aws marketplace

MarkLogic

Reviews, tips, and advice from real users






Powered by PeerSpot

Contents

Product Recap . . . . . 3 - 4

Valuable Features . . . . . 5 - 14

Other Solutions Considered . . . . . 15 - 18

ROI . . . . . 19 - 20

Use Case . . . . . 21 - 28

Setup . . . . . 29 - 30

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

Other Advice . . . . . 33 - 35

Trends . . . . . 36 - 37

About PeerSpot . . . . . 38 - 39

Product Recap

MarkLogic

MarkLogic Recap

MarkLogic is an operational and transactional Enterprise NoSQL database that is designed to integrate, store, manage and search SON and XML documents and semantic data (RDF Triples).

Valuable Features

Excerpts from real customer reviews on PeerSpot:

- "MarkLogic has had a tangible positive impact on our organization."

Verified 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 NSIT

SDE 2 at Virtusa

- "Overall, it reduced data transformation efforts, simplified architecture, and made it easier to build richer and more connected database models."

Varuns Ug

Senior software developer at Makemytrip

"For example, earlier search operations were slow and less flexible, but after using MarkLogic, we delivered near real-time results, improving both system efficiency and user satisfaction."

Ravi Raushan Kumar

Software Engineer at Netaji Subhash Engineering College

"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

"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."

Islam Md

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

"MarkLogic has positively impacted my organization by making our job easier because we can store a large amount of data, and the built-in search feature is great, including semantic data management."

Verified user

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

What users had to say about valuable features:

"The best features MarkLogic offers include multi-model flexibility, built-in search, data hub platform integration, and semantic data management. I personally appreciate the built-in search feature because it indexes all data immediately upon ingestion for rapid searching, so we can perform full-text, phrase, or geospatial searches.

"MarkLogic has positively impacted my organization by making our job easier, although we have yet to notice the full details. It made my job easier because we can store a large number of data, and the built-in search feature is excellent for semantic data management.."

Verified user

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

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

"In my experience, the best features of MarkLogic include its native support for XML and JSON, which makes working with hierarchical or semi-structured data easier than flattening it into relational tables. Additionally, its flexible schema and indexing capabilities allow me to index anything, including nested elements, which speeds up queries and reduces the need for custom code. Other significant aspects include built-in search capabilities that allow for full-text searches and complex queries directly on documents, and data integration and harmonization, which combines multiple sources into one logical view, simplifying manual processes.

"Finally, ACID transactions on semi-structured data give me confidence that updates, merges, and data integrity remain intact even in large or complex datasets.

"The flexible schema, indexing, and search capabilities of MarkLogic are incredibly useful for me. For example, the flexible schema and indexing mean I can load all the XML or JSON policy and claims data without having to predefine rigid table structures, saving considerable time during ingestion, especially since different legacy systems may have slightly different formats. I can simply map the fields I need and let MarkLogic handle the rest. The search capabilities are also very helpful; I can run queries across nested elements to find specific claims, policies, or attachments quickly. Previously, doing this in a relational database would require multiple joins and a great deal of transformation logic.

"MarkLogic has had a tangible positive impact on our organization. Thanks to its flexible schema and indexing capabilities, we ingest semi-structured and hierarchical data from multiple legacy systems 30 to 40 percent faster than before, which leads to faster ETL cycles and quicker delivery of analytics-ready datasets to downstream systems. Using MarkLogic's search and query capabilities, I also reduced the time to reconcile and validate policy and claims data across systems by approximately 35 percent, helping business teams gain insights much faster. This highlights the positive impact MarkLogic has made in our organization.."

Verified user

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

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 NSIT
SDE 2 at Virtusa

Read full review

“The best features of MarkLogic are its powerful search capabilities, flexible schema, built-in indexing, and high performance for XML or JSON data. MarkLogic provides a Google search-like capability, including full-text search, partial matching, and relevance scoring. Another feature is schema flexibility; since it is a document-based database, we can store XML or JSON without a strict schema constraint, which makes it easy to evolve data structures.

