


aws marketplace

MarkLogic

Reviews, tips, and advice from real users



Powered by  PeerSpot

Contents

Product Recap..... 3 - 4

Valuable Features..... 5 - 13

Other Solutions Considered..... 14 - 17

ROI..... 18 - 20

Use Case..... 21 - 28

Setup..... 29 - 30

Customer Service and Support..... 31 - 32

Other Advice..... 33 - 36

Trends..... 37 - 38

About PeerSpot..... 39 - 40

Product Recap



MarkLogic

MarkLogic Recap

MarkLogic offers robust capabilities for data storage and retrieval, supporting multiple formats like XML and JSON. Its built-in search and indexing facilitate rapid data querying, making it efficient for industries demanding quick data management solutions.

Boasting flexibility in data management, MarkLogic supports XML and JSON formats without strict schemas, integrating storage and search within a single platform to reduce complexity. This configuration enhances data handling, performance, and development speed. Industries like publishing, insurance, and healthcare benefit from its real-time processing, enabling tasks that range from creating PDFs to complex backend services. While users appreciate these capabilities, suggestions include interface modernization and better integration with tools like VS Code and IntelliJ.

What are MarkLogic's standout features?

- **Built-in Search and Indexing:** Facilitates rapid data querying and retrieval.
- **Multiple Data Format Support:** Offers flexibility with XML and JSON without strict schemas.
- **Integrated Storage and Search:** Combines functionalities to minimize system complexity.
- **ACID Transactions:** Enhances reliability and data integrity.

What benefits can be expected from MarkLogic?

- **Improved Efficiency:** Streamlines data handling with reduced complexity.
- **Enhanced Performance:** Supports faster development and system responsiveness.
- **Flexibility in Data Management:** Handles multiple formats effortlessly.
- **Scalable Solution:** Ensures performance across multiple nodes.

MarkLogic sees extensive use in publishing, insurance, and healthcare, where it aids in real-time processing, querying, and transformation of data. Its indexing and search capabilities allow efficient management of semi-structured data, smoothing tasks from document creation to backend solutions, without necessitating extensive migrations.

Valuable Features

Excerpts from real customer reviews on PeerSpot:



“MarkLogic has had a tangible positive impact on our organization.”



Validated user

Senior Data Engineer at a insurance company with 10,001+ employees



“Using MarkLogic has had a significant positive impact on our organization, especially in terms of performance, flexibility, and reliability.”



RituRaj

SDE 2 at Virtusa



“After moving that into MarkLogic pipelines, the same process became almost near real-time or at least reduced to minutes in most cases.”



Mampi Bhattacharya

Developer at a tech vendor with 10,001+ employees



“MarkLogic has impacted my organization positively because we have seen efficient results.”



VaibhavkawaleBhulai

Software Developer at tcs

- ✓ “We have definitely seen a return on investment, as the end customer really felt that things are going fast and they are able to get the data very fast, retrieve it, change it, transform it, and merge or discard many things they want to do in MarkLogic, saving three or four hours daily because we have enabled them to have the data in real time instead of waiting for another day.”



PramodChaudharyDarwha

Manager at TCS

- ✓ “We moved to MarkLogic and created the API using JavaScript server-side language, and we saw almost 60% improvement in the speed of the search.”



Dixit Singla

Staff Engineer at a tech vendor with 10,001+ employees

- ✓ “Compared to our previous approach, we are saving approximately two days per week, representing a forty percent reduction in time spent discussing or trying to understand data, and MarkLogic has delivered significant time savings.”



AkshayChaudhari

Data Specialist at TCS

What users had to say about valuable features:

“MarkLogic's built-in search capability is its best feature. We can utilize both CTS search and standard search functions, both of which drive faster query results. MarkLogic also includes indexes that further enhance performance by allowing us to know exactly where to search the data.

The data indexes and search capability have a significant positive impact on our work. When we perform transformations, search for data, or retrieve data, the search capability is a primary feature that enables us to efficiently locate whatever data we need. Even with ten million data records in MarkLogic, using the search capability alongside indexes makes it easy to retrieve data from even billions of records. This capability saves us considerable time when searching for and retrieving data..”

AkshayChaudhari

Data Specialist at TCS

[Read full review](#) 

“The best feature of MarkLogic is its search capability. It is really fast. It uses indexes and things are stored in document format. Whatever you want to query in the document, particular elements or properties are indexed by MarkLogic and we can then retrieve them. It is really quick in the way that we want to have it retrieved.

