accenture

Strategy | Consulting | Digital | Technology | Operations

High performance. Delivered.

Data Acceleration: Architecture for the Modern Data Supply Chain

Overview

Data technologies are evolving rapidly, but organizations have adopted most of these in piecemeal fashion. As a result, enterprise data—whether related to customer interactions, business performance, computer notifications, or external events in the business environment —is vastly underutilized. Moreover, companies' data ecosystems have become complex and littered with data silos. This makes the data more difficult to access, which in turn limits the value that organizations can get out of it. Indeed, according to a recent Gartner, Inc. report, 85 percent of *Fortune* 500 organizations will be unable to exploit Big Data for competitive advantage through 2015.[1] Furthermore, a recent Accenture study found that half of all companies have concerns about the accuracy of their data, and the majority of executives are unclear about the business outcomes they are getting from their data analytics programs.[2]

To unlock the value hidden in their data, companies must start treating data as a supply chain, enabling it to flow easily and usefully through the entire organization—and eventually throughout each company's ecosystem of partners, including suppliers and customers. The time is right for this approach. For one thing, new external data sources are becoming available, providing fresh opportunities for data insights. In addition, the tools and technology required to build a better data platform are available and in use. These provide a foundation on which companies can construct an integrated, end-to-end data supply chain.

1. "Big Data Business Benefits Are Hampered by 'Culture Clash'," Gartner, September 12, 2013.

2. "Journey to Analytics ROI," Accenture, February 27, 2013.

A modern data supply chain begins when data is created, imported, or combined with other data. The data moves through the links in the chain, incrementally acquiring value. The supply chain ends with actionable, valuable business insights—such as ideas for new product, service, or process innovations, marketing campaigns, or globalization strategies. Configured and managed effectively, a data supply chain enables organizations to discover their data, leverage more data sources, and accelerate data. These capabilities, in turn, position an organization to extract more value from its data through advanced computing techniques such as machine learning.

Data acceleration plays a major role in a robust data supply chain. In its simplest form, data acceleration stems from tools and techniques that enable massive amounts of data to be ingested (transported from its source into a system designed for data), stored, and accessed at lightning speeds. Specifically, with data acceleration, organizations gain quick access to valuable data—which enables them to perform analysis on the data, gain insights, and take actions in the sometimes very small window of opportunity available to businesses. Data acceleration thus helps organizations surmount three data-related challenges: movement, processing, and interactivity.

In this point of view, Accenture's Big Data practice in collaboration with Accenture Technology Labs closely examine those challenges and assess the landscape of architectural components available to address them. We then explore options for combining these components to create data platform solutions.

Three challenges that data acceleration can address

Data acceleration helps organizations address three challenges: how to move data swiftly from its source to places in the organization where it is needed, how to process it to gain actionable insights as quickly possible, and how to foster faster responses to queries submitted by users or applications—referred to as interactivity.

Movement

Traditionally, bringing data into an organization was a slow but fairly straightforward process: Data was collected in a staging area and then transformed in to the appropriate format. The data was then loaded to reside in one source, such as a mainframe or enterprise data warehouse. From there it was directly transferred in a point-to-point fashion to a data mart for users and applications to access. However, with the mammoth increase in data volumes and variety, such a traditional process no longer works effectively.

The Internet of Things (IoT) is playing a major role in driving new developments in data movement. In its simplest sense, the IoT comprises connected devices—ranging from refrigerators, smart meters, and video cameras to mobile phones and children's toys —that could be located anywhere in the world. According to Gartner, Inc., there will be as many as 26 billion devices on the IoT by 2020.[3] Every connected device generates data, each with its own format and idiosyncrasies.

Whether a business is deploying thousands of individual systems or simply trying to keep up with its own growth, having a modern data infrastructure in place that can collect relevant data can lead to differentiation by enabling data insights. But to extract valuable insights from data in this new world, organizations need to harness it from multiple sources without losing any of it, and deliver it for processing and storage. Some data exists as log files on external systems that have to be transported to an organization's data infrastructure for future use. Other sources provide streaming data, which is piped into the system in real time; that is, as the data is generated. Examples include power consumption information from smart electrical meters that is always updating.

Whatever the source and format, moving the data from its origin to where it is needed in the organization can seem like drinking from a fire hose while trying not to lose a single drop. Data acceleration helps organizations manage this feat by enabling multiple ways of bringing data into an organization's data infrastructure and ensuring that it can be referenced quickly.

Processing

Organizations have long been processing data in an effort to extract actionable insights from it. However, the volume and variety of data requiring processing have ballooned. To accommodate growth on those two fronts and generate faster but also accurate results, enterprises have to step up their processing capabilities. In particular, they must carry out three activities more speedily than ever: performing calculations on the data, creating and executing simulation models, and comparing statistics to derive new insights from the data.

The rise of real-time analytical technologies has presented new opportunities on this front. A good analytical technology pre-processes incoming data. For example, by monitoring a customer's location, an organization can deliver a promotion or discount to a customer's mobile device when he or she is near a likely place of purchase. But a better technology combines streaming data with historical (modeled) data to enable more intelligent decision making. For instance, by correlating a customer's location with his or her previous purchase history, the company can deliver a promotion that is tailored to that same customer, increasing the likelihood of conversion.

3. "Gartner Says the Internet of Things Installed Base Will Grow to 26 Billion Units By 2020,." Gartner, December 12, 2013.

To reap the full benefits of faster data processing, companies must make better use of computer clusters—organized sets of hundreds or thousands of computers working together to sift through large amounts of data. With the cost of random-access memory (RAM) at an all-time low, new solutions for extracting data from storage more quickly have bombarded the market, each with its own promise of speed, durability, and accuracy.

Data acceleration supports faster processing by leveraging advances in hardware and software for computer clusters, enabling them to operate more efficiently than ever.

Interactivity

Interactivity is about usability of the data infrastructure. Fundamentally, users or applications submit queries to the infrastructure and expect to receive responses to the queries within an acceptable amount of time. Traditional solutions have made it easy for people to submit queries to get the results they need to arrive at actionable insights. However, the rise of big data has led to new programming languages that discourage existing users from adopting the systems. Additionally, owing to the sheer volume of data, users may have to wait many minutes or even hours for results on a query.

The longer users have to wait, the more time it takes them to gain the insights required to make the business decision facing them and to satisfy their clients' expectations. That is the case whether clients are internal (such as a marketing director who wants to know which of the company's customers are most loyal and profitable) or external (for example, a business process outsourcing (BPO) client company that needs to know how performance of an outsourced process has changed over the term of the BPO engagement). Clients providing critical services to their own customers—such as retail transaction processing—might require response times in the sub-second (millisecond) range. With less critical business tasks, acceptable response times may be longer.

