

RELEX removes the limits on its big data ambitions with OpenIO Object Storage

Stakes

- **Highly time sensitive, data centric** business
- World's first operational software stacks equipped with **in-memory computing**
- **Storage scaling and reliability challenges to solve, while maintaining the high performance of their embedded solution**
- **High-SLA requirements**
- **Minimize vendor lock-in**



ABOUT RELEX SOLUTIONS

Founded in 2005 in Helsinki, Finland.

Sector: Retail optimization software.

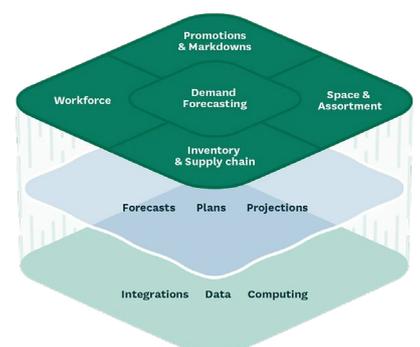
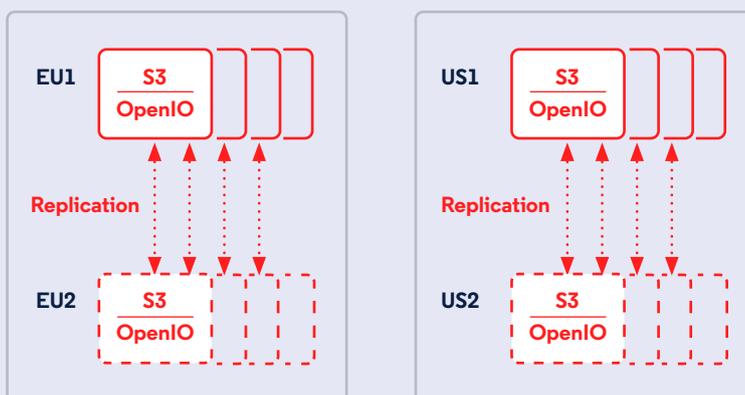
Clients: World's leading brands.

Activity: Analyzes big data to forecast demand per product, per store, per day. Daily stock and sales data for each product, store and warehouse is fed directly to RELEX from the customer's ERP system. Unique RELEX algorithms analyze those figures alongside historical sales data and external factors such as weather forecasts to calculate predicted demand over the coming days and weeks.

Success: 900+ employees in 11 countries

At a glance

- **4 Data Centers**
- **5 PB storage capacity**



www.relexsolutions.com



Many leading retail brands rely on Finnish company RELEX to keep their operations working at peak efficiency. Daily stock and sales data for each product, store and warehouse is fed directly to RELEX from the customer's ERP system. Unique RELEX algorithms analyze those figures alongside historical sales data and external factors such as weather forecasts to calculate predicted demand over the coming days and weeks. The RELEX platform tells the ERP system exactly how much of each product to order from suppliers for each store and warehouse.

**One place where OpenIO
outshines everybody else on
the client is the capability to
really scale from nothing to
petabytes of data using any
kind of hardware.**



Ilkka Niinivaara, IS Architecture Astronaut for RELEX sums up the company's value: "We use **a lot of data and hardcore number crunching** to make things better for retailers and their customers. We achieve some quite amazing numbers with 85% reductions in products being out of stock, 30% lower inventories, and a 40% reduction in waste, which means we're making the world a better place by throwing less stuff in the trash!"

The same forecast data is used to optimize other retail essentials such as store space and layout, the assortment offered in each branch, and even the workforce shift patterns.

In-memory computing delivers rapid data access

In the retail sphere, such analysis is **highly time sensitive**. "We might receive the data at 1am and have until 6am to send the order proposals back to the customer's ERP system," explains Mr. Niinivaara. "Our software usage patterns are very data intensive and data access really dominates the calculation time. Actual computation time for the algorithms is comparatively quick."

To accommodate these needs, RELEX developed one of the **world's first operational software stacks equipped with in-memory computing**. "To maximize the throughput of the core platform, we have implemented a completely custom, in-house, in-memory data platform," Mr Niinivaara adds. "It is employed as a **single java virtual machine, where all the algorithms, rules and business logic, the API layer - even our own proprietary in-memory database run in the same process**. There is **one separate installation of this process per customer** and initially it was not distributed in any way. The database access is as fast as local memory access, allowing us to encode data in memory in the most optimal order, guaranteeing it remains sequential and providing very good compression."

The challenge: scaling to support the world's largest retailers

On the deployment side, RELEX runs **bare metal** in co-location spaces. "Each customer runs a separate installation, single tenant, which was assigned to a specific server - basically an isolated island," says Mr Niinivaara. There were some obvious downsides to this megalithic approach. "Even though we achieved some amazing performance and customer satisfaction with our technology, it was obviously not limitless. What do you do when you can't buy a bigger server to host an even bigger customer? And you can imagine if there was a hardware failure on one of the servers, there was a huge scramble for the operations team to find available capacity on other servers and then restore the installations from cold backups. We knew a scale-out approach was needed."