“This speed helps us in many places because we are retrieving the data and somebody else is trying to consume it. Other source systems are waiting for our response from MarkLogic. If the response from MarkLogic is really fast and we are able to provide a response within quick time, the other source systems will be using this response and getting the end result in a very quick fashion. This is really helpful in making the customer user experience very fast.

“I think data is really secure in MarkLogic, and we can govern the data as well..”

PramodChaudharyDarwha

Manager at TCS

[Read full review](#) 

“In my opinion, the best features MarkLogic offers are that it is very easy to use and has a very fast response time.

“The fast response time and ease of use help me in my daily work because it is really helpful since we have to run a lot of jobs at the same time and then we want to make sure everything is running as expected. It always helps us to check whether our jobs are running fine.

“MarkLogic has positively impacted my organization as I think company-wise, it is one of the go-to tools for validating the jobs we are running. It is very helpful for us to deliver our products with quality and on time.

“Using MarkLogic has resulted in specific outcomes, as we run jobs for a long time, sometimes for a couple of days, sometimes for ten hours or twelve hours. It helps us a lot..”

Islam Md

Performance Test Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“The best feature of MarkLogic is the built-in search capability. There are two search capabilities available: CTS search and search:search. Both search capabilities drive faster query results in MarkLogic, and it is really fast. We also have indexes present in MarkLogic where we can make the performance of the query faster. Indexes allow us to know exactly where to search for the data.

“The data indexes and search capability impact my work in a very positive way. When we do any transformation, when we search any data, and when we retrieve the data, the search capability is a main thing in MarkLogic that enables us to find what we need. Even if we have 10 million data entries in our MarkLogic, when we search using the search capability and utilize the indexes to make the search faster, it is really easy to get the data we need from even billions of data. This is really saving us time when we want to search and retrieve the data.

“MarkLogic has impacted my organization positively because we have seen efficient results. Instead of keeping the data in a lot of databases and MDM databases with every source system trying to keep their data in some place and then fetch it and give results, now that we have started using MarkLogic, everybody wants to fetch the data from MarkLogic because this is the main place where the data is stored, it is easy for retrieval, and it is fast. The end-to-end applications in MarkLogic architecture are really fast when it comes to performance because MarkLogic gives responses in very quick real time. We have seen improved efficiency, and a lot of time is saved because we do not have to check with a lot of teams to get an understanding of the data. We can just check with MarkLogic and its relative data in MarkLogic itself..”

VaibhavkawaleBhulai

Software Developer at tcs

[Read full review](#) 

“MarkLogic offers several powerful features. First, there is universal indexing, in which it automatically indexes all the stored XML and JSON documents. Second, it can handle both XML and JSON unstructured data in a single database, which makes it flexible for complex and evolving data requirements. The third is the ACID property. Unlike many NoSQL databases, it provides strong consistency with ACID transactions, which is critical for real-time and reliable applications. It also supports horizontal scaling and clustering, which helps in handling large volumes of data and high traffic efficiently.

“One feature that stood out to me in our project is its ability to combine search and database capability on a single platform. The tight integration of full-text search with structured query makes it very powerful for building real-time search applications without relying on any external tools. It simplifies our architecture and reduces system complexity. I also appreciate its flexibility with data models, especially handling both XML and JSON seamlessly, which can be very useful in our use cases with multiple data resources.

“Using MarkLogic has had a significant positive impact on our organization, especially in terms of performance, flexibility, and reliability. With MarkLogic's universal indexing and built-in search, we have seen query response times improve noticeably. Complex searches dropped from a few seconds to sub-second response times in many cases. Users could perform combined keyword plus filter searches in real time, directly improving our application experience. Before implementing MarkLogic, we were using a relational database and NoSQL as separate search engines, requiring Elasticsearch and others. With MarkLogic providing a single platform for both storage and search, we reduced integration overhead, maintenance efforts, and failure points. The schema flexibility for XML and JSON allowed us to onboard new data sources faster. The ACID transactions that MarkLogic provides are crucial and something rarely supported by NoSQL databases. MarkLogic improved our system by enabling faster search, reducing the response time from seconds to sub-seconds, reducing architectural complexity by combining database and search, and improving reliability through ACID transactions and clustering..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

“A few features really stand out for me, especially when working on data integration projects. One of the most valuable ones is its ability to handle both structured and unstructured data in a single platform. For example, we deal with XML, JSON, and some semi-structured data from different systems. MarkLogic allows us to store and query everything together without needing separate pipelines. That simplified our architecture quite a bit.

