

The Journey to the Cloud

Hybrid multi-cloud performance and financial governance optimization



Introduction

For most organizations, the journey to the cloud is well underway. Unfortunately, many cloud migration decisions are made without options evaluation and realistic performance and financial expectations. Unrealistic expectations inevitably lead to unwelcome surprises. Financial stakeholders are worried about the annual budgeting process that might need un-planned revisions. IT stakeholders are concerned with selecting the right cloud data platforms and resource allocation to meet business service level goals (SLG) continuously.

BEZNext utilizes modeling and optimization to estimate the minimum required cloud configuration and cost to support an organization's workloads SLGs. Our approach does not require months of work, and we typically deliver the results in two to three weeks. We apply the same technology to optimize cloud migration and dynamic capacity management decisions. The actual measurement data are compared with predicted results. When significant anomalies and their root causes are determined, the model provides new recommendations to meet SLGs with the lowest cost continuously. ⁱ

In this paper, we review the BEZNext approach and case studies that show the value of modeling and optimization for any organization, no matter where they are on their journey to the cloud.

The Journey

In 2020, analyst firm IDC reported that spending on IT infrastructure in the cloud had exceeded investment on-premises for the first time. As a result, cloud computing has become the standard. As in prior IT eras, money flows freely towards this new investment. Indeed, Gartner estimated that 45% of all IT spending in 2021 went towards the cloud.

Once an organization has decided to pursue a cloud strategy, many decisions must be made to complete the journey from on-premises to the cloud or manage your existing cloud as it evolves. The following lists some significant destinations along the way at a high level:

- Selecting the correct cloud platform (i.e., vendors, services, configurations)
- Choosing the best migration strategy for on-premises workloads
- Managing the ongoing performance and cost as workloads grow and change

The journey becomes increasingly complex, considering that 92% of organizations utilize a hybrid multi-cloud environment. This complexity occurs as special interest groups form within organizations around specific technologies.

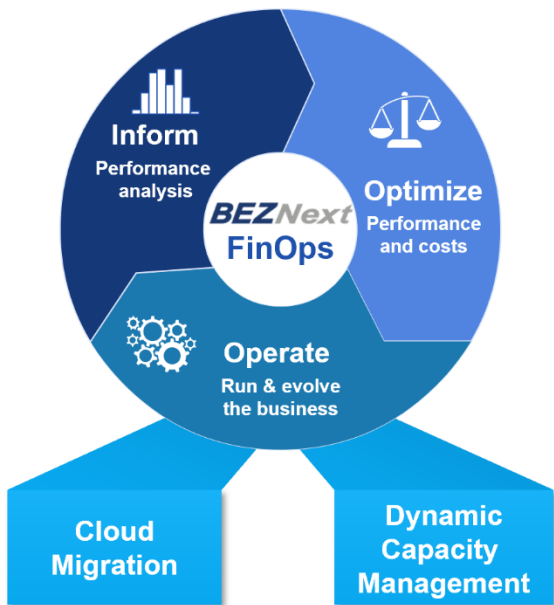
Unfortunately, many of the decisions made along the way are not based upon verifiable dataⁱⁱ. Indeed, some organizations make platform decisions before identifying the workloads that must be supported. A company will spend, on average, as much as \$12M annually for cloud services. What business leader would make such consequential decisions without data to back it up?

Gartner estimates that 80% of organizations exceed their planned cloud budgetsⁱⁱⁱ. But unfortunately, a poor decision-making process inevitably results in unrealistic goals, higher costs, and unhappy executives. Conversely, data-driven decision-making enables improved collaboration among IT, financial, and business stakeholders. In addition, it reduces

the risk of economic or performance surprises, providing stakeholders with realistic expectations faster and with greater accuracy. *BEZNext* FinOps optimization software and services deliver all these benefits.