The third feature is built-in indexing, as MarkLogic automatically maintains indexes, and we can configure the range indexes to specifically improve query performance. MarkLogic's XQuery support, which is native to the platform, allows efficient querying and transformation of XML data, while it even supports ACID properties. Unlike other NoSQL options, MarkLogic supports full ACID property compliance, ensuring data integrity and consistency.

MarkLogic's built-in indexing allows queries to run directly on indexes instead of scanning documents, which significantly improves performance. MarkLogic uses a universal index that automatically indexes all the content in the database, both structured and text, without requiring manual indexing as a traditional database would. We can configure range indexes for specific fields, such as policy number or customer name, allowing faster filtering and sorting of the results. In my workflow, this has helped tremendously because queries execute directly instead of scanning the XML data. Search performance improved significantly for partial and filtered searches, and it also reduced response time for user queries, even with a large database.

MarkLogic has improved our system performance, enabled flexibility in data handling, and specifically enhanced search efficiency. It improved search performance, provided flexibility in data modeling, and since it supports XML and JSON without a strict schema, we could easily adapt to changes in business requirements without any major database alterations. It even reduced development effort, as features such as built-in indexing and search reduce the need for external search systems, simplifying our architecture. It resulted in a better user experience with faster query responses and flexible searches. For example, earlier search operations were slow and less flexible, but after using

MarkLogic, we delivered near real-time results, improving both system efficiency and user satisfaction.."

Ravi Raushan Kumar
Software Engineer at Netaji Subhash Engineering College

Read full review

"MarkLogic for the hotel listing and search use case compared to other approaches made things easier than a traditional approach. If I compare it with using MySQL and Elasticsearch, typically you would need to maintain two systems: one for transaction storage and another for search. That introduces challenges regarding data synchronization, consistency, and operational overhead. With MarkLogic, since it natively supports both document storage and advanced search, I could avoid the dual-system complexity. It simplified the architecture because indexing and querying are built-in and tightly integrated. In terms of performance for semi-structured data and search-heavy queries, it was quite efficient because indexes are created automatically and queries are optimized around them. The schema flexibility was a significant advantage and it also helped reduce system complexity, improve development speed, and handle search use cases efficiently.

"One thing I found particularly useful was that MarkLogic handles indexing by default, unlike a traditional system where you have to explicitly define and manage indexes. MarkLogic automatically indexes documents, which made it easier to get started and integrate quickly. Another advantage is that it can handle both structured and unstructured data together, which is very useful in real-world scenarios where travel data has a mix of fixed fields and dynamic attributes. The fact that it supports flexible querying over nested data without needing complex joins made development simpler and queries more intuitive.

"Many features offered by MarkLogic are valuable. One of the standout features is its multi-modal capability. It can handle JSON, XML, and RDF data in a single database. That is particularly useful for applications dealing with diverse and evolving data formats. Another feature is the built-in search and indexing capability, along with schema flexibility since it is document-based. It handles semi-structured and nested data very naturally, which reduces the overhead of schema migration. An important feature is ACID transactions with NoSQL flexibility, so you can get reliability from a traditional database along with the scalability and flexibility of NoSQL.

"MarkLogic's multi-modal capabilities make things easier in scenarios where JSON can be used for application-facing data such as hotel details, XML comes into play with external APIs, and RDF can be used for representing relationships.

Instead of converting everything into one rigid format, MarkLogic allows you to store each in its native form and still query and access them. This opens up possibilities such as combining search data with relationships and searching in a single query, which would otherwise require multiple systems or a complex data pipeline. Overall, it reduces data transformation efforts, simplifies architecture, and makes it easier to build richer and more connected database models.

“One thing that surprised me about MarkLogic is how it has so many built-in traditional capabilities. Features such as search, indexing, and even data integration are natively available, so you do not have to rely on multiple external systems. That was unexpectedly useful because it simplified the overall architecture significantly. Another interesting aspect was its flexible querying over deeply nested data. In traditional databases, handling nested or hierarchical data often requires complex joins. I found it interesting to design data platforms rather than just a database, especially with complex capabilities around search and integration.