Another strong feature is the flexible document model. Unlike relational databases, we do not have to strictly define schemas upfront. In one project, requirements kept changing and it was dynamic. Instead of redesigning tables, we could adjust document structures and just continue with that. That saved a lot of rework. I also found the search and indexing capabilities very powerful. Even with large datasets, we could retrieve data quickly. We had use cases where we needed to search across multiple kinds of fields and documents, and MarkLogic handled that efficiently without much tuning.

One feature I personally relied on a lot is the built-in support for XQuery and JavaScript. It gave us flexibility in how we query and process data. For some transformations, XQuery worked really well, while for others we used JavaScript, depending on the complexity. Overall, what stands out is how all these capabilities are available in one platform, which reduces the need for multiple tools.

The flexible document model probably had the biggest impact on our projects. I can share a good example from a project where we were integrating customer data from multiple systems. Each system had a slightly different structure: some had extra fields, some were missing fields, and formats were not consistent at all. If we used a relational database, we would have had to keep altering schemas or creating multiple tables to handle those variations. But with MarkLogic, we could store each record as a document and gradually harmonize it into a common model. At one point, the business added new attributes to customer data midway through the project. Instead of redesigning anything, we just updated the transformation logic and started storing the new fields in the document. That change was handled pretty smoothly without impacting existing data. That flexibility saved a lot of time and avoided rework. It also made our system more adaptable to changing

requirements, which happens quite often in real projects.

One small but really useful feature is the built-in transaction and consistency handling. In one scenario, we were ingesting data in batches and there was a failure midway due to a data issue. MarkLogic handled it in a way that partial data was not left in an inconsistent state, which saved us from a lot of cleanup work. Additionally, the security and role-based access control is quite granular. We had cases where certain data needed to be restricted to specific roles. We managed that at a document level, which was really very helpful. These are not things you notice immediately, but in real-world usage, they make the system much more reliable and easier to manage.

MarkLogic has had a strong impact, especially in how we handle and use data across teams. Earlier, data was scattered across different systems, and getting a unified view required a lot of manual effort and coordination. With MarkLogic acting as a central data hub, that process became much more streamlined. Instead of pulling data from multiple sources every time, we had a consolidated dataset available, which made development and reporting faster. It reduced a lot of back and forth between teams and saved a few hours every week during development and testing stages. I think it also improved data quality. Since we had defined ingestion and transformation steps, inconsistencies were caught earlier rather than later in the pipeline. Overall, it helped improve efficiency, reduced duplication of work, and made it easier for teams to rely on consistent data.

In terms of metrics, we did see measurable improvements over time. Earlier, when we had to pull and reconcile data from multiple systems manually, it used to take a few hours for certain use cases. After moving that into MarkLogic pipelines, the same process became almost near real-time or at least reduced to minutes in most cases. In terms of development effort, we saved around 20 to 30 percent of time on data handling tasks, mainly because we did not have to build separate transformation layers or repeatedly fetch data from different sources. We also saw fewer data-related issues during testing since data was being standardized during ingestion, so inconsistency reduced. Improvements were noticeable in faster data access and also reduced manual effort and better data consistency..”

Other Solutions Considered

“We evaluated a few other options before choosing MarkLogic. We looked at traditional relational database solutions as well as a couple of NoSQL options like MongoDB. Relational databases were strong for structured data but did not handle changing schemas or mixed data formats well for our use case. MongoDB was closer to what we needed in terms of flexibility, but when it came to advanced search, indexing, and built-in support for both XML and JSON, MarkLogic had an edge..”

Mampi Bhattacharya

Developer at a tech vendor with 10,001+ employees

[Read full review](#) 

“Before finalizing MarkLogic, we evaluated a few alternatives. We looked at MongoDB, which is good for flexible document storage but required a separate search solution for advanced queries. We also considered using Oracle, which is strong and reliable but less flexible for semi-structured data. Therefore, we selected MarkLogic because it uniquely provides multi-model support along with built-in search and ACID transactions with real-time indexing..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

“Before moving to MarkLogic, we were mainly using a combination of relational databases and some custom ETL scripts for handling data integration. The challenge with that setup was that it became quite complex as the data volume and variety increased. For example, handling both XML and JSON data required separate processing logic, and maintaining those pipelines took a lot of effort. We switched to MarkLogic because it offered a more unified approach, allowing us to handle different data formats in one place and reducing the need for multiple tools and scripts. The flexibility of the document model and built-in capabilities for ingestion and querying made the overall system simpler and easier to manage compared to what we had earlier..”