BEZNext FinOps Optimization Software and Service

BEZNext software and services automate many of the critical FinOps functions required for a modern business to make informed, data-driven decisions supporting their ongoing journey to the cloud. Our software reflects our deep expertise in complex systems modeling and optimization^{iv}. Our team of PhDs has incorporated iterative queueing network modeling and gradient optimization to accurately predict the minimum configuration and cost required to meet the SLG of any organization's workloads.



One of the enormous benefits delivered by cloud computing is the ease at which you can scale underlying infrastructure. In 2020, Gartner listed seven workloads that should immediately move to the cloud. Among those identified were applications that benefit from horizontal scaling. Analytic workloads such as data lake and data warehouse applications are prime examples. *BEZNext* specializes in helping clients optimize their decision-making when migrating these applications to the cloud.

Data Collection

To make a data-driven decision, you need to begin with data. So *BEZNext* starts by collecting the actual performance data of your existing on-prem applications. Our software gathers measurement data continuously over time. Then, this data is aggregated into business workloads. Data warehouses and data lake environments tend to support the analytic needs of several lines of business (e.g., Sales, Marketing, Finance); each of these will have its own "character." Examples of information collected:


- How many queries are processed during different time intervals by each workload?
- What is the distribution of response times by workloads?
- How much CPU resources are used by different workloads?
- How I/O intensive are different workloads?

Workload Characterization

Once collected, the information is aggregated and correlated by line-of-business and business. We call it business workloads. Next, we work with stakeholders to formalize the business SLG of each business workload. If none has been assigned, our model will assume that the current on-prem response times are acceptable in the cloud. We also incorporate user and data growth expectations for forecasting purposes.

Modeling and Optimization

Workload characterization and performance characteristics of the cloud data platforms targets, including pricing data are input to our modeling technology. The iterative queueing network modeling and gradient optimization performs multi step automatic evaluation to predicts the minimum configuration and budget required to meet SLGs of each business workload on each of the Cloud Data Platform being evaluated.



Results Verification

BEZNext FinOps optimization software automatically compares predicted results with actual measurement data. After migration workloads to the cloud the software automatically determines performance and financial anomalies and their root causes. Our software uses identified anomalies and root causes to evaluate options and choose the most economical way to fix the problem.

Example Case Studies

To better understand the BEZNext solution, let's examine three case studies that represent BEZNext's distinct service offerings. A Fortune 500 client decided to shutter a data center and migrate on-premises data warehouse and Hadoop workloads to the cloud. The environment has been stable and well-performing for many years and is considered business-critical. In keeping with their corporate IT objectives, the decision has been made to move the workloads to the cloud. But where should they move it? How should they migrate all workloads on time and within the budget? How should they manage all business workloads in the cloud to continuously meet SLGs with the lowest cost?


Cloud Selection Decision Optimization

The client attempted to select the cloud platform by engaging with several large system integrators and the cloud data platform (CDP) vendors to evaluate potential target platforms. More than nine months were spent performing proof-of-concept exercises and benchmark tests to select the right cloud platform. Still, they were unable to answer two fundamental questions:

- What is the minimal configuration required to meet their workload's SLGs?
- How much would it cost them?

BEZNext was hired. We collected performance data for applications slated for migration. The client wanted to evaluate the price/performance of their on-premises workloads on Snowflake, Teradata Vantage, and Amazon Redshift.

BEZNext software aggregated the performance measurement data by the line of business and developed the performance, resource utilization, and data usage profiles for each business workload.



Daily workloads' resource utilization by cloud migration candidates

When predicting performance and costs for each workload on different cloud data platforms, the BEZNext software weighs the cloud platform's architecture, underlying database, and optimizer differences, including the ability to utilize indexes. The model also considers the cloud platform's underlying instance type configurations. Many instances have different vCPU counts, memory sizes, storage, and network bandwidth.

The chart below compares the monthly and yearly costs for running the client's business workloads on the different cloud data platforms. Cloud data platform vendor names are anonymized. In this chart, the findings highly depend on workloads' profiles and SLGs^v.

