IBM System S Overview

IBM is developing platforms to address the requirement to efficiently extract knowledge and information from potentially enormous volumes and varieties of continuous data streams. The project known as System S is designed to scale from systems that acquire, analyze, interpret, and organize continuous streams on a single processing node, to high performance clusters of hundreds of processing nodes.

System S Overview:
System S provides an execution platform and services for user-developed applications that ingest, filter, analyze, and correlate potentially massive volumes of continuous data streams. It supports the composition of new applications in the form of stream processing graphs that can be created on the fly, mapped to a variety of hardware configurations, and adapted as requests come and go, and relative priorities shift.

System S supports high volume, structured & unstructured streaming data sources such as: images, audio, voice, VoIP, video, TV, financial news, radio, police scanners, web traffic, email, chat, GPS data, financial transaction data, satellite data, sensors, badge swipes, etc.

System S was designed to address the following data management platform objectives:

- Parallel and high performance stream processing software platform capable of scaling over a range of hardware capability
- Agile and automated reconfiguration in response to changing user objectives, available data, and the intrinsic variability of system resource availability
- Incremental tasking in the face of rapidly changing data forms and types
- Secure, privacy-compliant, and auditable execution environment.

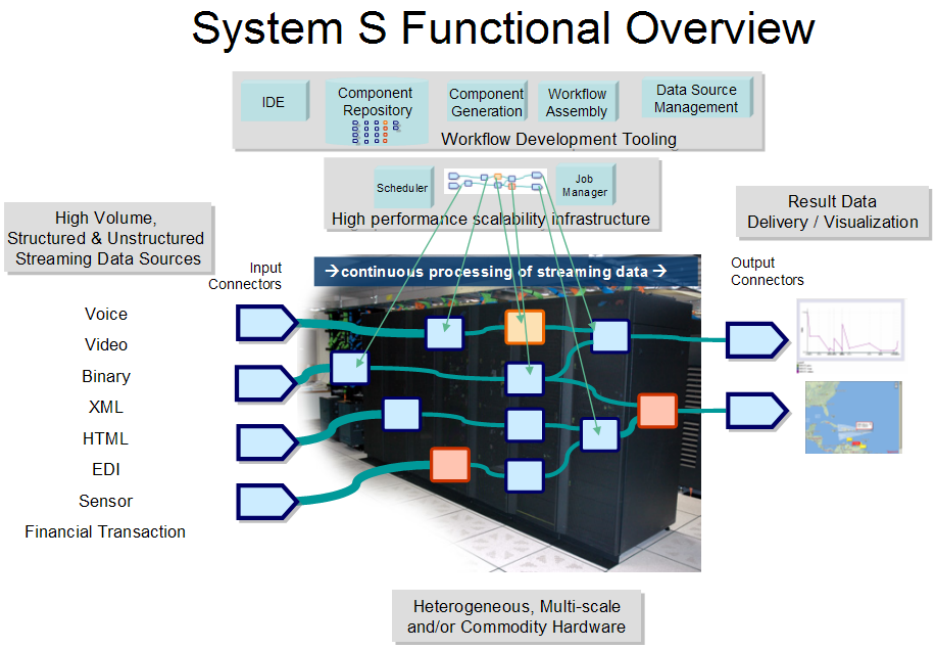

References:

http://domino.research.ibm.com/comm/research_projects.nsf/pages/esps.index.html