Mampi Bhattacharya

Developer at a tech vendor with 10,001+ employees

[Read full review](#) 

“Before adopting MarkLogic, we were using a combination of traditional relational databases such as Oracle along with a separate search solution, such as Elasticsearch.

“The main reason for switching from Oracle and Elasticsearch to MarkLogic was simplifying our architecture by consolidating database and search into a single platform. With Oracle and Elasticsearch, we had two separate systems, and syncing between them was complex and error-prone. MarkLogic allowed us to manage these components on one platform. Given that our data was semi-structured, managing it in a relational database was tough, but MarkLogic's document model made schema evolution easier without extensive migration..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

“Before MarkLogic, I primarily relied on relational databases such as SQL servers and some Hive-based data warehouses to manage my data. The reason for switching or adding MarkLogic into the mix is that a great deal of my incoming data is semi-structured, such as XML or JSON from legacy systems, which was cumbersome to flatten into relational tables. It was slow, prone to errors, and complicated ETL pipelines.

“MarkLogic made it much easier to ingest, query, and integrate semi-structured data directly without all that extra transformation overhead. It was not a complete replacement but rather a specialized tool for a specific pain point in my data flows..”

Validated user

[Read full review](#) 

Senior Data Engineer at a insurance company with 10,001+ employees

“I evaluated multiple options before choosing MarkLogic, including alternatives such as MongoDB, Couchbase, and PostgreSQL with JSON support. While MongoDB excels in JSON handling and NoSQL flexibility, it lacks ACID compliance across complex transactions. Couchbase is good for key-value and document storage but did not fit my hierarchical XML-heavy use cases. PostgreSQL with JSON support could handle some semi-structured data, but its performance on large nested XML datasets was not great, and ETL complexity remained high.

“Ultimately, MarkLogic's native XML support, ACID compliance, and built-in search and indexing made it the best fit for my insurance project..”

Validated user

[Read full review](#) 

Senior Data Engineer at a insurance company with 10,001+ employees

ROI

Real user quotes about their ROI:

“I do not have metrics to share about return on investment as I am not the right person for this question; that is calculated by our Chief Financial Officer..”

Validated user

[Read full review](#) 

Non IT Recruiter at a computer software company with 11-50 employees

“Compared to our previous approach, we are saving approximately two days per week. This represents a forty percent reduction in time spent discussing or trying to understand data. MarkLogic has delivered significant time savings..”

AkshayChaudhari

[Read full review](#) 

Data Specialist at TCS

“We have seen a return on investment with MarkLogic. We have saved 20% to 40% of time in data retrieval. Also, the discussions have been reduced when it comes to finding out how the data behaves because the data is stored in one place now..”

VaibhavkawaleBhulai

[Read full review](#) 

Software Developer at tcs

“We saw a clear return on the investment after implementing MarkLogic in terms of saving and personnel efficiency. Since we did not need a separate database and search system, we avoided building and maintaining integrations. This led to roughly a thirty to forty percent reduction in backend development effort. With flexible schema and universal indexing, new features and data sources were onboarded faster, reducing feature delivery time by around forty to fifty percent. In terms of infrastructure and maintenance, we also achieved twenty to thirty percent savings in infrastructure and maintenance overhead..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

“We have definitely seen a return on investment. The end customer really felt that things are going fast and they are able to make more work and get more work done because MarkLogic is fast and they are able to search within the system or everywhere. They are able to get the data very fast. They are able to retrieve, change it, transform it, and merge or discard many things they want to do in MarkLogic. It is really helpful for them to have MarkLogic in the system that they use. In metrics, I think they save three or four hours now daily because we have really enabled them to have the data in real time instead of waiting for another day..”

PramodChaudharyDarwha

Manager at TCS

[Read full review](#) 

“From my perspective, there has definitely been a return on investment with MarkLogic, even if it is not always easy to quantify in exact dollar terms. For example, by using MarkLogic to handle semi-structured data directly, I have reduced ETL prep and transformation time by roughly 30 to 40 percent, freeing up engineers to focus on more value-added tasks instead of manual data cleaning.

