

ZettaFlow: Towards High-Performance ML-based Analytics across the Digital Continuum

Gabriel Antoniu, Alexandru Costan, Ovidiu Marcu

Univ Rennes, Inria, CNRS, IRISA

October 12, 2019

At the previous BDEC2 workshops we introduced the Sigma data processing architecture [Antoniou2019], which aims to enable unified data processing across hybrid infrastructures combining Edge, Cloud and HPC systems. In this white paper, we discuss how this architecture could be leveraged to enable the seamless execution of machine learning algorithms at the intersection of these domains.

The Sigma architecture combines batch-based and stream-based Big Data processing techniques (i.e., by extending the traditional Lambda architecture) with in situ/in transit data processing techniques inspired from the HPC area (Figure 1). It allows to collect, process and analyze extreme volumes of past data, real-time data, jointly with simulated data both in situ (where the data is generated) and in transit (e.g., on specific resources dedicated to analytics), before the data is shipped to long-term persistent storage.

To provide a reference implementation of the Sigma architecture, our approach consisted of jointly leveraging two existing software components: the Damaris [Damaris,Dorier2013] middleware for scalable in situ/in transit processing and the KerA [Marcu2018] unified system for ingestion and storage of data for scalable stream processing.

Data shifts to the Edge. By 2022 Gartner predicts that 75% of enterprise-generated data will be created and processed outside of the data center and cloud infrastructures, compared with 10% today [Gartner2018]. Recent developments brought by **5G networks** will be a key enabler of the future digital world. They will bring new service capabilities for industrial/research stakeholders thanks to the unprecedented on-demand performance and real-time reactivity. For example, energy, transport, manufacturing and water utilities will be capable of connecting to millions of networked devices, taking real-time, intelligent and autonomous decisions (roundtrip latency in 1ms range) [Ericsson2019]. All these data will be continuously **streamed** at high rates to the processing sites (at the Edge, in the Fog or in the Cloud), requiring frameworks able to cope with their **real-time processing requirements**.

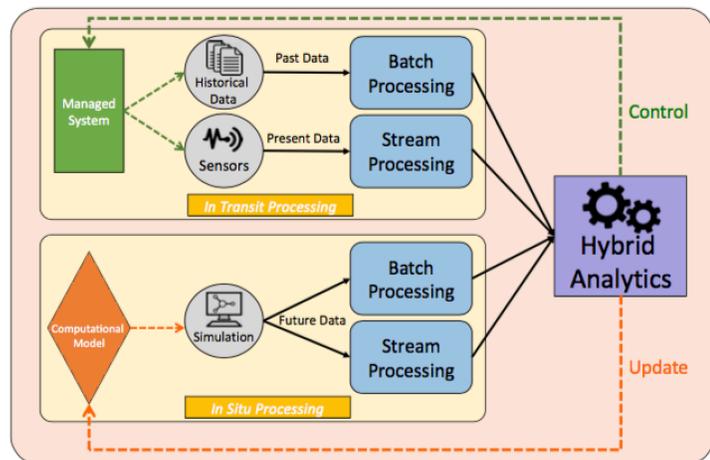


Figure 1. The Sigma data processing architecture.

Unified framework for stream storage, analytics and machine learning. Our vision to support the next generation of analytics and machine learning applications relies on a **unified approach for stream storage and processing**, spanning from the Edge (where data is generated) to the Fog/Cloud and HPC centers (where it is processed).

To illustrate this vision, we introduce ZettaFlow, a high-performance multi-model analytics-oriented storage system, that extends the KerA approach by proposing a unified multi-model storage engine (i.e., supporting *streams*, *key-value*, *columnar* APIs). To do so, it relies on RAMCloud [Ousterhout2015] (a low-latency key-value store) and Apache Arrow [Arrow2019] (a high-performance in-memory columnar storage system).

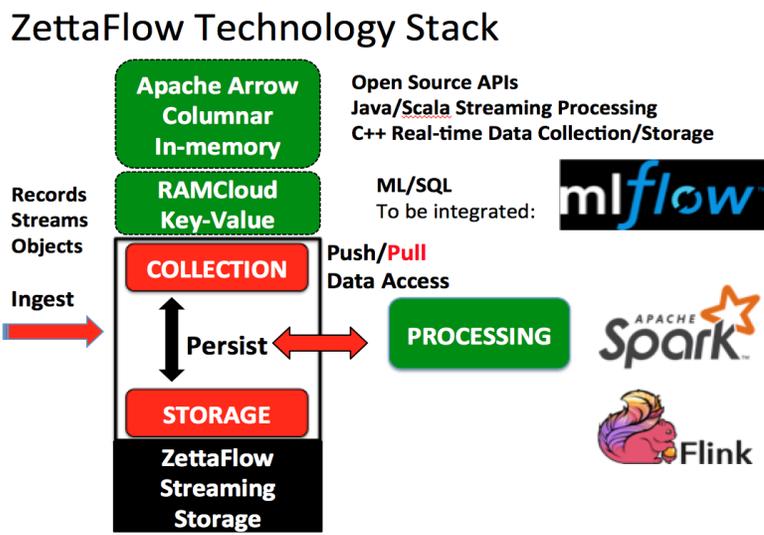


Figure 2. The ZettaFlow architectural overview.