Data acceleration supports faster interactivity by enabling users and applications to connect to the data infrastructure in universally acceptable ways and by ensuring that query results are delivered as quickly as required.

Understanding the architecture landscape

Organizations can choose from many different data technology components to build the architecture needed to support data acceleration. These include big data platforms, complex event processing, ingestion, in-memory databases, cache clusters, and appliances. Each component can address data movement, processing, and/or interactivity, and each has distinctive technology features. In the sections that follow, we take a closer look at these components.

Big data platform

A big data platform (BDP) is a distributed file system and compute engine that can be used to facilitate data movement and processing. BDPs contain what we call a big data core—a computer cluster with distributed data storage and computing power. Advancements in big data technologies have enabled BDCs to function as a platform for additional types of computing, some of which (like query engines) can specifically support data interactivity.

Traditionally, the big data core file system can use techniques such as replication and sharding (database partitioning that separates very large databases into smaller, faster, more easily managed parts) to accelerate and scale data storage. Additionally, these techniques can help strengthen processing capabilities. Newer additions enable more powerful use of the core memory as a high-speed datastore, supporting improved data movement, processing, and interactivity. These improvements allow for in-memory computing on an existing computer cluster. Moreover, streaming technologies added to a the core can enable real-time complex event processing, and in-memory analytics technologies support better data interactivity.

Additional enhancements to big data core's focus on creating fast and familiar interfaces with data on the cluster. Typically, the core stores semi-structured data (such as XML and JSON) and unstructured data (for instance, word documents, pdfs, audio files, and videos) and requires map/reduce functionality to read. Query engine software enables the creation of structured data tables in the core and common query functionality (such as SQL).

Figure 1: Big data platform

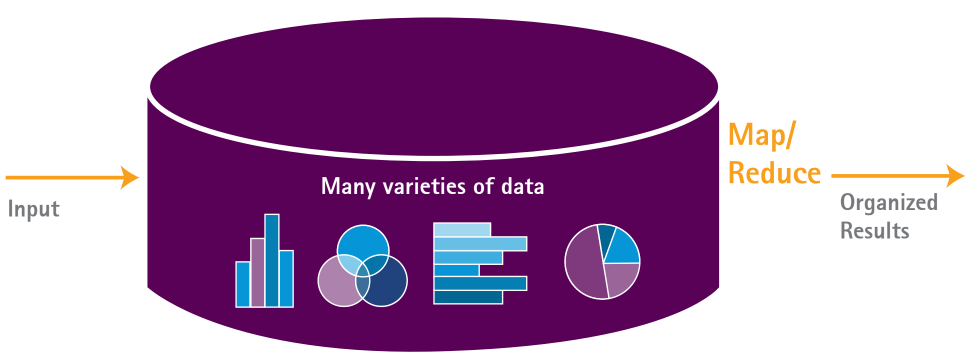

Ingestion

Ingestion is all about collecting, capturing, and moving data from its sources to underlying repositories where users can process it. Traditional ingestion was done in an extract-transform-load (ETL) method aimed at ensuring organized and complete data. Modern data infrastructure is less concerned about the structure of the data as it enters the system and more about making sure the data is collected. Modern techniques act on streaming data, such as continuous clicks on a website, and involves queues (processing of the data in the appropriate order).

As noted earlier, organizations need a mechanism for capturing data from multiple external sources (each of which might deliver data in different formats and might have different requirements) and quickly transporting the data to a place where users can access it for processing. The data can be static and reside in a repository external to the organization's data infrastructure—or it may be generated in real time by the external source. Ingestion solutions offer mechanisms for accessing and using data in both scenarios.

In this "pub-sub" system, the producer of the data publishes it from the source to a buffer or channel (data holding area). The subscriber (user or consumer) of the data picks it up from there. A queuing mechanism allows data to be buffered while the system waits for producers and consumers to take their respective actions. The speed of data producers' and consumers' actions determines the size of the buffer and the queue.

Robust ingestion supports data acceleration by enabling large amounts of data to be collected and stored quickly.

Figure 2: Ingestion

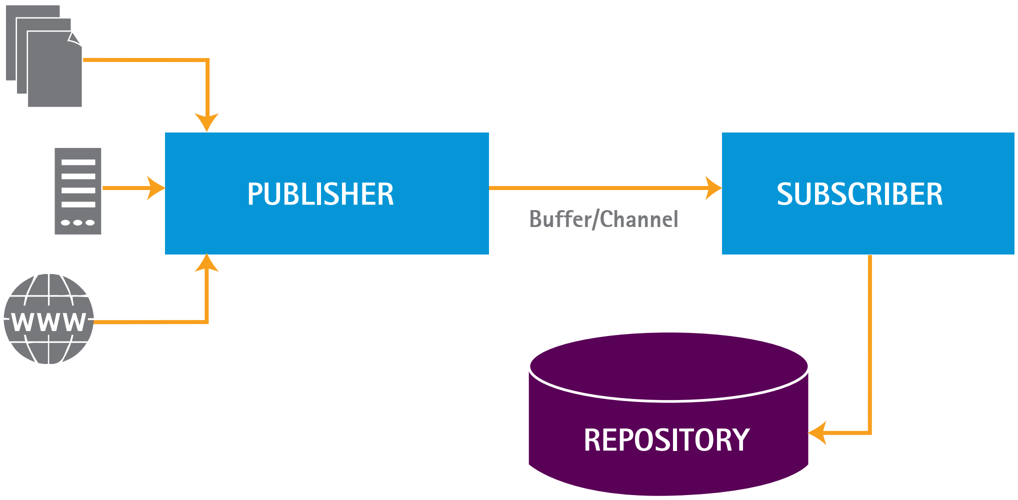

Complex event processing

Complex event processing (CEP) is a method of tracking and analyzing (processing) streams of data about events (such as click streams or video feeds) and deriving a conclusion from them. A quick example is validating security events against previously verified breaches of information in real time to assess new threats. Complex event processing combines data from multiple sources to infer events or patterns that suggest more complicated circumstances. It aims to identify meaningful events (such as opportunities or threats) and enable organizations to respond to them as quickly as possible.

Complex event processing is particularly valuable for performing real-time analytics and driving real-time insights. As fresh data streams in from its sources, these engines perform the initial pre-processing and transformations to:

- Count parts of the data and use the totals to expedite future processing of data batches, by combining historical data with fresh data.
- Match the data against pre-determined patterns as well as infer new patterns in the data.
- Trigger events and actions based on the detected patterns, delivering real-time insights to decision makers.