“Additionally, faster validation and reconciliation of policy and claims data mean business teams can generate reports and insights 35 percent faster..”

Validated user

[Read full review](#) 

Senior Data Engineer at a insurance company with 10,001+ employees

Use Case

“MarkLogic serves as our enterprise-level database with multiple applications. We have numerous source systems that dump data into MarkLogic. DHF flows manage transformation, harmonization, mapping, and curation of this data, creating final records in MarkLogic. We have built REST APIs that consume data from MarkLogic and also accept data from external sources, allowing for both data modification and retrieval. We also have TDEs built in, which enable us to present existing data in MarkLogic in SQL tabular format for use by other systems like Qlik and Power BI.

Our team handles both data ingestion and API development while another team manages the UI and other tools. From our customer's perspective, MarkLogic is deployed at a manufacturing and processing plant facility. Data from different sources is ingested and stored in MarkLogic through DHF. Multiple customers interact with this data through the UI, and they can make changes or add new data by interacting through APIs..”

AkshayChaudhari

Data Specialist at TCS

[Read full review](#) 

“My main use case for MarkLogic involves running queries to check some of the jobs. I run batch jobs and then I want to check whether the batch jobs are running fine. I check the data on MarkLogic by running the query on the query logic portal.

“Regarding my main use case with MarkLogic, I find it very handy because every time I run a job, I go and run the query. I go to different databases and then see whether it's running fine. It is enjoyable working with MarkLogic.

“A recent task where MarkLogic was especially helpful involved trying to check the number of batch jobs, DES or PDM jobs, and different jobs. We always check the number and then based on the number, we compare with other tools and then see whether it's matching. It is a comparison with multiple tools. If, for example, PG Admin was not working with PostgreSQL, but MarkLogic was working fine, we were able to fix the issue on the other tools which were not working..”

Islam Md

Performance Test Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“Our main use case for MarkLogic is harmonization for a lot of data.

“We have a framework that is DHF 5 and it is based on Java. It helps us gather the data from many source systems and then put it into one system that is MarkLogic. We have control over this data coming from many sources. We do not have to wait for many source systems and be dependent on many source systems. We have accomplished this where all the important data is placed at one place and MarkLogic is that place. MarkLogic will be syncing the data every day, and it will help us keep the data in one location. My main use case is harmonization, and then data is available for all the other source systems, whichever want to use this data.

“MarkLogic has become a central piece for us now. We use APIs, TensorFlow, and many other things that MarkLogic provides. We also take help of AWS to store the data and then use MLCP to ingest it into MarkLogic..”

PramodChaudharyDarwha

Manager at TCS

[Read full review](#) 

“My main use case for MarkLogic is that we have MarkLogic as an enterprise-level database. We have many source systems that dump data into MarkLogic. Then we have DHF flows, which handle transformation, harmonization, mapping, and curation of this data and create a final record in MarkLogic. We also have REST APIs built in, so these REST APIs can consume the data that is in MarkLogic and can also send data from outside to MarkLogic. This can change the data that is already present in MarkLogic, and users can retrieve the data in MarkLogic. We also have TDEs built in. We have data that is already in MarkLogic, and the NoSQL data can be presented in the form of SQL in a tabular format using TDE, and this can be used by other systems such as Qlik and Power BI.

“MarkLogic fits into my daily workflow as we frequently deal with both data ingestion and the API. There is another team that takes care of UI and other tools. In my daily workflow, from a customer perspective, MarkLogic is used in a client who is a manufacturer and processing plant. Things are deployed on MarkLogic, and this data is stored in MarkLogic. Using DHF, the data is stored in MarkLogic from different sources. We get the data and store it in MarkLogic. Then from the UI, many customers can interact with the data that is present, and if they want to make a change or add new things into the data, they can interact using APIs..”

VaibhavkawaleBhulai

Software Developer at tcs

[Read full review](#) 

“Our main use case for MarkLogic is as a data hub for integrating and managing data from multiple sources. I mostly work on ingesting data from different systems and transforming it and storing it in a structured way so it could be used by downstream applications. For example, in one of our projects, we were pulling data from a couple of legacy systems and normalizing it into a common format using a MarkLogic pipeline. I also use it for writing queries, mainly XQuery or JavaScript, to fetch and process data based on day-to-day business needs. Sometimes, it is about building APIs on top of MarkLogic to expose that data to other services. Overall, it is a mix of data ingestion, transformation, and querying, making that data usable for applications.

