WHITEPAPER

How fast storage built for modern databases can drive ecommerce success

The Register® | lightbits

Ecommerce operators are on track to take almost half of retail sales worldwide by 2027, as shoppers continue to be entranced by seamless browsing and purchasing of a virtually unlimited product selection online.

That 'oh so easy' shopping experience is all enabled by a complex ecosystem of new, emerging databases and analytic and AI systems - which in turn rely on powerful and resilient architecture. This has traditionally meant that complex database architecture was tightly coupled to Direct-Attached Storage (DAS).

This white paper details the challenges that this approach represents as ecommerce providers continually scale upwards, and how their complexity can potentially compromise not only that critical user experience, but also the overall operational resilience of ecommerce operators.

The paper also explains how NVMe/TCP - an open-source storage standard developed by Lightbits - can cut through this complexity by slashing latency, boosting performance, and improving resilience both in the datacenter and in the cloud.

Why ecommerce is a question of experience and scale

It can be hard to grasp just how large and fast-growing the global ecommerce market is. According to figures cited by the US' International Trade Commission, global B2C ecommerce stood at $1.45 trillion in 2017, but will hit $4.2 trillion in 2024, and $5.6 trillion in 2017.

And while ecommerce will show a CAGR of more than 11 percent in the US and the rest of the world, the largest growth will come in India, with a CAGR of 14.1 percent, followed by Brazil, Argentina, and Turkey.

According to the Boston Consulting Group, by 2027, ecommerce could account for 41 percent of global retail sales, up from 18 percent in 2017 - with the US and Asia setting the pace.

Maintaining this growth presents a massive technical and engineering challenge that will only become more complex.

The platform and infrastructure engineers within ecommerce operators must ensure that millions of users can browse a vast product range and navigate detailed information and imagery of products at will, seamlessly, instantly, and simultaneously.

Those users also benefit from guidance, in the shape of tailored recommendations and offers, informed by the experiences of other customers. Increasingly, this is driven by cutting-edge AI techniques, requiring ever - more sophisticated data management.

And it's a given that purchases are handled instantly and securely, while the consumer's personal and other information is protected - raising the stakes even further.

Ecommerce operators will also be serving up advertisements or offers on behalf of key sponsors or partners. This again requires the analysis of vast amounts of data and integration with real-time bidding operations for ad slots.

Operators need to keep track of not only what they are selling, but what they're not selling too, tweaking their offers

and forecasts accordingly. Simple inventory replenishment means updating stock counts, adding new imagery and other information to the online store, and keeping track of physical stock. Launching new lines, or even developing own label products, means detailed analysis of customer preferences and forecasting.

In addition to addressing all of these challenges, infrastructure engineers need to take regional and local factors into account. This could be in the form of consumers' preference for wireless communications in some regions, or data protection and confidentiality regulations that vary significantly from country to country.

The consequences of falling down on any of these challenges are stark – disrupted customer journeys, items languishing in baskets, uncompleted purchases, and ultimately dwindling sales.

Keeping up means a massive commitment to cutting-edge recommender algorithms and analysis systems, as well as to the robust, scalable architecture needed to run them. But it also means storing and processing vast amounts of data. And that is increasingly becoming the weakest link.

How the databases used in eCommerce have evolved

The real-time, seamless customer experiences at the heart of ecommerce rely on the ability to load, store, manipulate, and analyze vast amounts of data.

The product catalogs, orders, and payments databases are only the start. These datastores will be called upon for various operations throughout the sales process. User data might need to be accessed for fraud analysis, for example, while an operator might run its own database for security purposes.

Elsewhere, cutting-edge recommender systems will pull together information from the product database and the user database, before the analytics system does its work. The wide range of users and applications means that ecommerce platforms often call on a wide range of database formats.

Traditional relational databases

Relational databases are at the heart of ecommerce operations, handling critical structured data such as product information, user details, or orders. This could be something like Oracle, though ecommerce operations will often turn to platforms like MySQL or PostgreSQL as well.

Document databases

Recommendation systems and other analytical processes usually rely on less structured data, meaning that they will look to NoSQL databases including MongoDB and Couchbase. Because they support more flexible schema, these document databases are used for managing the wide range of content needed in ecommerce - including images, video, and review content. But they are also used for crucial business processes, such as managing orders and payment information.

Key value stores

Key value stores - such as Cassandra - are typically used for session management and shipping management, and can be harnessed to accelerate the serving up of product data. They can play a central role in delivering the personalization that is part and parcel of the user experience in ecommerce, storing user histories and fuelling recommendation applications. They are also used for storing shopping cart data, even when users leave a site - a critical feature for any ecommerce operator. And their caching capabilities also lighten the load on other datastores, such as the product database.

Vector databases

Vector databases - such as Pinecone, Milvus, and Postgres with pgvector - are used to store mathematical representations of raw data, or high dimensional vectors. That raw data can include unstructured types such as documents, or images and audio. Vector databases often underpin the AI and machine learning applications that support recommendation systems, as they make it easier for a model to identify items with similar attributes or identify relationships between items.

Other database types

In-memory systems are used to reduce latency, and speed personalization or credit decisions. Fraud detection can also rely on in-memory systems. Time-series systems - such as Kafka or Prometheus - are used for monitoring system performance. Downtime will be extremely expensive – as well as tracking trends over time - in terms of sales or user behaviors.

Each of these databases presents its own challenges. But it's usually platform engineers and site reliability engineers (SREs) who must figure out how to support them all.

The storage challenge

While each datastore will have expert developers or database administrators that know how to optimize performance and speed for their respective application, it's the platform engineers who have the task of architecting, optimizing, and managing the underlying infrastructure that ensures this database ecosystem, including storage. That's a formidable challenge, both in terms of cost and in making sure everything runs quickly and efficiently without degrading or disrupting the end user experience.