The primary advantage of CEP is the immediacy of the insights and actions it facilitates, compared to users having to wait for an overnight batch-processing job to be completed. The increased processing speed stems from the fact that data movement and processing occur in parallel, backed by in-memory computations. These solutions differ from ingestion solutions in that they have added processing power to perform calculations on the initial data before it is consumed by the datastore or file system.

Figure 3: Complex event processing

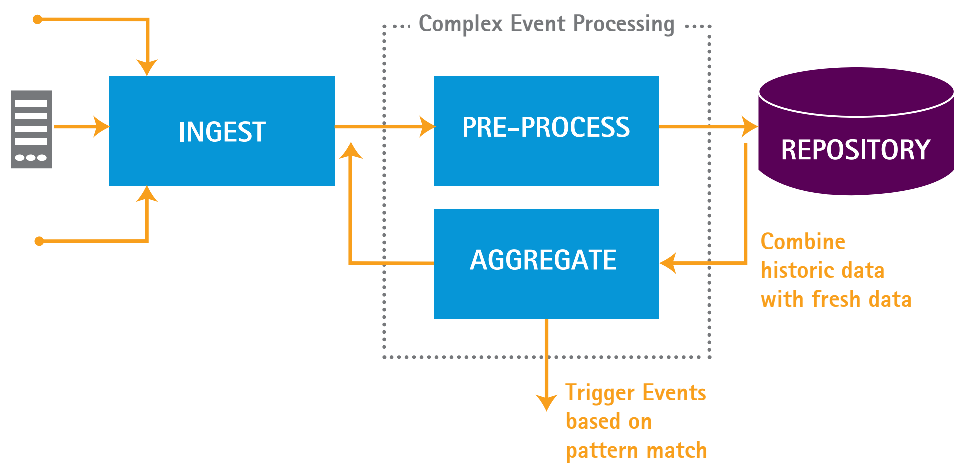

In-memory database

An in-memory database (IMDB) is a database management system that relies primarily on main memory for computer data storage. It differs from database management systems that use a disk storage mechanism. In-memory databases are faster because the internal algorithms are simpler and execute fewer central processing unit instructions. Moreover, accessing data in memory eliminates the "seek time" involved in querying data on disk storage, thus providing speedier and more predictable performance.

Because IMDBs constrain the entire database and the applications to a single address space, they reduce the complexity of data management. Any data can be accessed within just microseconds. IMDBs are not new, but decreases in RAM prices and steady increases in server RAM capacity have made them highly cost-effective options.

Figure 4: In-memory database

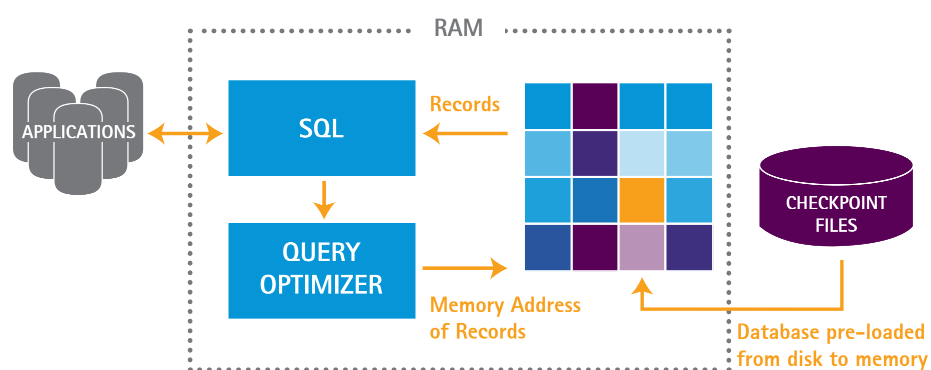

Cache clusters

Cache clusters are clusters of servers in which memory is managed by central software designed to transfer the load from upstream data sources, such as databases, to applications and users. Cache clusters are typically maintained in memory and can offer high-speed access to frequently accessed data. They sit between the data source and data consumer. The clusters are used when there is an extremely high volume of reads from multiple sources of data that does not change often, or when a database is stored on disk where seek time can be sub-optimal.

Figure 5: Cache cluster

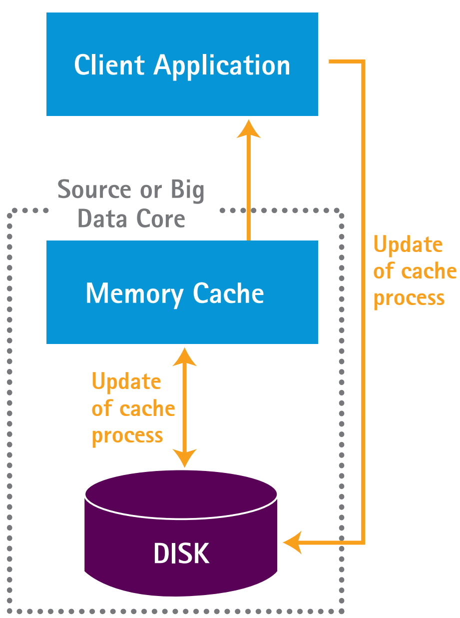

Cache clusters perform caching operations on a large scale. Traditionally they accommodate simple operations such as reading and writing values. Often they are populated when a query is sent from a data consumer to a data source. The results from the data source are then stored in the cache cluster. That way, if the same query comes in again, it does not have to go all the way back to the data source for retrieval by the data consumer. Query "receipts" build up over time in the cluster. When a data consumer requests data stored in the cluster, then the cluster responds by going to the data source—unless specific parameters are met (typically, time since the last refresh). Pre-populating data into a cache cluster (also known as "warming it") with data that is known to be frequently accessed can decrease stress on underlying systems after a system restart. Data grids take caching a step forward by adding support for more complex query operations and certain types of massively parallel processing (MPP) computations.

Appliance

An appliance is a prepackaged or preconfigured set of hardware (servers, memory, storage, and input/output channels), software (operating system, database management system, and administrative management software), and support services. It is sold as a unit, typically with built-in hardware redundancy, which helps confirm that the device will remain available in the event of a component failure. An appliance may have a common database for both online transaction processing and online analytical processing). This reduces delays in data movement, processing, and interactivity—thus supporting data acceleration.