When considering that use case, the biggest difference I notice with MarkLogic compared to traditional approaches or manual processes is how much it reduces the effort of handling different data formats. Before this, in some projects, we had to write separate pipelines or scripts to process XML, JSON, and other formats. This became quite fragmented. With MarkLogic, we could handle both structured and unstructured data in the same place, which simplified things a lot. For example, in one use case, we were getting data from two systems, one sending XML and another JSON. Instead of building separate transformation layers, we ingested both into MarkLogic and then harmonized them into a common model. That saved a lot of development time and also reduced complexity. Another thing that stands out is how quickly we could query and retrieve data. Since everything is stored in a flexible document model, we did not have to redesign schemas every time whenever requirements changed. That made iterations much faster compared to traditional setups. It is not always simple. There is a learning curve, especially with data modeling and writing efficient queries. Once the team gets comfortable, it becomes a very powerful platform for data integration.

One thing I did not expect initially was how useful MarkLogic is for tracking data lineage and debugging data issues. When something looked wrong in the final output, we could actually trace back how the data was ingested and transformed step-by-step. Since everything is stored as documents and transformations are applied in stages, it became easier to identify where things went wrong. Regarding one of the cases I remember, a field was getting overwritten during harmonization,

and instead of digging through multiple systems, we tracked it directly within MarkLogic and fixed it quickly. The flexibility of the data model helped a lot. When requirements kept changing, it really helped. Instead of redesigning schemas, we could adjust the document structure and update transformations, which made our workflow more adaptable. Overall, it helped with debugging and maintaining data quality, which was a nice advantage..”

Mampi Bhattacharya

Developer at a tech vendor with 10,001+ employees

[Read full review](#) 

“I used MarkLogic for a total of two years.

“When it comes to main use cases, I used MarkLogic as a backend service for handling complex structured data such as XML or JSON. I have REST API services and modules using XQuery where the system needed efficient storage, query, and transforms of large volumes of data. Additionally, I worked with real-time ingestion pipelines where data from multiple sources were processed and stored in MarkLogic, enabling real-time access and updates.

“MarkLogic is designed for multi-handle multi-model data, which means it can natively store and query XML documents, JSON documents, and unstructured and semi-structured data. Instead of normal database joins, MarkLogic works by querying inside a document efficiently using indexes. In one of my projects, we used MarkLogic to manage a large-scale document processing system, where we ingested data from multiple upstream systems in XML and JSON format, such as product or property-related data. As soon as the data was ingested, it became immediately searchable due to MarkLogic indexing. MarkLogic handles semi-structured data by storing it as a document, automatically indexing it, and allowing real-time query and updates using XQuery and strong consistency.

“My experience with MarkLogic demonstrates how we leveraged its features beyond just data storage. For example, I worked on optimizing queries written in multiple modules, mostly related to searching with text and applying structured filters, which significantly improved query accuracy and performance. Apart from basic features, I have worked on performance tuning, indexing strategies, and combining full-text search with structured query. I also used MarkLogic as both a database and search engine, which helped to simplify our architecture.

“In our use cases, MarkLogic's universal indexing and clustering have a direct impact on performance and scalability, and it has helped us significantly. In normal databases, we need to define indexes up front, and if a new query comes in, we often need a schema or indexes. In MarkLogic, all data such as XML and JSON were automatically indexed, and we did not need to pre-plan any query patterns. In real time, we had a dynamic search requirement with filters, pricing, location, and keywords, and instead of creating multiple indexes manually, we leveraged our universal index plus range index. For example, when a user searches with multiple

filters plus keywords, queries are still fast because MarkLogic uses its internal index instead of scanning documents. Regarding clustering, we have our MarkLogic clustered environment. When multiple nodes work together, horizontal scaling is part of it, as we could add more nodes if data grew, ensuring high availability. For instance, if one node failed, another would handle that traffic. During high traffic, the system stayed stable, and we handled the large data volume without performance degradation. Universal indexing helps us avoid manual indexing management while still providing fast queries for dynamic searches, and clustering allows us to scale horizontally and handle high traffic with no latency and high availability..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

Setup

The setup process involves configuring and preparing the product or service for use, which may include tasks such as installation, account creation, initial configuration, and troubleshooting any issues that may arise. Below you can find real user quotes about the setup process.