Dedicated support for machine learning. ZettaFlow is currently integrated with:

- The Apache Flink data analytics framework, through a shared memory approach effectively co-locating storage and processing;
- The Damaris middleware through streaming RPC (asynchronous interfaces). The implementation relies on native C++ integration, to avoid Big Data Java serialization overheads, and to efficiently support high-performance networks like Infiniband and Intel DPDK.

We are now considering integration of ZettaFlow with **MLflow** [Zaharia2019] (an open-source platform for the end-to-end machine learning lifecycle) and Apache Spark. The overall goal is to build a **reference high-performance architecture for unified real-time analytics and machine learning**. This unified approach will provide machine learning workloads a high performance experience similar to the one currently enabled for stream-oriented analytics.

Questions and Answers

- 1) *What innovative capabilities/functionalities will the proposed candidate platform demonstrate (e.g. transcontinuum workflow, edge computing, data logistics, distributed reduction engine, etc.)?*

The above integration provides the possibility to explore the possibility to **dynamically push to storage** the critical aspects of machine learning workloads (i.e., *in-storage processing*). As a step further, we will investigate specific optimizations by studying machine learning data access patterns to effectively optimize the real-time flow of data that traverses both ZettaFlow and machine learning frameworks.

- 2) *What applications/communities would/could be addressed?*

5G applications like smart vehicles and transport, critical services and infrastructure control, human-machine interaction and critical control of remote devices.

- 3) *What is the “platform vision,” i.e. what kind of shared cyberinfrastructure (CI) for science would further research/design/development of this platform lead to?*

A reference high-performance multi-model analytics storage platform supporting workflows that **combine machine learning and real-time analytics**.

- 4) *How available/ready/complete is the set of software components to be used to build the demonstrator?*

We are working towards an efficient integration of ZettaFlow with machine learning frameworks.

- 5) *As far as one can tell at this early date, to what extent can this be done with existing and/or otherwise available hardware/software/human resources?*

Building an operational AI demonstrator requires interaction with industrial and research partners to identify machine learning access patterns and develop AI models that will help ZettaFlow self-optimize and reach bare-metal performance, while considering the variety of use cases, effectively optimizing the critical machine learning aspects. We are also in the process of building a startup that will leverage the ZettaFlow platform for real-time edge/cloud data analytics.

References

- [Antoniou2019] Gabriel Antoniu, Alexandru Costan, Ovidiu Marcu, Maria S. Pérez, Nenad Stojanovic: Towards a demonstrator of the Sigma Data Processing Architecture for BDEC2, BDEC2 workshop, Poznan, May, 2019. White paper.
- [Gartner2018] Smarter with Gartner: What Edge Computing Means for Infrastructure and Operations, 2018.
- [Ericsson2019] 5G use cases, Ericsson, <https://www.ericsson.com/en/5g/use-cases>.
- [Marcu2018] Ovidiu Marcu, KerA: A Unified Ingestion and Storage System for Scalable Big Data Processing, PhD thesis, 2018. <https://tel.archives-ouvertes.fr/tel-01972280>.
- [Ousterhout2015] John Ousterhout, Arjun Gopalan, Ashish Gupta, Ankita Kejriwal, Collin Lee, Behnam Montazeri, Diego Ongaro, Seo Jin Park, Henry Qin, Mendel Rosenblum, Stephen Rumble, Ryan Stutsman, and Stephen Yang. “The RAMCloud Storage System”. In: ACM Trans. Comput. Syst. 33.3 (Aug. 2015), 7:1–7:55
- [Arrow2019] Apache Arrow, <http://arrow.apache.org/>
- [Zaharia2019] Matei Zaharia, Andrew Chen, Aaron Davidson, Ali Ghodsi, Sue Ann Hong, Andy Konwinski, Siddharth Murching, Tomas Nykodym, Paul Ogilvie, Mani Parkhe, Fen Xie, Corey Zumar: Accelerating the Machine Learning Lifecycle with MLflow, USENIX’19, Santa Clara.

- [Dorier 2013] M. Dorier, R. Sisneros, T. Peterka, G. Antoniu, D. Semeraro, “Damaris/Viz: a Nonintrusive, Adaptable and User-Friendly In Situ Visualization Framework”, Proc. LDAV – IEEE Symposium on Large-Scale Data Analysis and Visualization, Oct 2013, Atlanta, USA. URL: <https://hal.inria.fr/hal-00859603>
- [Damaris] The Damaris project. <https://project.inria.fr/damaris/>