DBMS Vendor	Cost per	Jan	Feb	...	Nov	Dec	Annual cost
A	Month	\$234,778	\$241,453	...	\$261,479	\$261,479	\$2,964,189
B	Month	\$807,206	\$813,466	...	\$919,872	\$926,131	\$10,468,877
C	Month	\$1,255,660	\$1,255,660	...	\$1,301,740	\$1,301,740	\$15,528,720
A	Query	\$0.0040	\$0.0040	...	\$0.0040	\$0.0040	\$0.0040
B	Query	\$0.0139	\$0.0139	...	\$0.0143	\$0.0143	\$0.0141
C	Query	\$0.0287	\$0.0286	...	\$0.0198	\$0.0196	\$0.0205

BEZNext Model Estimate the Monthly and Yearly Costs for customer's workloads for each CDP Vendor

Our client was given specific answers to the two most essential questions mentioned earlier. In addition, the entire project was completed in three weeks versus the nine months spent on the prior proof-of-concept approach. As a result, the client now has solid, verifiable predictions comparing options to make the right decision. One year later, BEZNext was brought back to verify the predictions. The results showed that the actual outcome is within 2% of predictions. The confidence they gained in the process now enables the client to optimize cloud migration and management decisions, knowing that they have the tools to reduce the risk of future cloud-related financial and performance surprises.

Cloud Migration Decision Optimization

After the selection of the cloud data platform in a year, the client realized that they have to support a hybrid multi-cloud environment. The different lines of business requested different cloud data platforms. During the cloud migration, there is a need to evaluate many options of supporting workloads in a hybrid multi-cloud environment, including:

- Finding workloads suitable for each cloud data platform with minimum data replication
- How to reduce ETL load time
- How to take into consideration the impact of data tokenization and detokenization, and data compression during the configuration planning
- How to select the right cloud data platform for migrating Hadoop workloads to the cloud?

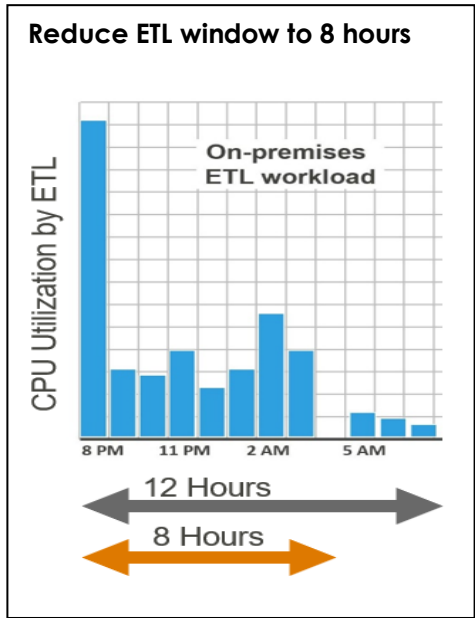
Below are several cloud migration challenges addressed during this project:

Select right cloud data platform for each workload in a Hybrid Multi-Cloud Environment

Finding workloads suitable for each cloud data platform and minimize data replication

Determine resurces needed to support security Tokenization and Data Compression Overhed

Deterimne the minimum configuration and budget needed to support SLGs with data detokenization and compression,



Cloud Platform for Hadoop Migration

Predict the minimum configuration and budget needed to support SLG for Hadoop workloads in AWS EMR and Snowflake environments.

Cloud migration decisions optimization fostered improved collaboration between IT leaders and the business stakeholders enabling a more informed, data-driven decision rather than blindly going down a path and failing to meet initially unrealistic objectives.

Dynamic Capacity Management Decision Optimization

The ability to meet SLGs for business workloads with predictable costs in a hybrid multi-cloud environment is a challenge. While this reality has been unchanged since the on-premises era, what has changed is the ability of the cloud to quickly and automatically add resources to solve performance problems. That ability, however, results in large surprise billings from the cloud data platform providers. FinOps was envisioned to address performance and cost issues and reduce the risk of financial surprises.