Figure 6: Appliances

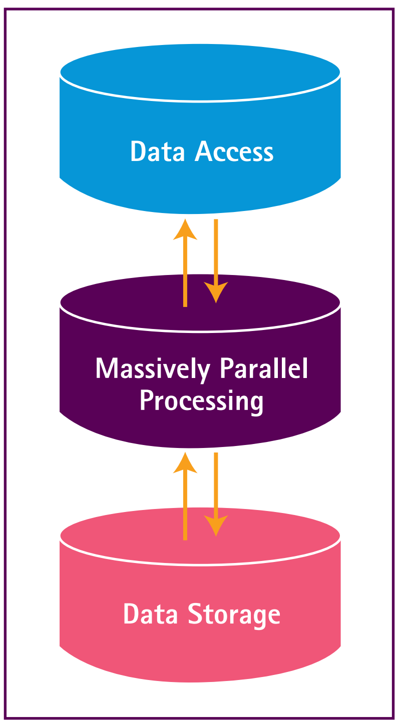

Appliances leverage technologies similar to big data cores that provide processing parallelism. By using MPP architecture, appliances can support high-performing, faster running databases and scale up when loads increase.

High-performing databases running on a cluster of servers are complicated to implement and require specialized knowledge of system, database, and storage management. In organizations that lack such knowledge, business or IT staff may be leery of adopting such databases. System maintenance and software updating are also highly time consuming for system administrators working with such databases. For such organizations, appliances offer an easier way to get the benefits of high-performing databases while avoiding the challenges. Most appliances provide the infrastructure and tools needed to build high-performance applications—including anything from core database technology and real-time replication services to lifecycle management and data provisioning.

On the hardware side of appliances, "custom silicon"—circuit boards that are not available for use outside of the appliance—offers valuable benefits. An example is the use of custom silicon in application-specific integrated circuits, which enable developers to create unique solutions tailored to their specific needs. Custom silicon also enables development on devices optimized for specific use cases without the cost of individually developing all the underlying intellectual property. For example, custom silicon for network optimization provides a unique solution that integrates embedded logic, memory, serializer/deserializer technology, networking cores, and processor cores—all of which can squeeze additional performance gains out of the appliance, providing advantages over non-custom solutions.

Thanks to these advanced capabilities, appliances can support and perform complex calculations on massive amounts of data from across an enterprise. Decision-makers can thus analyze huge volumes of data at unprecedented response times with remarkable flexibility, without the need for constant support and hand-holding from vendors. For many organizations, this "plug-and-play" aspect of appliances holds considerable appeal.

Architecture components and technology features

Component	Technology Features
Big data platform	• Distributed computing • In-memory • Streaming • Optimized network
Ingestion	• Distributed computing • In-memory • Streaming
Complex event processing	• Distributed computing • In-memory • Streaming
In-memory database	• Distributed computing • In-memory
Cache clusters	• In-memory
Appliances	• Distributed computing • In-memory • Optimized network • Custom silicon

Combining components to create solutions

The architecture components described above cannot function in isolation to support data acceleration. Instead, they must "play well" with each other, capitalizing on one another's advantages. In this section, we explore four fundamental technology stacks that meet these imperatives. We use an incremental, add-on approach to show how these stacks (all of which include common layers) are built to enable data movement, processing, and interactivity.

Figure 7: Solutions landscape

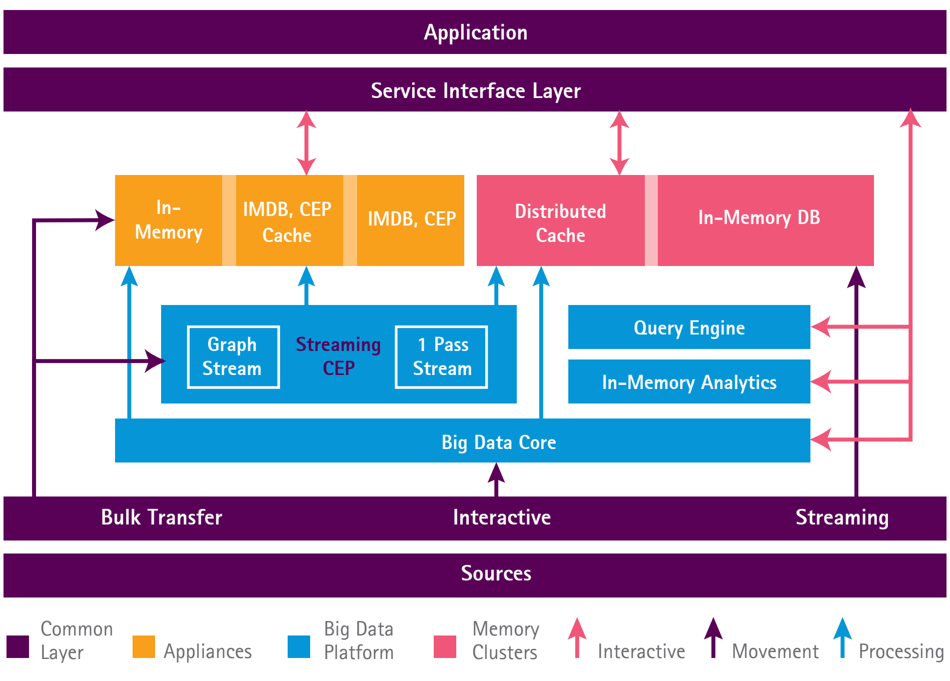

Common layers

For any use of enterprise data, the data comes into the organization's data infrastructure, it is processed, and users interact with it by submitting queries and receiving responses with which to make decisions. Where the data originates and where it leaves the system are what we call common layers that can be envisioned as above and below the architecture components discussed earlier. Data sources feed the ingestion component at the bottom of the system. At the top of the system, users and applications interact with the data.

Problem Types

	Movement "Drinking from the fire hose and can't miss a drop"		Processing "Need to process large amounts of data faster"	Interactivity "Urgent question that must be answered immediately"
	ETL	Streaming		
1. Appliance Only	Basic	Enhanced	Enhanced	Enhanced
2. BDP to Appliance	Basic	Enhanced	Enhanced	Enhanced
3. Streaming to Appliance	Basic	Enhanced+	Enhanced+	Enhanced
4. BDP Only	Basic	Enhanced	Basic	Basic
5. Streaming to BDP	Basic	Enhanced+	Enhanced	Basic
6. BDP with In-Memory Analytics	Basic	Enhanced	Enhanced	Enhanced
7. Streaming to BDP with In-Mem Analytics	Basic	Enhanced+	Enhanced+	Enhanced
8. BDP with Query Engine	Basic	Enhanced	Basic	Enhanced
9. Streaming to BDP with Query Engine	Basic	Enhanced+	Enhanced+	Enhanced
10. Distributed Cache Cluster Only	Basic	Enhanced	Basic	Enhanced
11. BDP to Cache Cluster	Basic	Enhanced	Basic	Enhanced
12. In-Memory Database Cluster Only	Basic	Enhanced	Basic	Enhanced
13. BDP to In-Memory Database Cluster	Basic	Enhanced	Basic	Enhanced
14. Streaming to In-Memory Database Cluster	Basic	Enhanced+	Enhanced+	Enhanced
	Complex event processing can enhance streaming ingestion		Complex event processing can increase speed by pre-processing data	Caches and In-Memory databases can enable real-time interactivity

Big data platform

Big data core only

In this scenario, data typically enters the computer cluster through a batch or streaming process. However, events are not processed immediately. The big data core is job based—calculations are scheduled to run at a certain interval rather than done in real time. It leverages replication and distributed parallel processing on large datasets, which enables advanced analytics. Applications and services may access the core directly and deliver improved performance on large, unstructured datasets. It is quickly becoming the de facto standard; therefore, we consider this technology the benchmark for exceptional data movement, processing, and interactivity.

