

MarkLogic Query Service

Every company wants to leverage the cloud for the agility it provides. Running in the cloud can also provide significant cost savings. But, not always. The reason is that when running large workloads in the cloud, they are not predictable. You have to either over-provision query capacity to handle spikes in demand, or provision for the average and risk disappointing customers during those black Friday events.

The MarkLogic® Query Service provides immediate query capacity to handle spikes in demand. And, when demand slows, that extra capacity is removed. Because it is a service, you just have to define what capacity to start with and what spikes will be expected. The service keeps a close watch on your system and adds or removes capacity as needed. It is automatic, all provisioned, and all handled by the service running in AWS.



Key Benefits

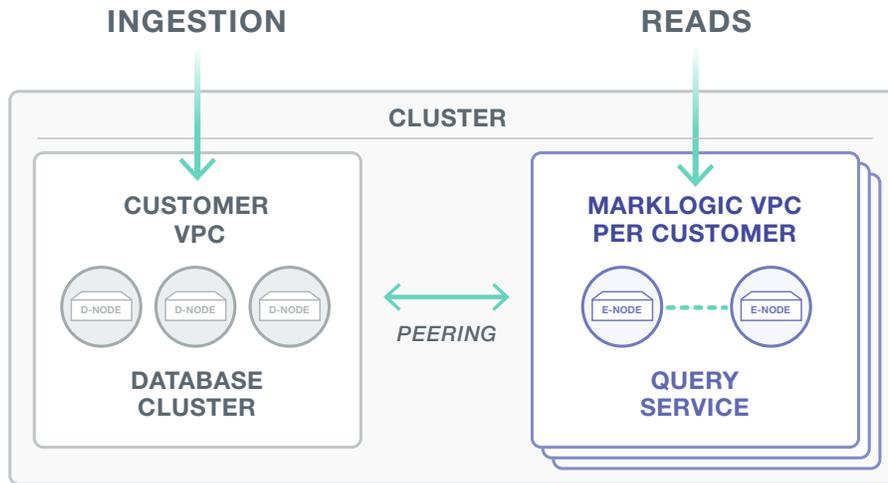
- **Provides Immediate Capacity** – The service adds capacity automatically to handle spikes in demand, made possible by MarkLogic’s unique architecture that avoids the typical hassle required by other databases.
- **Cuts cloud costs** – Keep your query capacity in lock-step with demand, and never pay for more than required. MarkLogic is more cost-effective than other databases (even the one offered by the cloud provider itself).
- **Keeps peace of mind** – It is brought to you by the same company that builds the MarkLogic database, and is built and managed with the same enterprise-grade quality and reliability trusted by thousands of large enterprises.

Use Cases

- **Applications that have hard to predict query demand** – For example, when releasing a novel API to end-users with significant adoption, or there is an application with major spikes in demand (e.g., promotions, social-media events)
- **Applications that have known demand that varies throughout the day** – For example, publishing important, highly consumed content every morning at 10:00 AM or a customer 360 app that is heavily accessed every afternoon
- **Operational Data Hubs with multiple applications that should be isolated** – For example, a reporting application that should be isolated from a transactional application for performance reasons



MarkLogic's Unique Architecture



Scaling query capacity in most databases is challenging because it requires migrating data, re-partitioning, re-balancing, and more. Their architectures are not designed to handle spikes in demand.

The MarkLogic Query Service takes advantage of MarkLogic's architecture to handle the spikes in demand by relying on the unique separation between MarkLogic's E- and D-Nodes. MarkLogic's E-Nodes (Evaluator Nodes) process queries and D-Nodes (Data Manager Nodes) handle data storage and retrieval (including things like journaling, data recovery, and backup).

This unique separation makes it is easy to add or remove E-Nodes as needed for additional query capacity, without having to worry about provisioning D-Nodes. The MarkLogic Query Service automatically adds and removes E-Nodes in the cloud to handle reads while the D-Nodes continue to reside in the AWS Virtual Private Cloud that you manage, providing a separation that is connected through peering within the cluster.

About MarkLogic

MarkLogic is the world's best database for integrating data from silos, providing an operational and transactional Enterprise NoSQL database platform that integrates data better, faster, with less cost. Visit www.marklogic.com for more information.

© 2018 MARKLOGIC CORPORATION. ALL RIGHTS RESERVED. This technology is protected by U.S. Patent No. 7,127,469B2, U.S. Patent No. 7,171,404B2, U.S. Patent No. 7,756,858 B2, and U.S. Patent No 7,962,474 B2. MarkLogic is a trademark or registered trademark of MarkLogic Corporation in the United States and/or other countries. All other trademarks mentioned are the property of their respective owners.

MARKLOGIC CORPORATION
999 Skyway Road, Suite 200 San Carlos, CA 94070
+1 650 655 2300 | +1 877 992 8885 | www.marklogic.com | sales@marklogic.com