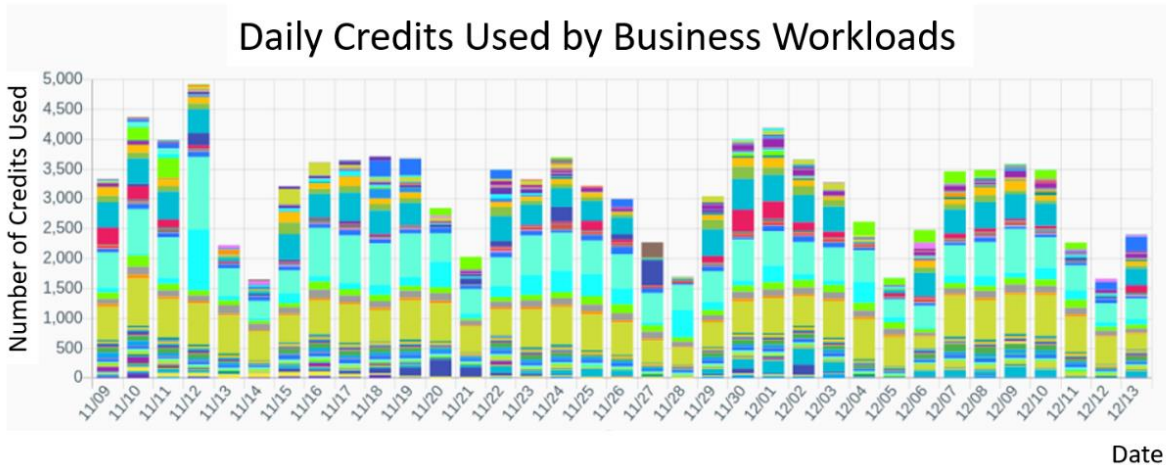
After starting the migration and moving several workloads to the cloud data platform, the client realized that migration to the cloud takes longer than they expected and they have to manage the complex environment effectively. IT leaders

were focused on maintaining performance service levels, even as new application workloads were added. In addition, the CFO and the finance team were concerned about the annual budgeting process. Without accurate and timely information, they would be blindly guessing. This was a situation they preferred to avoid and had learned from experience that BEZNext could solve for them.

BEZNext organized a continuous dynamic capacity management process for hybrid multi-cloud environments to meet SLGs for all workloads on all platforms. One of the goals is to consider the expected increase in the number of users and volume of data and planned deployment of new applications going through the DevOps process.

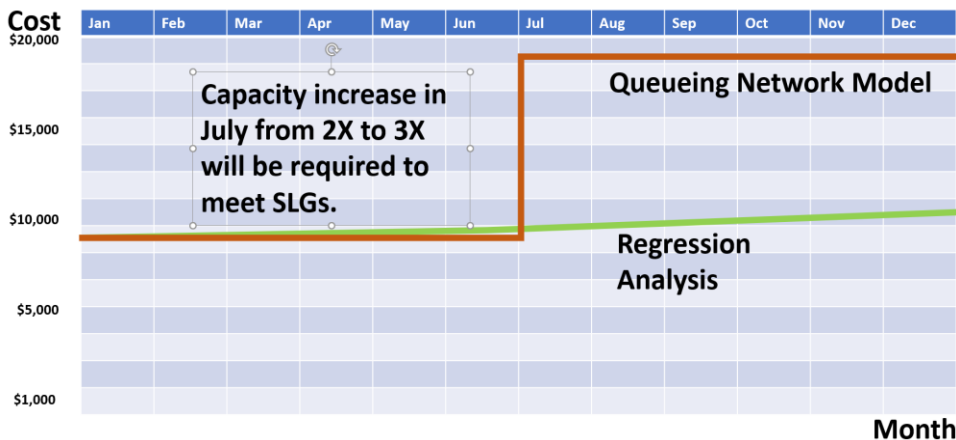
Dynamic capacity management employs BEZNext performance and finance governance software optimizing the resource allocation and predicting the cost required to meet SLGs for all workloads.

