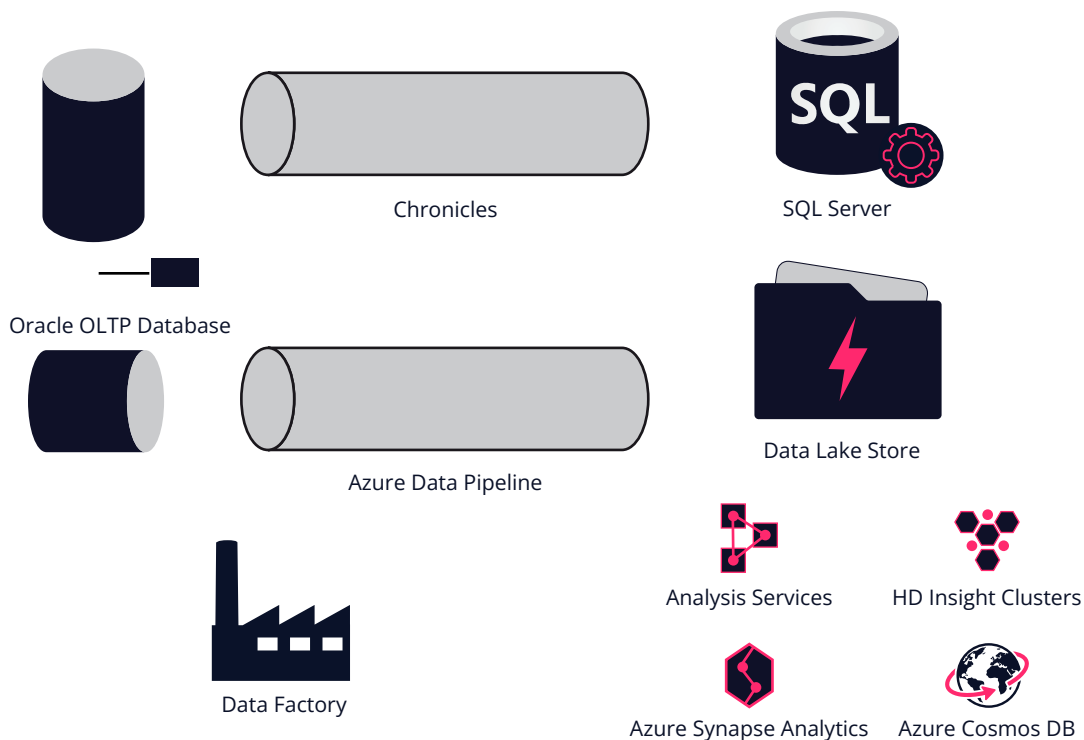
In the example above, the Silk platform helps the customer save 38% on licensing costs, which equates to \$262,800 annually on just one Clarity database.

With the ability to dynamically scale environments on-demand in Microsoft Azure, futureproofing is easier than ever before. This means adding computing power and storage is now done on-demand. With Microsoft Azure, there are no up-front hardware costs, there is no need to budget for replacement hardware, and licensing is built into your monthly fee when taking the Cloud Native approach. For the purposes of budgeting, this activity is called “FinOps.” Similarly, to DevOps, FinOps is a cultural practice. This is the practice of further managing cloud costs to gain more financial control and predictability. When you layer your databases on top of Silk in the Azure cloud, the Silk and Azure combination allows organizations to control costs, reduce cloud waste and eliminate overspending.

While Silk may help to right size cloud infrastructure, additional features that Silk provides include the ability to deduplicate and compress data in real time, further improving overall data utilization. Coupled with Silk’s ability to provide thin, instantaneous snapshots for refreshing lower-level environments, reporting and operational recovery, Silk may also increase the return on cloud investment exponentially.