High-performance storage, in the form of solid state disks (SSDs), is expensive in itself. But legacy approaches to provisioning storage compound issues. Traditionally, each database would have had its own DAS resources. Given that each database is essential to overall performance, this tends to lead to capacity over-provisioning and under-utilization – which can be extremely costly at the petabyte scale large ecommerce firms tend to operate at.

Infrastructure architects looking to increase utilization and scalability might look to a distributed, software-defined storage solution to address these problems. However, this can present its own issues. Platforms such as Ceph can introduce complexities, while provisioning the required number of nodes' massive installations compounds complexity and cost challenges. Likewise, the breadth and diversity of the systems required can mean latency issues become more apparent as installations are scaled up.

The decision of on-premises infrastructure vs public cloud

Some ecommerce providers choose on-premises datacenters to the cloud, largely because they offer more control over the infrastructure and its costs. But this can raise additional issues, not the least of which is the challenge to keep electricity consumption within acceptable limits, even as compute – and storage – becomes more dense and power-hungry.

Many subsequently harness the cloud instead, either to supplement or replace on-premises capacity. This is a hybrid approach that comes with its own complexity, whether companies are going completely cloud native, or looking to burst into the cloud during sales events or other periods of high demand.

It's a situation that often leaves ecommerce companies trying to solve the dilemma of using cloud providers' standard instances which might not offer the optimal balance of storage and compute power, leading to further overprovisioning and cost. Alternatively, they might feel forced to opt for cloud providers' premium storage offerings, once again leading to increased costs.

The cloud can cause further complications when it comes to quelling data sovereignty and privacy concerns and handling punishing ingress/egress fees. These are sometimes incurred when moving large amounts of data between different on and off premises systems.

The need for 24/7/365 availability and resiliency

On top of this, platform engineers face all of the usual challenges associated with complex storage in any industry, including the unerring and unforgiving requirement to

maintain system resiliency and reliability, and ensuring a robust backup and recovery regime. An outage can cost millions of dollars per hour in lost online sales and create equally as critical repercussions for customer confidence and loyalty.

Even short of an outright outage, storage issues can cause serious problems for an ecommerce vendor. Latency in any part of the ecosystem can snowball and impact the customer experience, leading to missing or slow-loading information, or delayed checkouts or security checks. These can all delay completion of sales or lead customers to abandon a purchase altogether.

While there's no easy way for a big ecommerce provider to avoid the need to own and run a complex ecosystem of databases, the engineers keeping all of them ticking are often desperate for a storage system that combines the resilience, speed, and performance needed to keep them operating at an optimal level.

How NVMe/TCP changes the rules

The introduction of NVMe/TCP has boosted storage performance by harnessing the power of parallelism. Unlike traditional storage protocols designed for hard drives, NVMe/TCP enables multiple parallel data paths between compute and storage resources over standard TCP/IP networks. This parallel architecture allows for a high level of concurrency, with support for up to 64K I/O queues, each containing up to 64K entries.

By leveraging this massive parallelism, NVMe/TCP delivers exceptional throughput and ultra-low latency, with each connection capable of achieving up to one million IOPS and latencies under 100 microseconds. This parallel processing not only maximizes CPU utilization but also enhances overall power efficiency - addressing critical concerns at both the rack and datacenter levels.

Introducing Lightbits for eCommerce

Lightbits helped develop the NVMe over TCP protocol before handing over the spec to the open source community. The company's own implementation of the standard is a software-defined data platform that allows ecommerce vendors to deploy the power of NVMe for block-based storage, while benefiting from built-in enterprise data services. These include snapshotting, cloning, restore and compression services - as well as replication across or within zones.

It operates both on-premises and in the cloud, giving ecommerce operators the ability to build out and optimize their own infrastructure, while bursting onto the cloud as necessary - whether to handle peak selling seasons or to allow for ongoing migration to the cloud.

The Lightbits platform supports the full range of SQL and NoSQL databases that engineers need to provision to support massive scale ecommerce. It handles up to one million IOPs per volume, while offering sub-millisecond latency, even under intense workloads. A three-server cluster can deliver two million IOPs, for a 50/50 load delivering read latency of 254 microseconds and 398 microseconds for writes, for example.

Moreover, latency does not increase as systems are scaled up. This gives infrastructure architects the confidence that they can scale up without compromising user experience or operational performance. This helps deliver a TCO, which can come in at least 50 percent lower than DAS, SAN, and HCI storage - as well as public cloud options. At the scale at which ecommerce vendors operate, that represents significant potential savings.

Conclusion

Ecommerce providers have a clear set of objectives: make it as easy as possible for customers to find and purchase the products they need; generate additional revenue through advertisers; manage inventory and other operations as efficiently as possible; and never stop optimizing it all.

This can only be achieved with an incredible amount of data, a dizzying array of esoteric databases, and the skills to get the best out of both. But it also requires underlying infrastructure that is up to the job - and storage is critical to this.

Deploying Lightbits NVMe/TCP gives platform engineers and infrastructure architects the speed they need - on premises or in the cloud - and the ability to scale up without increasing latency, while simultaneously reducing complexity and cost, and solidifying resilience.

Sponsored by Lightbits.

About Lightbits Labs

Lightbits Labs® (Lightbits), offers a complete, secure, and efficient data platform for eCommerce that accelerates data collection and online transactions. The NMVe/TCP and clustered storage architecture, coupled with essential data services solve the common eCommerce platform challenges of availability, reliability, performance, scalability, efficiency, and cost. With Lightbits disaggregated, software-defined storage you can build a robust eCommerce platform with unmatched performance, efficiency, agility, and flexibility while ensuring your customers get an exceptional digital experience.
Learn more about Lightbits for eCommerce.