

MATRIX X TECHNOLOGY CORE

Proven, Ground-Breaking Technology

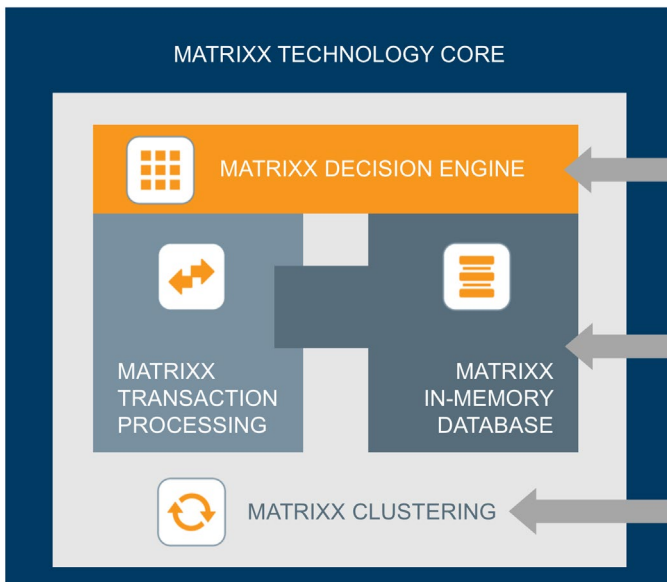
The launch of 5G will reset the performance requirements for telco systems. Legacy online charging systems, which dominate the BSS landscape today, will not be able to deliver the ultra-low latency or 100x throughput that 5G will unleash on back-office systems. Meeting these demands will require a single converged charging system (CCS) for all segments and services and not disparate systems and databases all glued together with middleware. Only by moving to an entirely new IT architecture, one built to handle the agility, scalability and performance demands of new network capabilities, will help operators realize the potential of such things as enhanced mobile broadband and ultra-low latency performance. Making that happen will require a new technology core capable of handling transaction volume and complexity at the speed of 5G.

At its foundation, the MATRIX X digital monetization solution is based on the patented MATRIX X Technology Core that was designed and built from the ground up to solve the fundamental issues uncovered with existing real-time processing and database

technologies. The goal was to create a single, unified solution with the network-grade reliability to support all digital monetization and network service delivery while providing more flexibility and configurability than even the most robust enterprise business software solutions. The solution required the efficiency and scalability to effortlessly handle the enormous, unpredictable workloads of tens/hundreds of millions of users purchasing and consuming a myriad of products and services at will.

Solving this technical paradox is not easy. Entirely new approaches for online transaction processing, business rule execution and clustering were needed to:

- Process real-time transaction volumes that are orders of magnitude greater than in previous generations
- Enable the flexibility to easily introduce new product concepts efficiently
- Deliver a complete digital capability on a small operational footprint



Patented at Every Layer

- Visual business rules configuration
- Translated to math for execution
- Extremely efficient & predictable
- Fully ACID compliant transactions
- No locking or blocking
- Solves transactional chaos
- All-master synchronous clustering
- Full site-redundancy replication
- Zero downtime maintenance & upgrades

Key MATRIXX Innovations Include:

The MATRIXX Technology Core is what enables the transformative capabilities of our monetization solution. With innovations in business rule execution, transaction processing, data management and clustering, the MATRIXX Technology Core provides the market agility and operational scalability critical to digital-first operating environments.

In-Memory Database

A non-blocking database with ground-breaking low latency and 100x the throughput of traditional relational databases, but with the full transactional integrity lacking in NoSql and web-scale datastores. To achieve ACID compliance with this efficiency, the database removes the massive processing costs of data locking, deadlock detection and lock escalation, and instead applies a patented algorithm to avoid transactional collisions and guarantee a precise answer every time with nearly zero overhead. Database reads and writes complete in nanoseconds, and transactions commit in microseconds. The database architecture also ensures that all data required to process an event is available in local memory, eliminating the unpredictable and expensive network hops associated with other distributed databases.

Transaction Processing

A new transaction management approach that supports fully parallel data access without compromising transactional integrity. The solution eliminates virtually all transactional interference at any scale by supporting completely parallel updates to “high velocity” data such as balances and meters. The overall outcome is consistently super-low latency (<10-millisecond average API request-to-response) end-to-end across all use cases. The system is not sensitive to business logic complexity or total throughput, and predictably handles even the most complex use cases such as balances shared across a large enterprise.

Decision Engine

A high-speed decision engine for business rule execution that operates predictably and consistently across complex, multi-variant decision processing scenarios (rating, policy, GL account assignment, offer prioritization, etc.). Traditional decision engines have use cases coded or scripted into them using serially processed conditional logic (IF-THEN-ELSE statements). Additional coding/scripting to implement new use cases is often required, resulting in poor, and worse, unpredictable, performance due to the inefficiency of executing this type of workflow on modern processors.

MATRIXX focused on the two fundamental components of any business decision — parametric analysis, necessary to determine a “circumstance” of the decision, and execution, the configurable action based on that circumstance. The MATRIXX Decision Engine optimizes each of these stages to maximize configurability and replace conditional logic with multi-dimensional mathematical algorithms.