“The setup process of MarkLogic is pretty straightforward. We have to install it and there are very few configurations to be made, and it is really easy if you have admin experience..”

VaibhavkawaleBhulai

Software Developer at tcs

[Read full review](#) 

“The initial setup of MarkLogic can be difficult depending on the environment. For example, having to handle government applications, VPNs and high-level security can add to the complexity of the implementation. The MarkLogic server installation is straightforward..”

Beverly R. Jamison

IT Solutions Architect / Computer Scientist at Practical Semantics

[Read full review](#) 

“MarkLogic has really good deployment processes. If you have a big MarkLogic architecture with hundreds of nodes in the cluster, many databases, forests, and many app servers, their deployment features like ML-Gradle can automate everything. You can create a new cluster within minutes through that Gradle. That is a very good feature they have provided..”

Dixit Singla

Staff Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

Customer Service and Support

“The support resources and documentation for MarkLogic available on the website are really good, and I am satisfied. On a scale of 1 to 10, I will recommend 10 for the support..”

VaibhavkawaleBhulai

Software Developer at tcs

[Read full review](#) 

“Customer support is really good. I think they are very reachable and they provide a solution whatever we need. We mostly use their support for issues with recovering data, and they tell us how we can recover.

“I would rate customer support 10 out of 10, but we did not have much conversation. It is just one time we needed help, otherwise MarkLogic itself is good. There is a lot of documentation..”

PramodChaudharyDarwha

Manager at TCS

[Read full review](#) 

“Regarding customer support for MarkLogic, in our experience, we did not rely heavily on customer support on a day-to-day basis because once the system was set up, it was quite stable, and most issues could be handled internally. That said, there were occasions, especially early on, where we needed help around configuration and performance tuning. In those cases, the support team was helpful and guided us in the right direction, although it sometimes took back and forth to fully resolve the issue. Overall, the support is considered above average, with good responsiveness and helpful guidance when needed. We mostly relied on internal expertise and documentation and reached out for complex issues..”

Mampi Bhattacharya

Developer at a tech vendor with 10,001+ employees

[Read full review](#) 

“MarkLogic has been generally good and reliable in my experience. When I connect with them, it is very responsive. I have gone through support tickets, and proper tracking is available, so overall, it is a good support system, and I would rate it slightly higher than average.

“I would rate MarkLogic's customer support an eight due to its responsiveness, especially for higher priority issues. Support engineers demonstrate good product expertise, and the structure of the ticketing and enterprise support models work well. If someone inquires, I would suggest looking for alternatives if their team is small or they have cost constraints, but if there are no budget issues and their team is large, MarkLogic is reliable and comfortable, providing scalability..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

Other Advice

“If someone is looking into using MarkLogic, I would say MarkLogic is very helpful for providing the monitoring with detailed features. Running the query is very easy. I rate this product an eight out of ten..”

Islam Md

Performance Test Engineer at a tech vendor with 10,001+ employees

[Read full review](#) 

“MarkLogic's governance of data is excellent, and its security features are strong. However, there is very limited documentation about MarkLogic's AI capabilities. I would recommend that anyone considering purchasing MarkLogic should utilize the free format available to test it on their local machine before making a purchase decision. This review rates MarkLogic a nine out of ten..”

AkshayChaudhari

Data Specialist at TCS

[Read full review](#) 

“My experience with MarkLogic has been very positive. It is a powerful platform, especially for data-driven and search-driven applications where handling complex XML and JSON data and real-time querying is important. The combination of database and search capabilities along with strong consistency and scalability make it an excellent choice for enterprise use cases. However, there are areas such as developer experience, ecosystems, and the learning curve that could be improved to enhance accessibility. I would rate MarkLogic an eight overall..”

RituRaj

SDE 2 at Virtusa

[Read full review](#) 

“About the features, there is also graphs, but we have not used it.

“There are features in MarkLogic that are really fast and really good. It takes care of all of the things. It is an enterprise-level database which is helping us eliminate the need of purchasing other tools. It is really good.

“My advice for others would be to try to purchase it first. They can try to learn and get experience on how to work with MarkLogic and then they can buy.

“I would rate MarkLogic nine out of 10..”