BEZNext measures the cost of running workloads as it is shown on the chart below:



The predicted budget needed for supporting SLGs for expected Workloads and Data Volumes growth

BEZNext modeling considers the expected increase in the number of users and volume of data to predict the budget required to meet SLGs for all workloads during the next 12 months. Below is a comparison of the predicted budget requirements using the regression analysis and queuing network models.




The budget based on the predicted minimum configuration needed to meet SLGs has a jump from 2XL to 3XL and the corresponding increase in cost in the middle of the year. This jump is caused by requirements to scale up to meet SLGs. It accurately predicts the cost and reduces the risk of surprises. The predicted budget based on regression analysis is inaccurate. Regression analysis does not take into consideration SLG requirements.

We also take into consideration the plan of new applications deployment. We apply our FinOps optimization technology during the DevOps process to provide value to application developers and operations before deployment of a new application:

- Application developers get information about anomalies detected during testing
- Operations get information about the cloud platforms' configuration and budget required to meet SLG for the new application.

During this project, BEZNext collected measurement data during testing of new Snowflake applications on the small virtual warehouse with the size M. BEZNext software predicted that the virtual warehouse size should be after deploying the new application from September to December increased to 2*4XL. Another increase in instances would be needed during January to August of the following year.

The minimum configuration would be 5*3XL to continue to meet the SLG for a new workload. BEZNext then automatically compares the actual performance measurement with expected and detects performance and financial anomalies for each workload.



Comparing the actual performance with predicted



Performance and Financial Anomalies detection

If actual performance, resource utilization, and financial results are significantly different from predicted, BEZNext technology detects the frequency and severity of anomalies and their root causes. Then, BEZVision performance and optimization engines determine the most effective way to fix the anomalies.

Predicted performance and financial expectations enable verification of results, anomalies, and root causes and organize continuous performance and financial governance.

The Value of BEZNext FinOps Optimization

BEZNext FinOps optimization software and services provide any organization with the following value:

- A vendor-neutral, workload-specific performance and financial governance process
- Access to a team of PhDs, experts in modeling, and cost/performance optimization
- Performance and financial anomaly detection
- A data-driven closed-loop modeling and optimization approach with feedback control
- Optimized platform selection
- Optimized migration decision making
- Continuous dynamic capacity management and expert analysis

No modern organization should feel comfortable or confident about managing its journey to the cloud without the tools to optimize and verify decisions. Likewise, critical decisions that impact large numbers of users and the company's bottom line should not be made without analysis underpinning those decisions. The value of BEZNext FinOps optimization software and services is eliminating guesswork and risk by providing financial and IT stakeholders with valuable, data-driven monitoring, forecasting, and control. If your organization wants to engage our team, please contact us at inquiry@benext.com.



Boris Zibitsker, PhD
CEO [BEZNext](http://benext.com)



Alex Lupersolsky, PhD
CTO [BEZNext](http://benext.com)

For more information about our offerings, please contact us at inquiry@benext.com, visit us at www.benext.com and [linkedin.com/benext.com](https://www.linkedin.com/company/benext.com)

ⁱ CIO Review 2022 <https://cloud.cioreview.com/vendor/2022/benext>

ⁱⁱ 2021 State of the Cloud Report - Flexera

ⁱⁱⁱ [7 Workloads Your Customers Should Move to the Public Cloud Now](#) – Gartner July 2020, ID: G00747075

^{iv} B. Zibitsker, White Paper "[Which Platform is Best for your Cloud Data Warehouse?](#)"

^vB. Zibitsker, A. Lupersolsky, P. Pratasevich, H.Khan "[Anomalies Detection and Cloud Platform Selection During DevOps](#)", CMG Impact 2021