### Future State: Cloud Native Services

As stated earlier, a migration from Oracle to SQL Server may allow for a more seamless transition to the cloud, which will allow organizations to take advantage of cloud native services. [Azure Data Lake](#) is a common cloud native service that Epic organizations, especially academic and research organizations, leverage in the cloud. In this example, it is assumed that some of the newly migrated SQL Server DBs supporting Cogito will now become a tributary to a larger [Azure Data Lake](#). Using Azure services like HDInsight, Data Lake Analytics and Azure Data Lake Storage, allows organizations to build a state-of-the-art Big Data Analytics platform. The customer is also able to utilize existing operational datastores that may still exist within the legacy Oracle database. The diagram below depicts this type of configuration:



There are two things to note here. First, we have retained the MS SQL DSS environment. Presumably progress is being made to port/migrate other PL/SQL based logic to MS SQL stored logic so that the workload on the primary Oracle OLTP server is diminishing. The MSSQL environment is also still being used for canned Line of Business (LOB) reporting. Meanwhile we have also added a plethora of new SQL/ NON-SQL analytics and Big Data options, thus increasing the value of the legacy datastore. Below is a list of cloud native data services that may be useful to Epic customers.

## Analysis Services

This PaaS solution offers data modeling capabilities far beyond relational models and allows you dynamically refactor schemas and data models.

## Azure Synapse Analytics

Azure Synapse Analytics is a Data Engineering design studio that offers integration, management, monitoring, and security services as a PaaS offering.

## Azure HDInsight

Azure HDInsight is a managed Apache Hadoop service that lets you run Apache Spark, Apache Hive, Apache Kafka, Apache HBase, and more on Azure.

## Azure Cosmos DB

Azure Cosmos DB is a managed, serverless NoSQL database for high-performance applications for large scale data sets.

The 'final re factoring' of an Oracle legacy database may be its eventual retirement as a datastore, which can be done incrementally. This is the best way to mitigate risk and disruption. Azure services can be added to enhance the value of the datastore or eventually replace it.

In our scenario we took an Oracle Cogito installation and moved it to an Azure + Silk IaaS solution. Its ancillary database used primarily for LOB reporting was initially ported to MS SQL Server as a cost saving measure. This is a low risk, low cost, change that can yield immediate savings.

Subsequently we added an Azure Data Lake to our architecture. Multiple non-Oracle – SQL/NO SQL/Big Data options are also available to the end user. That is the benefit of moving these workloads to Azure + Silk – quickly lift and shift databases then begin to evolve in the cloud.

## **Conclusion**

Epic's announcement to no longer support Oracle for their customers' Cogito environments is a great opportunity to drastically reduce costs by modernizing and accelerating your cloud journey. We have shown herein that there are many benefits that come with re-platforming from Oracle to SQL Server: a 72% cost reduction, the ability to take advantage of rich data services, and workflow improvement. In the cost exercise above, we have shown that a transition to SQL Server will save 72% on licensing costs. When expanded out across all Clarity databases within an Epic environment, the cost savings become substantial. The ability to take advantage of native cloud data services like AI and ML allows Epic organizations to make their data more useful. A migration to SQL Server and a migration to the cloud also helps to streamline future update and upgrade procedures.

Patient care is the sole focus of every health system. Making your data more useful by re-platforming from Oracle to SQL Server, moving to the cloud, streamlining your upgrade and update procedures, and taking advantage of native Azure cloud services will allow technical staff members to focus more on patient-centric initiatives thus improving patient care.

### **About Silk**

The Silk Cloud DB Virtualization Platform gives demanding workloads 10x faster performance on the cloud compared to native cloud alone. It is a virtualization layer that sits between the underlying cloud infrastructure and customers' workloads. Without refactoring, workloads such as Oracle, Microsoft SQL Server, and industry-specific applications can move onto the GCP and Azure cloud while massively improving user experience. Industry leaders in e-commerce, software publishing, FinTech, and healthcare trust Silk with their mission-critical workloads to get the ultra-fast speeds their customers demand. Silk is headquartered in Needham, MA. To learn more, visit [silk.us](https://silk.us).