

LSEG Markets Technology Analytics



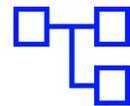
The LSEG Markets Technology suite provides robust, scalable and high-performance solutions for operators of trading venues, clearing houses, central securities depositories and related businesses. The products are used by LSEG to operate its own businesses and are proven in demanding global markets and regulatory environments.

Unified, high-performance, and cloud-native analytics system

The LSEG Markets Technology Analytics product empowers financial institutions to make fast and effective decisions by transforming financial markets data streams into actionable insights in real time. Integrated with other products in the LSEG Markets Technology suite, the Analytics system supports use cases including (but not limited to) market performance evaluation, market maker monitoring, participant performance/behaviour analysis, regulatory compliance, and offering analytics services to customers to create further value out of data.



Built on a modern stack of big data technologies, the system facilitates a unified workflow, with capabilities to handle all stages of the analytics life cycle, with a proven track record of converting terabytes of data into consumer-ready insights daily. The cloud native design of the system enables highly scalable, cost-effective solutions which grow with data and usage. The modularity of the components allows co-existence of commercial and open-source technologies to achieve an optimal balance between functionality and cost, while providing the flexibility to upgrade technology and the toolchain to leverage advancements of big data technology to be future-ready.



Business teams with data engineering and data science skills can leverage the self-serviceable elements of the Analytics system to build data pipelines, analytics, and dashboards to cater to business demands with low turn-around time. The product supports deployment of a sand-box environment(s) to perform R&D and paths to promote workloads to production environments upon being validated by business teams for accuracy. This enables data scientists to perform rapid 'build-deploy' cycles to support dynamic business needs independently.



Key facts

01

Unified system for a complete analytics workflow:

Modernised, big data technology stack, covering all stages of data science: ingestion, processing, storage and visualisation. Modular architecture enables the use of leading-edge tools for a robust analytics process, while providing flexibility to upgrade components independently.

02

Real-time analytics engine: Convert high velocity trading data into consumer-ready analytics suites in real time. The built-in complex event processing engine enables multi-dimensional computation and analysis of trading data to support business needs with descriptive and predictive insights.

03

Multi-tier data lake: High performance, secure storage for structured and unstructured data, tiered based on data temperature to achieve the right balance between performance and cost. Unified catalogue and interface enable centralised access and management of all data, with comprehensive governance to control data exposure based on entitlements.

04

Self-serviceability for data and analytics: Empower in-house data scientists to take complete control of analytics processes to serve needs of business teams. Sandbox oriented workflows bring leading edge development tools into the hands of data scientists to drive innovation and growth, while eliminating the overheads around managing infrastructure and setting up build pipelines.

05

Hardened for extreme market conditions: Proven in generating analytics from billions of data records, amounting to terabytes of data from trading venues, market data providers and news feeds.

Functional highlights

01

Data ingestion, storage and distribution: Facilitates acquisition of data in real time from streaming sources and in periodic batches for data files. Ingestion of data is facilitated by using data pipelines which contain business logic to clean and transform data, with necessary quality checks to ensure accuracy and integrity. Ingested data are stored in the multi-tier data lake, with the ability to manage data within tiers to achieve the required balance between fast access and cost effectiveness. Data lake provides a unified (across tiers), SQL based query interface for consumer applications to retrieve data.

02

Data governance and security: Data is encrypted both at rest and in-transit where access is governed by data entitlements which control exposure of data tables, columns and records, preventing unauthorised access of restricted data. Data lake is equipped with a data catalogue which acts as the central point of meta-data for all consumers. The Analytics system has an interface with LSEG's trading system for orders, trades, reference data, and market data, enabled off-the shelf, with the option to enable data flows from rest of the products in the LSEG Markets Technology suite to serve the need for a central data lake.

03

Real-time and periodic analytics: Supports real-time and periodic computation of analytic insights using incoming data streams and data stored in the data lake. Produced analytics are made available to consumers in form of streams (for real-time consumption) and stored in the data lake (for periodic consumption). Computation of analytics is supported by complex event processors which aggregate data in real-time to build and maintain multi-dimensional states. Off-the-shelf support is available with utility functions for common operations such as constructing and analysing order books in real-time to derive insights from trading activity. Technologies and tool chain used in analytic computation workflow can be coupled with development tools and libraries specialised in machine learning to facilitate development of predictive analytics.

04

Data and analytic visualisation: Provides means for no-code development of dashboards to retrieve, aggregate and visualise data, to deliver analytic insights to business teams. Visualisation options can be selected from a variety of charts, options including but not limited to, data grids, cartesian (bar, line and scatter) charts, heat/tree maps and candlestick charts. Dashboards can be composed of different types of visualisations which can be linked with interactive, selection-based data filtering to enable visual storytelling on analytics. Dashboard designer enables non-technical users to effectively and efficiently compose dashboards, while providing flexibility for more advanced users to take control of visualisations in depth for complex scenarios. Dashboard viewer is automatically scaled in/out based on the number of users, facilitating solutions which deliver analytic services to a large, growing number of consumers.