“MarkLogic has impacted my organization positively. Since my exposure to MarkLogic has been more on the exploration and evaluation side, I did not see full production-scale impact. However, even in the use case I worked on, a few clear benefits stood out. One was simplifying the architecture instead of thinking in terms of separate systems such as MySQL for storage and Elasticsearch for search. MarkLogic allowed both in a single system, which reduced integration overhead and potential consistency issues. Another benefit was fast deployment and integration because of its schema flexibility and automatic indexing. It was easier to onboard new data fields and quickly test different query patterns without heavy schema changes. I noticed that the search-heavy queries on semi-structured data performed quite well, which really helped in reducing system complexity and speeding up development for search-heavy, semi-structured data use cases..”

Varuns Ug
Senior software developer at Makemytrip

Read full review

Other Solutions Considered

"Before MarkLogic, I used a combination of MySQL for storage and sometimes Elasticsearch for search-heavy use cases. While exploring MarkLogic, the approach was not an immediate switch in production but more of an evaluation to see how it compared to the traditional approach of using MySQL and Elasticsearch.."

Varuns Ug
Senior software developer at Makemytrip

Read full review

"Before choosing MarkLogic, I explored some alternatives, primarily comparing it with the combination of MySQL and Elasticsearch, and I also considered MongoDB since it provides document-based storage and schema flexibility. The key difference I found was that while MongoDB handles flexible data effectively, it does not offer the same level of integrated search capability as MarkLogic.."

Varuns Ug
Senior software developer at Makemytrip

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 NSIT
SDE 2 at Virtusa

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 NSIT
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.."

Verified user

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

Read full review

"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.."

Verified user

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

Read full review

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.."

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

Read full review

"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 NSIT
SDE 2 at Virtusa

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.."

Verified user

Read full review

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

"Since my experience with MarkLogic is more focused on exploration, I have not seen production-level ROI metrics such as cost and team size reduction. However, even during the evaluation, I could see potential in reduced development effort. The performance of search and filtering queries on semi-structured data felt more efficient compared to a traditional approach using MySQL and Elasticsearch. While I cannot quote exact numbers since it was not in production, it definitely showed potential for reduced development time, simplified architecture, and fast search use cases. Ultimately, it reduced development complexity and effort noticeably, especially by eliminating the need to manage multiple systems.

"Since my experience is more focused on exploration, I have not seen production-level ROI metrics such as cost or team size reduction. However, even during evaluation, I could gauge potential ROI in terms of what it generates and the additional benefits it provides.."

Varuns Ug

Read full review

Senior software developer at Makemytrip

Use Case

“My main use case for MarkLogic is for storing data. A quick specific example of the kind of data I store in MarkLogic is that we have a data warehouse, and we use it as a NoSQL database to store, manage and search complex heterogeneous data..”

Verified user

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

Read full review

"I use MarkLogic for performing many operations such as handling semi-structured data and building search-driven use cases. For example, I examined how it can handle and store XML documents and leverage its powerful indexing and search capabilities for fast querying. It is particularly useful in scenarios where you need flexible schemas along with advanced search capabilities including filtering, full-text search, and aggregations.

"A task where MarkLogic was central to my work was a travel search and filtering system, similar to a hotel listing use case. In this project, the data was semi-structured, with elements such as hotel details, amenities, and pricing that vary significantly across different entities. I used MarkLogic to store this data as JSON documents. What made MarkLogic central was its built-in indexing and search capabilities. I configured indexing on fields such as location, price range, and amenities and then used its query capabilities to perform fast filtering and full-text search. Instead of relying on a separate search system, MarkLogic handled both data storage and search efficiently, which simplified my architecture and improved query performance.."

Varuns Ug
Senior software developer at Makemytrip

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

“My main use case for MarkLogic is for a specific client where I handle XML data in the backend. I use XQuery, and we utilize MarkLogic mostly for querying the data in the backend.