Figure 8: Big data core only

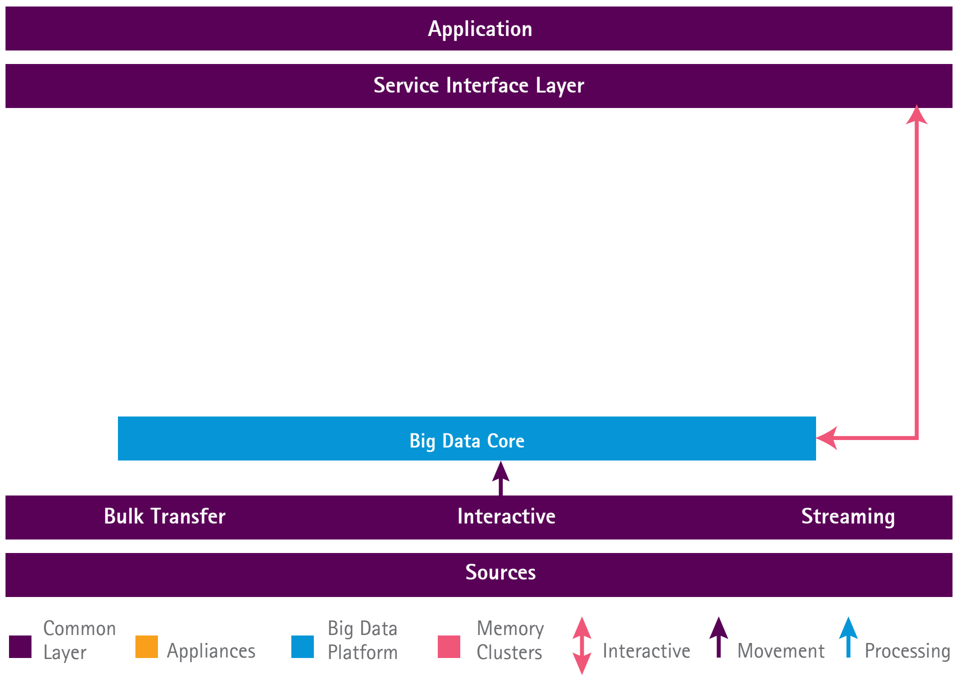

Big data core and complex event processing

Adding CEP enhances a big data core's processing capabilities, as ingesting data through a CEP enables real-time detection of patterns in the data and event triggering. This functionality is useful for correlating real-time information with an analytic model; for example, when an organization wants to be alerted to a security event in real-time. By leveraging processing capabilities on an existing dataset on the core, data scientists can create a machine learning model and transfer it to the CEP unit. Rather than waiting for the core's jobs to execute, the CEP can take action immediately, drawing on criteria generated in the model. It thus enhances the big data core's processing capabilities and augments interactivity components by enabling real-time animated dashboards.

Figure 9: Big data core and complex event processing

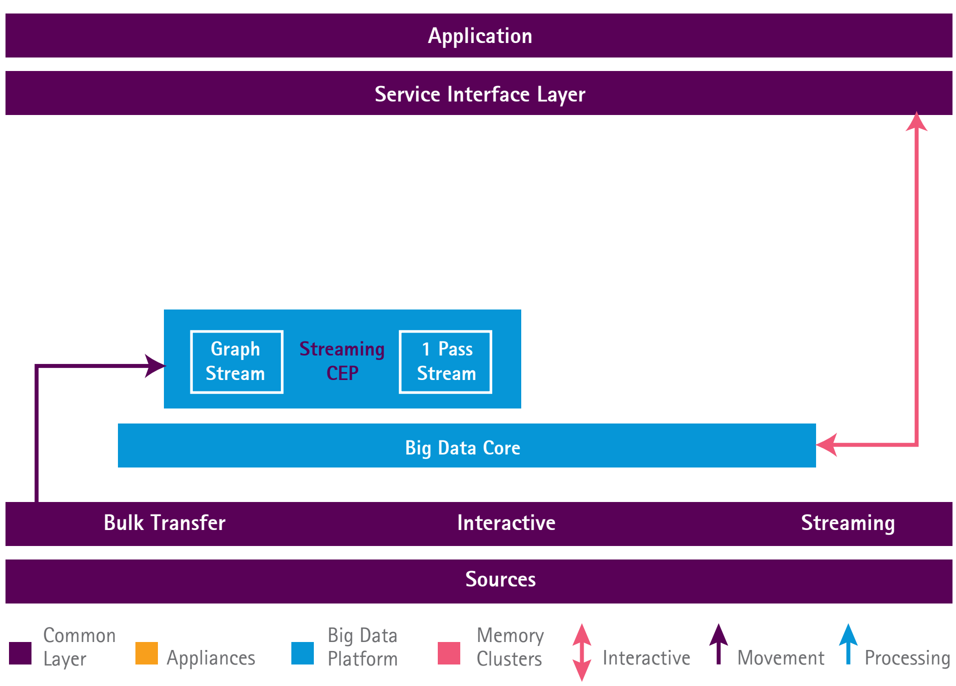

Big data core and in-memory database

A big data core's traditional analytical capabilities stem from the ability to leverage the distributed computing power of commodity hardware. As such computing power has strengthened over time, so have the applications that use this hardware. For example, IMDB analytics software can be added to the big data core to improve computations by placing key data in RAM on nodes in the cluster, avoiding the problem of slow disk operations. Additionally, new software products promise to help reduce required computing time by several orders of magnitude.