05

End-to-end self-serviceability: Enables data engineers/scientists to acquire data from new sources, build analytics, compose dashboards in a contained sandbox environment and seamlessly promote them into production after passing the required quality gates. Provides off-the-shelf support for development (sandbox) infrastructure management, repository to manage created artefacts and promotion paths to promote artefacts from sandbox to production environment. Daily synchronisation of data between production and sandbox environments can be enabled optionally if sandbox needs replicas of certain production data stores for a defined retention period.

Key business benefits



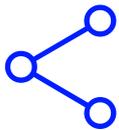
Fast, effective decisions with real-time insights: Accelerate decision making of business teams by converting data streams into actionable insights. Leap beyond the traditional approach to analytics by harnessing the power of machine learning combined with big data technology, to create a powerful mix of descriptive and predictive insights.



Create extra value out of data: Monetise petabytes of trading data to generate revenue by creating analytics services to customers. Leverage auto-scaling to manage economies of scale for effective growth.



Short time-to-market: Self-serviceability enables institutions to cultivate data science skills within organisation to support analytics needs of business teams.



Scalability and cost effectiveness: The cloud native design of the system enables highly scalable, cost-effective solutions which grow with data and usage.



Future-ready, modern technology stack: Built on a modern stack of big data technologies, the product facilitates a unified workflow, with capabilities to handle all stages of analytics life cycle. Modularity of the components allows co-existence of commercial and open-source technologies to achieve optimal balance between functionality and cost, while providing the flexibility to upgrade technology and toolchain to leverage advancements of big data technology and trends in AI/ML.

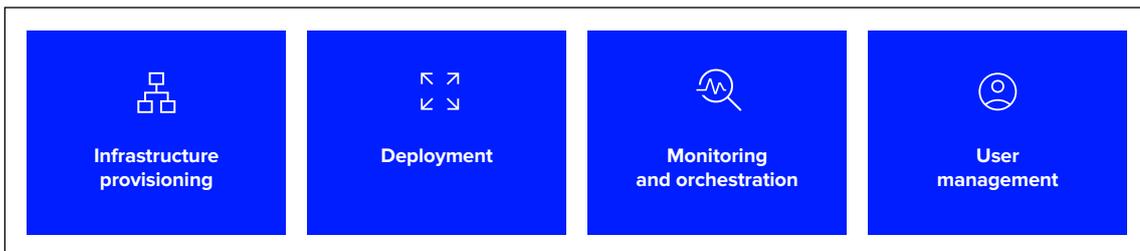
Key components

The technology agnostic, modular design of the analytics product allows the co-existence of commercial and open-source technologies to achieve the optimal balance between features and cost, with the option to select between cloud and on-premises deployment.

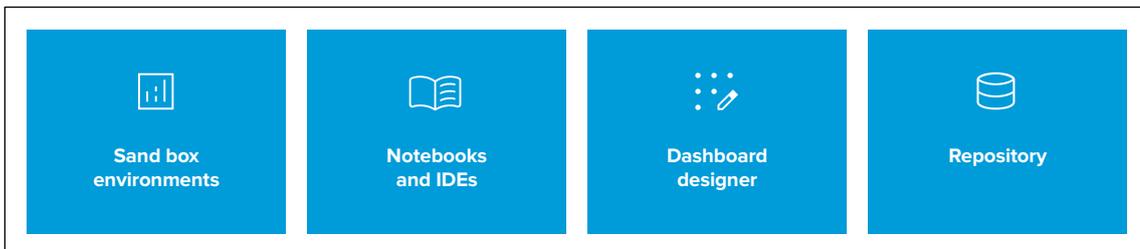
Core capabilities



Management operations



Self-serviceability



Technical overview



Leveraging a proven technology platform: Built on AWS, LSEG Markets Technology Analytics combines streaming analytics (Amazon Managed Flink) and batch processing (AWS EMR Spark) with orchestration via AWS Managed Airflow. The architecture ensures flexibility and efficiency across ingestion, transformation, storage, and visualisation workflows.



High performance and extensive scalability: The system supports dynamic scaling of ingestion pipelines, EMR clusters, and GUI microservices to handle growing workloads and up to 1,000 concurrent users. Latency-sensitive pipelines are isolated using Kafka-based message brokering to prevent backpressure, ensuring consistent performance.



Reliability and fault tolerance: Multi-AZ deployments, exactly-once processing guarantees, and durable backups ensure resilience. Managed Flink supports multi-AZ Kubernetes pod and Zookeeper deployments for high availability. Application logs are captured via Datadog agents for proactive monitoring.

Technical platform

- **Deployment:** AWS Cloud (supports hybrid if required)
- **Architecture:** Distributed, message-based, microservices design

Contact us

Get in touch with us at markets.technology@lseg.com

Visit our [website](#) →

Disclaimer

This document contains text, data, graphics, photographs, illustrations, artwork, names, logos, trademarks, service marks and information ("Information") connected with LSEG.

Information contained in this document is proprietary to LSEG. Unless with the prior written approval of LSEG, this document may not be reproduced or disclosed, whether as a whole or in part, for any party (other than to any individual who has a need to peruse the content of this document in connection with the purpose for which it is submitted) or use for any purpose, other than the purpose for which it is submitted or disclosed by LSEG.