In one of my use cases, we stored customer policy records as an XML document in MarkLogic. Each document contained details such as policy number, customer name, and other metadata. A common requirement was that users often did not remember the full policy number, so they searched using partial input, such as part of the policy number or their name. I implemented XQuery-based search logic along with proper indexing, so MarkLogic could efficiently return the matching document even with incomplete input. This significantly improved search performance and user experience.

Apart from the search cases, one key advantage we leveraged with MarkLogic's schema flexibility was that since our data was stored in XML, we could easily accommodate changes in the structure without major migration. I also worked on optimizing query performance by configuring indexes properly, which reduced query response time significantly. Additionally, we used MarkLogic as a central data store, integrated with the backend service through APIs, ensuring fast and reliable data access. We also ensured that the queries were written efficiently and aligned with the index configurations to avoid full document scans, which is critical for performance in MarkLogic..”

Ravi Raushan Kumar
Software Engineer at Netaji Subhash Engineering College

Read full review

“MarkLogic's primary use case in my experience is handling semi-structured and hierarchical data that does not fit neatly into traditional relational tables. In my particular project, I work with NoSQL data, meaning I handle semi-structured and hierarchical data. For example, in one of my insurance projects at ValueMomentum, I worked on an initiative where I had policy and claims data in XML format coming from different legacy systems. I used MarkLogic to ingest and normalize the data and integrate it with our Hive data warehouse for reporting and analytics. A specific example I would share is that I created a 360-degree view of customer policies and claims using MarkLogic, which allowed me to merge XML claim documents with relational customer information, perform queries across different nodes, and feed clean aggregated data into the ETL pipeline for downstream analytics.

“This significantly saved time compared to flattening everything manually into SQL tables.

“Working with XML and integrating that with my ETL processes is quite interesting, as MarkLogic makes it far easier to handle hierarchical data instead of attempting to force it all into relational tables. The built-in support for XML and JSON, combined with the indexing and searching capabilities, allows me to query deeply nested structures without writing extensive custom parsing code. The challenging part is mostly around integration with our existing ETL pipelines, ensuring that transformed data flows correctly into Hive and matches the relational schemas. Sometimes I have to be careful with data types and schema evaluation because MarkLogic is schema-flexible, but the downstream systems are strict. Overall, though, it speeds up handling complex semi-structured data and reduces manual transformation work significantly. I would say it is easier than many other XML handling approaches I have tried.

“I would add that I really appreciate how MarkLogic handles hierarchical relationships naturally, especially in insurance data where nested information such as policies contain multiple coverages, each with different claims and documents. This aspect of the insurance domain is really cumbersome, and MarkLogic allows me to query across these nested structures directly without having to flatten everything at the beginning. Its search and indexing features

make it easier to identify anomalies or missing information in the semi-structured data before it reaches the ETL pipelines, saving considerable time during validation and reconciliation. Overall, I find it a very practical tool for bridging legacy data formats with modern data warehouses..”

Verified user

Senior Data Engineer at a insurance company 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 NSIT
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.

"MarkLogic is deployed as a cluster environment on the server, enabling scalability, high availability, and fault tolerance. My organization has MarkLogic deployed in a cluster setup, which helps with scalability and high availability.."

Ravi Raushan Kumar

Read full review

Software Engineer at Netaji Subhash Engineering College

"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

Read full review

IT Solutions Architect / Computer Scientist at Practical Semantics

"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

"Customer support for MarkLogic provides strong enterprise-level assistance through direct interactions. It is usually handled by a specific team that is very responsive to a variety of issues we encounter.."

Ravi Raushan Kumar
Software Engineer at Netaji Subhash Engineering College

Read full review

"I have faced some situations where I needed help. While I have not interacted directly with MarkLogic support in a production environment, my understanding is based on industry feedback, which suggests it has enterprise-grade support, including ticketing systems and dedicated support channels for customers.."

Varuns Ug
Senior software developer at Makemytrip

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 NSIT
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 is great, and my advice for others looking into using MarkLogic is that it is a great database with awesome features that you should consider. I would rate this product a nine out of ten..”

