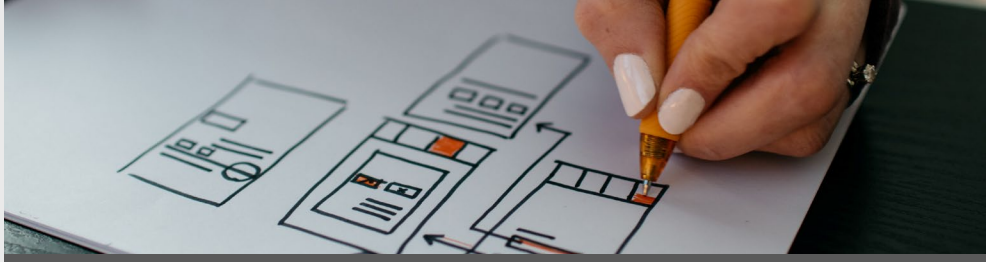

Sparkling Logic SMARTS™

Decision Management Platform

- Intuitive authoring (RedPen, fluid metaphors)
- Data-based decision testing and impact analysis
- Lifecycle and deployment management

Sparkling Logic SMARTS™ is an end-to-end, low-code/no-code decision management platform that spans the entire business decision lifecycle, from data import to decision modeling to application production.

SMARTS was designed by a team of experts who wanted to deliver a powerful yet simple product so that a business analyst could start with data and build decision logic with built-in predictive data analytics and execution decision analytics.

At the highest-level, SMARTS is the sum of four capabilities:

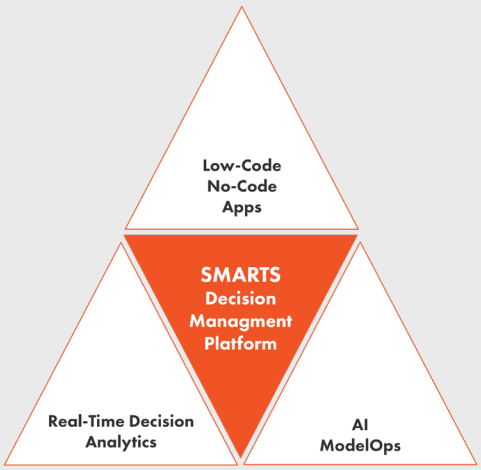

A decision management platform which leverages data samples to intuitively author decision logic, test in place, and measure the impact analysis. The platform features all the capabilities needed for enterprise-level lifecycle and deployment management

A low-code/no-code apps environment to design decision services, augment them with an intelligent front-end, and allow untrained business users to configure the decision logic while enforcing governance in a no-code business app

An AI & ModelOps environment to cover the full spectrum of operationalizing predictive models, from importing models built by data scientists to launching learning jobs to augment the existing strategy

A real-time decision analytics environment to define dashboards, measure the quality of the business performance, and to seamlessly associate the metrics to the real-time execution environment

Having access to concrete transactions to provide context, business analysts can immediately start importing data, capturing their decision logic, test how it applies to each transaction, and measure the collective impact of making changes.

Our patented technology RedPen™ accelerates rule authoring without a need-to-know rule syntax. We also support the Decision Model and Notation (DMN) through our technology Pencil. DMN provides a standard-based approach to discover the requirements for the decision logic that can then be made directly executable in SMARTS.

The platform can cater to the needs of a stand-alone project for a start-up to an enterprise deployment for a Fortune 100 company. For the latter, the decision logic can be shared and distributed through a library, or segmented per product, client, region, or any other dimension.

The decision management platform allows end-to-end lifecycle management, including release management. For instance, business analysts can also view the history of their decision logic and revert to an earlier release at any time. SMARTS lifecycle management and governance features allow changes to be published all the way to production, without involving IT. Not only, the resulting decision services deploy seamlessly into cloud infrastructure or on-premise, but the same application can be moved from on-premise to the cloud and vice versa, taking full advantage of elastic capabilities:

- **Support for secure service invocations in an authenticated context:** Many users may invoke the service concurrently, using any client technology that can interact with services
- **Support for scalability:** Decision engines executing within the decision service will leverage all cores and instances available to them within the installation. Adding more cores or instances statically or elastically results in the ability for the engine to support more concurrent executions
- **Support for high availability through redundancy:** Having multiple instances with replicated repositories removes single points of failure. If one an instance is taken out, the rest of the replicated instances can continue with the load
- **Support for no-downtime hot swap of decision logic with full traceability:** Without any IT intervention, users can change the release of the decision logic being executed with one click, and no downtime. SMARTS loads the new release and hot swaps it automatically
- **Support for ready-to-execute decision logic:** SMARTS allows users to specify when a decision service is declared to be ready to receive invocations. Typically, when users want to make sure that the decision logic is called only when it is actually loaded, compiled, and cached in memory