Figure 10: Big data core and in-memory database

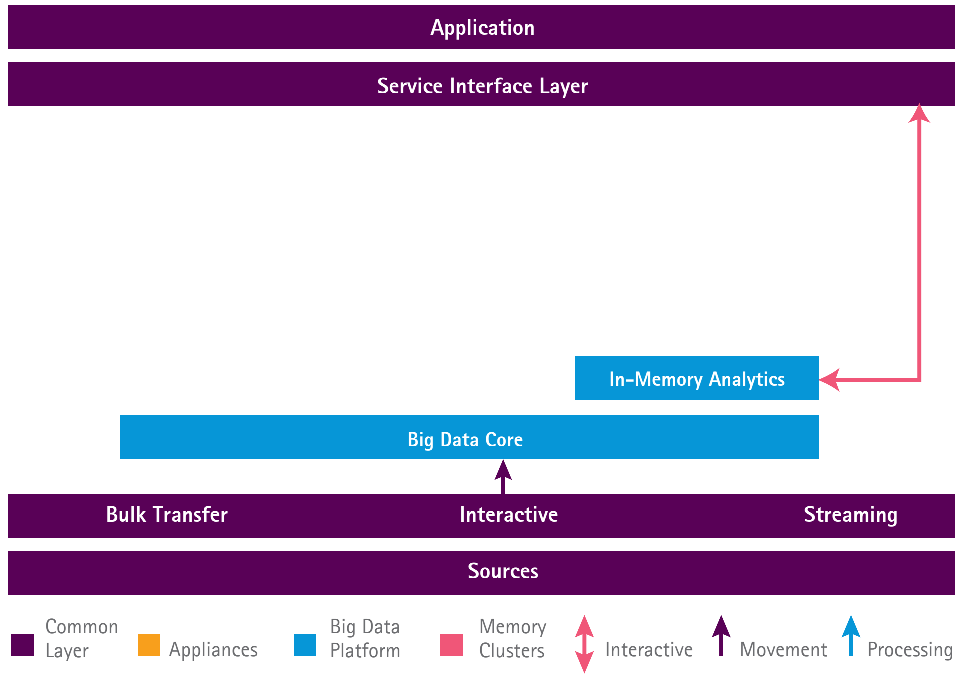

Figure 11: Big data core, complex event processing and in-memory database analytics

Joining these three enables faster processing and interactivity.

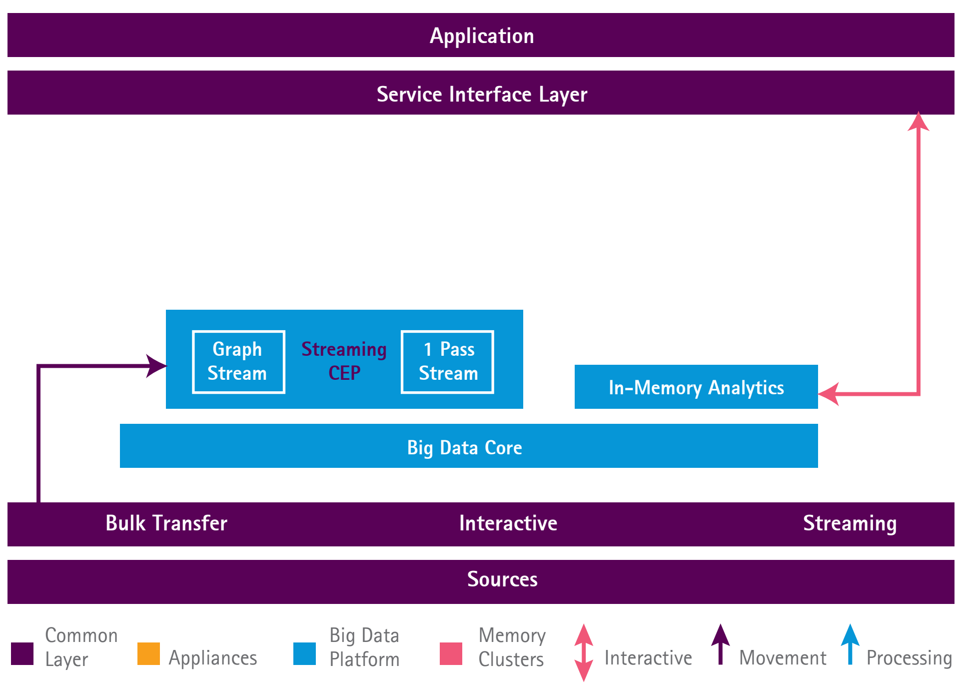

Big data core with query engine

Adding query engine technology to a BDC opens common interfaces for applications to access data with less delay. This makes big data more immediately accessible to users and applications.

Figure 12: Big data core with query engine

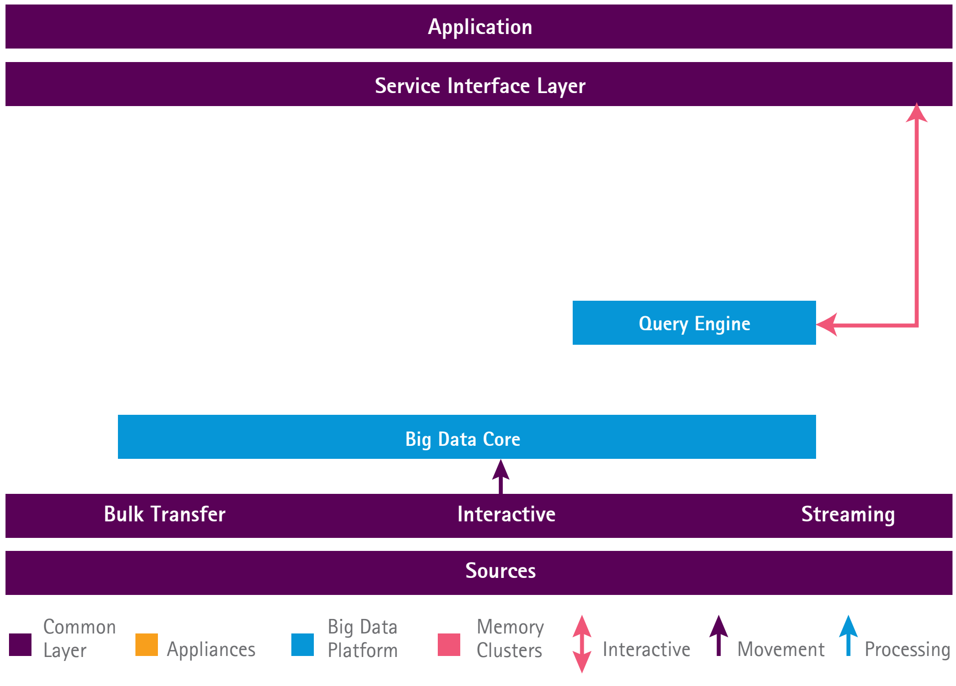

Big data core, complex event processing and query engine

With certain technologies, CEP results can be accessed directly from query engine technologies, fostering improved data movement, processing, and interactivity.