PramodChaudharyDarwha

Manager at TCS

[Read full review](#) 

“I suggest starting with a clear understanding of your data use case before adopting MarkLogic, especially if you are dealing with multiple data formats or large volumes of data. That is where it really adds value. From our experience, it is

important to invest time early in data modeling and indexing. We initially underestimated that, leading to performance issues later. Once we optimized indexes properly, things improved a lot. Also, make sure your team gets comfortable with the basics of XQuery or JavaScript in MarkLogic. There is a learning curve, but once you get past it, development becomes much smoother. I recommend starting with a smaller use case or pilot project, which helped us understand how to structure data and pipelines before scaling it across the system. It is a powerful platform, but you get the best results when you plan the foundation properly and gradually build on it.

Regarding additional thoughts on MarkLogic, one additional thought I would add is how MarkLogic fits as a long-term solution. From my experience, it is not just a tool you use for a short-term project. Once it is set up properly, it becomes a core part of your data architecture. More teams gradually started relying on it once they saw the value of having centralized and consistent data. It is important to have the right expertise in the team. It is a powerful platform, but to get the best out of it, you need people who understand how to design data models and optimize queries properly. It is a strong platform for enterprise use when you treat it as a long-term investment rather than just a quick solution.

My overall rating for MarkLogic is 8 out of 10..”

Mampi Bhattacharya

Developer at a tech vendor with 10,001+ employees

[Read full review](#) 

“In comparison, for about one week, we are almost saving two days now. So two of five means we are saving 40% of time compared to discussing or trying to understand the data. So 40% of time is reduced.

“The learning curve for new users or developers on MarkLogic is a bit tough at first, but afterwards it is really easy to understand for developers new to MarkLogic.

“MarkLogic is very flexible when it comes to integrating with other tools or systems because they have a lot of authentication systems and connection systems available. We also have a lot of app servers with HTTP auth, OAuth 2.0, and basic digest authentication. MarkLogic is really flexible when it comes to integrating with other tools or systems.

“MarkLogic is very good at handling larger data sets. As we already have billions of data, it is really handling it well. We just need to take care of the CPU, RAM, and storage.

“MarkLogic handles heavy loads at peak times and performs well. It automatically manages itself with the configurations we can have, parallel processing, and load balancers in place. This takes care of the performance during heavy loads at peak times.

“MarkLogic is also available in a free format, so you can deploy MarkLogic into your own systems and try it out to see how it works, how fast it is, and how it can solve your own data problems. It would be beneficial if you give it a try before purchasing.

“Regarding MarkLogic's governance and security, I think data governance is really good. We can keep the data private to some roles, and we have user-based access control. We can give restrictions on the data and govern the data. Security-wise, MarkLogic is really good because without proper credentials, MarkLogic does not have any ways of getting hacked.


“I would rate this review a 9 out of 10..”

VaibhavkawaleBhulai
Software Developer at tcs

[Read full review](#) 


Top Industries

by visitors reading reviews




Company Size

by reviewers



by visitors reading reviews



Large Enterprise Midsized Enterprise Small Business

About this buyer's guide

Thanks for downloading this PeerSpot report.

The summaries, overviews and recaps in this report are all based on real user feedback and reviews collected by PeerSpot's team. Every reviewer on PeerSpot has been authenticated with our triple authentication process. This is done to ensure that every review provided is an unbiased review from a real user.

Get a custom version of this report... Personalized for you!

Please note that this is a generic report based on reviews and opinions from the collective PeerSpot community. We offer a [customized report](#) of solutions recommended for you based on:

- Your industry
- Company size
- Which solutions you're already considering

The customized report will include recommendations for you based on what other people like you are using and researching.

Answer a few questions in our short wizard to get your customized report.

[Get your personalized report here](#)

About PeerSpot

PeerSpot is the leading review site for cloud, AI, and business software. We created PeerSpot to provide a trusted platform to share information about software, applications, and services. Since 2012, over 22 million people have used PeerSpot to choose the right software for their business.

PeerSpot helps tech professionals by providing:

- A list of products recommended by real users
- In-depth reviews, including pros and cons
- Specific information to help you choose the best vendor for your needs

Use PeerSpot to:

- Read and post reviews of products
- Access over 30,000 buyer's guides and comparison reports
- Request or share information about functionality, quality, and pricing

Join PeerSpot to connect with peers to help you:

- Get immediate answers to questions
- Validate vendor claims
- Exchange tips for getting the best deals with vendor

Visit PeerSpot: www.peerspot.com

PeerSpot

244 5th Avenue, Suite R-230 • New York, NY 10001

reports@peerspot.com

+1 646.328.1944