Making the case for object storage

The RELEX team knew that to **resolve their scaling and reliability challenges while maintaining the high performance of their embedded solution**, they would need to **add distributed features** to their existing data platform. At the same time, they would gain replication for their database by externalizing the storage layer.

“There is an interesting implementation detail of our in-memory database,” reveals Mr Niinivaara. “Each transaction is committed as a new file that only contains the changes from that transaction. For any previous data, it just links to the previous commit files on disc. So, it’s an **ever-growing list of immutable files that reference each other**. So long as you know the name of the first one you can find all the others you need. This makes it **very suitable for object storage**.”

Searching for a partner, not just a storage solution

From an original list of 40 potential object storage vendors, RELEX chose OpenIO. “Our selection process was completely driven by the needs of the product development and operations teams - their technical needs and the cultural fit between our teams and the vendors’ team,” Mr Niinivaara recalls. “OpenIO were a great cultural fit. Working with them was really easy from the get-go and their values match our company values well.”

Object storage powers the RELEX private cloud

RELEX migrated to OpenIO Object Storage at their **primary data center in Finland** in 2018 and then brought online a **secondary site, also in Finland**, where OpenIO was installed from day one. These facilities are connected by a **private 40GB link** and **OpenIO’s replication features** are used to replicate data between the OpenIO clusters.

Since then, the company has gone on to open **two further data centers in the United States**, both based on OpenIO storage. Together the 4 datacenters have been architected to store up to **5 petabytes of operational data** as demand rises, with **further capacity used to maintain highly reliable back-ups**.

The RELEX private cloud operates across these 4 datacenters, delivering a **high-SLA** software as a service to customers around the world. There is also an option for on-premises installation of the RELEX solution at customer request. For maximum interoperability, RELEX uses the **S3 API** to link its OpenIO clusters with all other systems.

“The big difference between us and other companies is that we use object storage as our **primary storage** rather than just as a backup,” says Mr Niinivaara. “**OpenIO Object Storage is really at the core of our distributed features**. All the data between servers in a cluster is transferred to the object storage system. **Whenever there is a new transaction committed, the new file is uploaded to the object storage system and from there it is distributed through the other nodes**. This gives us **persistence for the file at the same time as replication to all of the nodes**.”



OpenIO were a great cultural fit. Working with them was really easy from the get-go and their values match our company values well.



Seamless scalability with full flexibility

To achieve the RELEX goal of handling billions of different datapoints for the world's largest retailers, **scalability** was always a top priority. "One place where OpenIO outshines everybody else is the capability to **really scale from nothing to petabytes of data using any kind of hardware, mixing and matching hard drives, servers, architectures**, whatever you want," Mr Niinivaara confirms. "You can just chuck more storage servers into an OpenIO cluster and it scales much more easily than the competitors, on all three axes of scalability. You get more speed, more capacity and more resilience just by adding more servers."

Flexibility was also key for file sizes, because RELEX commit **files can range from a few kilobytes to tens of gigabytes**. "It was essential that we find a system that could support a range of different sized objects in the same system. And we also wanted a solution that was flexible enough, in extreme cases, to allow for things like **converging the object storage layer on the same hardware as the compute**."

Open source and open minds

OpenIO's broader approach was also appealing: "The core parts of the solution are open source, which is good, particularly from a **business continuity and risk-management** perspective. Another attractive thing is always to **minimize vendor lock-in**. Using the **S3 API** allows that, but I think it's even more important that it allows us to use **standard SDKs** in all our clients."

For a company used to pushing the boundaries of what's possible, Mr Niinivaara says it was important to find a storage provider that was open to true partnership: "They are prepared to answer any kind of technical question – no matter how stupid or deep. I don't see any end in sight to the RELEX and OpenIO co-operation. Quite contrary, I think it will grow deeper and stronger."

You can just chuck more storage servers into an OpenIO cluster and it scales much more easily than the competitors, on all three axes of scalability. You get more speed, more capacity and more resilience just by adding more servers.

Key Takeaways

Scalability - OpenIO's innovative architecture is perfect for organizations seeking to use Artificial Intelligence to extract big value from their big data. It overcomes the limits of traditional object stores to support objects of all sizes and to enable seamless expansion and reconfiguration without service disruption. The capacity, speed and resilience of additional resources are available as soon as they are added to the cluster.

Flexibility – OpenIO's hardware-agnostic approach and commitment to open source software ensure that organizations have the freedom to expand and innovate without technical constraints or vendor lock-in.

Partnership – OpenIO's customers value the company's collaborative and open approach to projects that makes them easy to work with.

Learn more:
www.openio.io