Figure 13: Big data core, complex event processing and query engine

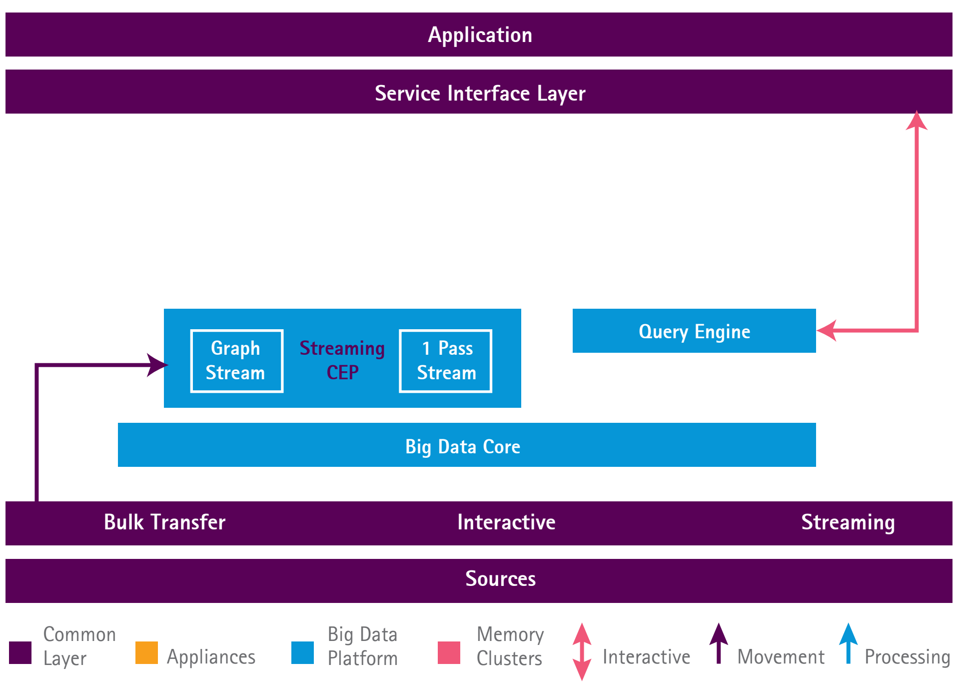

In-memory database

In-memory database cluster only

To facilitate data movement, processing, and interactivity, data from different external sources is either streamed-in or bulk transferred directly to the IMDB. The processing includes simple to complex calculations, model executions, and statistical comparisons—all of which take place in-memory within the database. Without the need to page information in or out of memory, the IMDB improves read and write performance, speeding up data processing. Users and applications can directly query the IMDB as they would any other database for specific information. These queries typically use SQL-like structures, making the data easily accessible. Additionally, queries are optimized in-memory. For instance, when returning data, the computers in the cluster with the most resources available will be selected to respond. Such optimization provides faster response times.

Figure 14: In-memory database cluster only

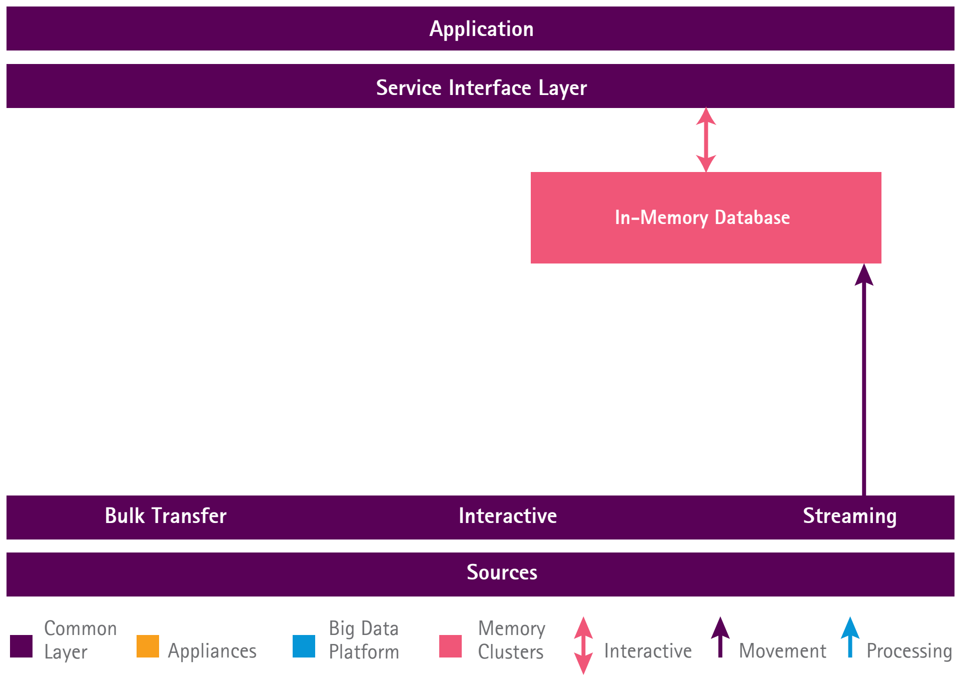

In-memory database cluster and big data platform

Data is first ingested into the system either as a bulk transfer or as a streaming process through the platform. The data is stored on the platform's distributed file system. This approach enables some pre-processing to take place on the platform before data is transferred to the IMDB. Such pre-calculation speeds up future processing. The database does the bulk of the analytical processing completely in-memory, delivering faster read and write performance. As with IMDB cluster only, queries requested by an application are optimized and executed in the in-memory database, and results are quickly returned to the application.