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

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 NSIT
SDE 2 at Virtusa

Read full review

"MarkLogic's built-in indexing allows queries to run directly on indexes instead of scanning documents, which significantly improves performance. MarkLogic uses a universal index that automatically indexes all the content in the database, both structured and text, without requiring manual indexing as a traditional database would. We can configure range indexes for specific fields, such as policy number or customer name, allowing faster filtering and sorting of the results. In my workflow, this has helped tremendously because queries execute directly instead of scanning the XML data. Search performance improved significantly for partial and filtered searches, and it also reduced response time for user queries, even with a large database.."

Ravi Raushan Kumar
Software Engineer at Netaji Subhash Engineering College

Read full review

"I would suggest first clearly evaluating whether your use case truly benefits from MarkLogic's strengths. It works particularly well for search-heavy and semi-

structured data use cases where flexible and powerful querying is needed. At the same time, I would recommend comparing it with alternatives such as MySQL, MongoDB, and Elasticsearch for trade-offs. Additionally, it is important to plan for the learning curve, especially around concepts such as indexing and querying.

“Overall, I think MarkLogic is a very powerful platform, especially when involving semi-structured data and advanced searches. I would rate this review an 8 out of 10..”

Varuns Ug
Senior software developer at Makemytrip

Read full review

“MarkLogic is solid for its main features such as handling XML, flexible indexing, search, and ACID compliance on semi-structured data. However, the learning curve is steep, integration with other systems can be tricky, and the community and tooling are not as extensive as mainstream platforms. It is strong, but I would not say it is perfect.

“My advice for others looking into using MarkLogic is three-fold: First, understand your data upfront. If you have a great deal of semi-structured or hierarchical data such as XML or JSON, MarkLogic can save you a great deal of time, but you need to understand your indexing and querying strategies. Second, plan your cluster and indexing strategies carefully as it pays off in performance and scalability later. Third, take full advantage of built-in search and ACID features, which are real-time savers for validation, reconciliation, and governance.

“Additionally, do not underestimate the importance of training and onboarding; there is a learning curve, so make sure your team becomes comfortable with queries, transforms, and integrations before relying on it in production. If used correctly, it is a very strong tool for bridging legacy semi-structured data with modern analytics pipelines. I would rate MarkLogic at a seven on a scale of one to ten..”

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

Read full review

Top Industries

by visitors reading reviews

Educational Organization	29%
Financial Services Firm	14%
Recreational Facilities/Services Company	7%
Manufacturing Company	7%

Company Size

by reviewers

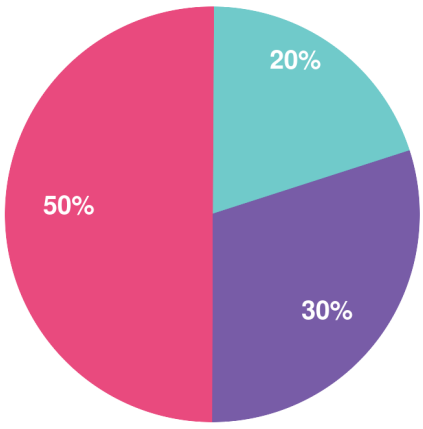

by visitors reading reviews

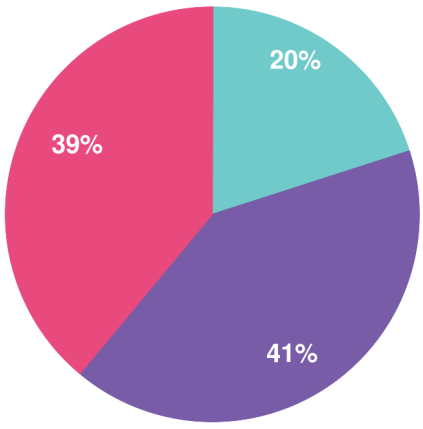

Large Enterprise Midsize Enterprise Small Business

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 software running on AWS and other platforms. 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