The first step, called Normalization, determines the business circumstance relative to the decision being made. The MATRIXX Decision Engine architecture makes no assumptions around the type of decision, so the Normalization stage allows analysis to be configured to any set of parameters for any given decision. For instance, the MyMATRIXX configuration GUI makes it easy to configure and analyze any of these transformational examples, as stand-alone or combined:

- The time of a usage event to determine peak time vs. off-peak time
- A balance to determine if the user qualifies for a discount based on their spend-to-date this cycle
- Various parameters about owned offers and current sponsorships to determine if data traffic should be charged or zero-rated
- Offer attributes to determine which GL category should be applied to a purchase

Each analysis has a configurable set of input parameters, a configurable set of result values and a specified Normalizer that defines the math library used for the analysis. Any number of Normalizers can be configured for making a specific decision. A rating decision could, for instance, depend on a time-of-day analysis, a spend-to-date analysis and a network-zone analysis, allowing each decision to use only and exactly the libraries required to determine the circumstance with each math library executing nearly instantly on modern processors.

The second step, called Action Execution, has been designed to be equally efficient. The actions for each circumstance in a decision are also configured in MyMATRIXX using an intuitive, excel-like table

format. These actions are then translated from their “human-friendly” form into a machine-friendly, multi-dimensional matrix of mathematical equations (thus the name of the company). Not only does this format enable extremely proficient lookup and execution of the appropriate action, but it allows for highly efficient recalculation of a new circumstance upon a boundary condition (e.g., a balance crosses the discount trigger point during the rating process, the session spans from peak-time to off-peak, etc.).

The net result of the patented Decision Engine architecture is that any data available can be analyzed in any reasonable way and used to trigger an action all without coding or scripting. The latency of the entire process, even for the most complex rating scenarios, is sub-millisecond.

Real-Time Clustering

An entirely new, patented approach to local clustering and geo-redundancy, legacy lock-based transaction control requires that each piece of data in the database is owned by one of the servers in the cluster. Ownership can be distributed across the cluster for different sets of data; however, to support locking, only one server can initiate an update to a given piece of data. The impact is extremely unpredictable overhead and latency for complex transactions like rating, where many pieces of data are accessed and updated. As a result, locks may end up distributed across multiple servers.

MATRIXX leverages non-locking database technology to instead create a cluster of atomically synchronized servers in an all-master configuration. Any event can be fully processed by any server, and all the data required is locally available for instant access. Transactions are committed across the entire cluster as a single atomic action, so all the servers are guaranteed to have consistent views of the database at all times (N-way redundancy). All transactions are also logged to multiple storage locations for additional resiliency.

Immediately upon commit, the transaction data is published to one or more secondary sites where another cluster updates all the database servers

within milliseconds and again logs the transaction to multiple stores. All types of data are site-replicated instantly in this manner, including in-session data and balances, so even upon a full site failure, the MATRIXX digital monetization solution can automatically reroute incoming requests and continue providing full service at full performance.

Adaptive Quota Management (AQM)

AQM is a new approach to usage quota management that measures the rate at which a specific session is consuming a balance or asset, and automatically adjusts the quota sizes and expiration times returned to the network to meet configured business and network requirements.

AQM solves the puzzle of how to provide highly accurate, timely usage threshold notifications, even across shared enterprise balances, without resorting to tiny quotas that will escalate network signaling loads out of control. The patented algorithm allows an operator to configure desired timing accuracy curves for each type of session, specifying not only the desired average message load for the session but also a geometric decay curve. It is followed as a notification threshold approaches, efficiently reducing quota sizes in a fair allocation scheme, so all sessions sharing the same balance cooperate to hit the threshold without quota starvation or hoarding. It is the only quota

management solution that automatically adapts the quota size decisions based on the actual consumption rates of all sessions involved.

AQM allows customers to provide to-the-second notifications to their subscribers without service jitter while maintaining an overall control message load that is a fraction of what would be required otherwise.

The MATRIXX Technology Core

The MATRIXX Technology Core is the proprietary high-speed decision engine and in-memory database that enables the transformative capabilities of our monetization solution. Its patented decision-table architecture enables a click-not-code capability, eliminating the need for custom development to support new business requirements. Its high-speed decision engine has a redesigned transaction processing model that eradicates performance bottlenecks so that every customer interaction is instant and precise. Its in-memory database puts an end to data silos, providing a single source of truth for all customer and partner channels. With innovations in business rule execution, transaction processing and data management, the MATRIXX Technology Core provides the market agility and operational scalability critical to digital-first operating environments.

About MATRIXX Software

MATRIXX Software delivers a modern converged charging and digital monetization solution proven at scale. Global operators like Telefónica and Telstra, IoT providers like Tata Communications and network-as-a-service (NaaS) providers like DISH rely on the platform to overcome the limitations of traditional Business Support Systems (BSS). With MATRIXX, service providers can rapidly configure, deploy and monetize personalized, innovative offerings. Its cloud native platform delivers accurate, real-time information that improves customer engagement. MATRIXX enables commercial innovation and real-time customer experiences that drive revenue and growth opportunities across multiple markets.

matrixx.com