Figure 15: In-memory database cluster and big data platform

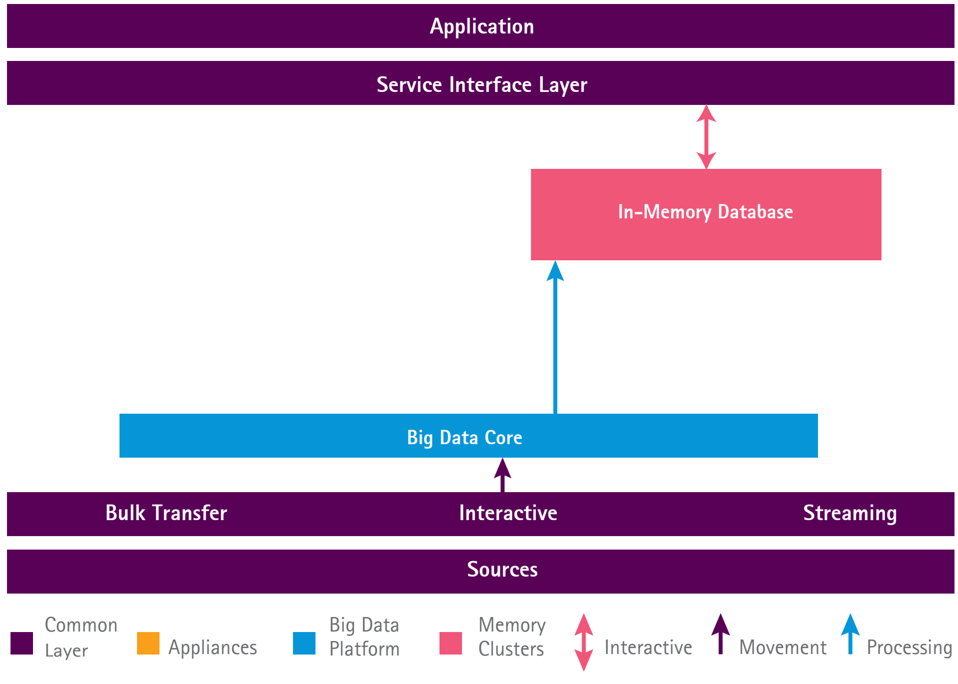

In-Memory database cluster and complex event processing

Data from sources is first ingested into the system through a CEP engine. The bulk of the analytical processing, including model execution and statistical comparison, takes place in the IMDB. Queries requested by an application are executed in the database and returned to the application for faster interactivity.

Figure 16: In-Memory database cluster and complex event processing

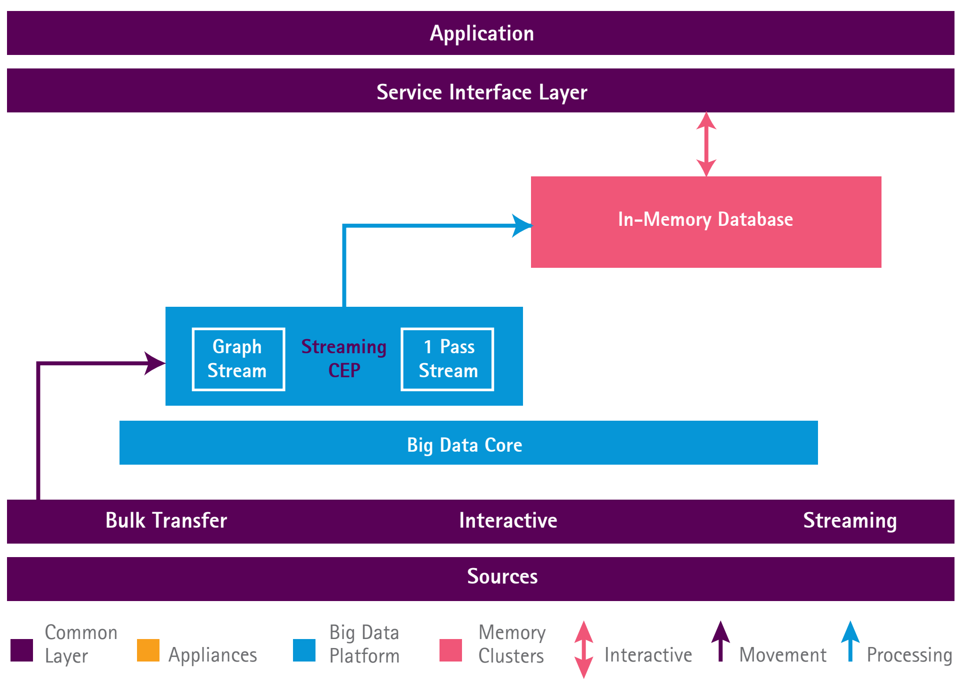

Distributed cache

Cache Only

This stack consists of a simple caching framework sitting atop the data source repository and connected to an application. The application retrieves data. To optimize query time, the cache must be "tuned" such that data subsets that are most relevant to the application are placed in the cache. Since the cache simply stores data, processing of data falls to the application, which may cause slower processing speeds and delays.

Figure 17: Cache Only

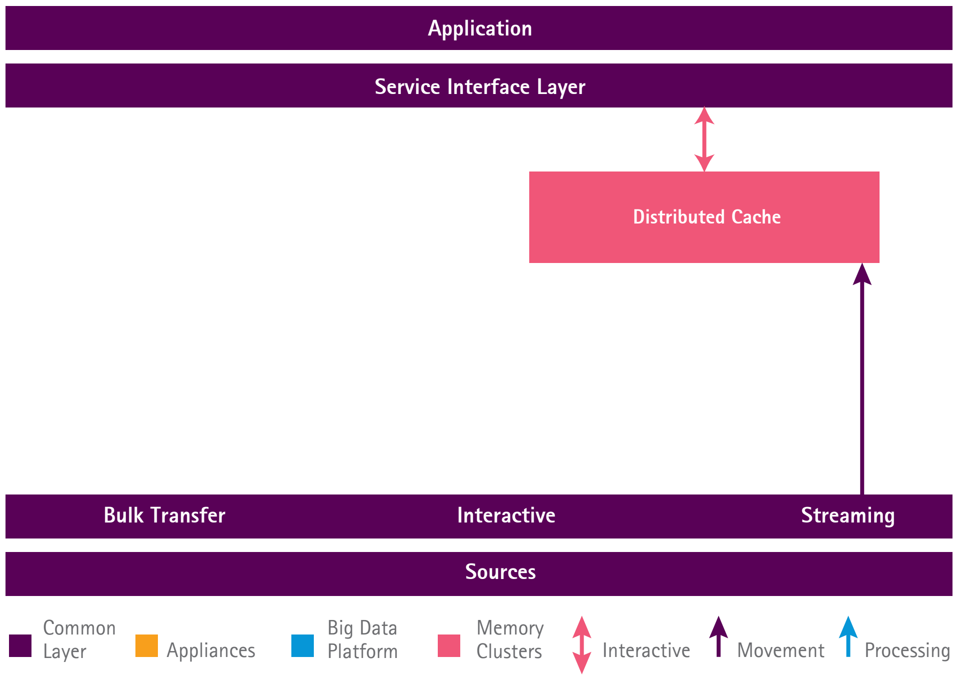

Cache, application and big data platform

The platform ingests data from the source and does the bulk of the processing before loading a data subset into the cache. This moves the burden of data processing from the application to the the platform, which can run complex analytic processes on large datasets more efficiently. A cache sits atop the the platform, which feeds the application's query results.

Figure 18: Cache, application and big data platform

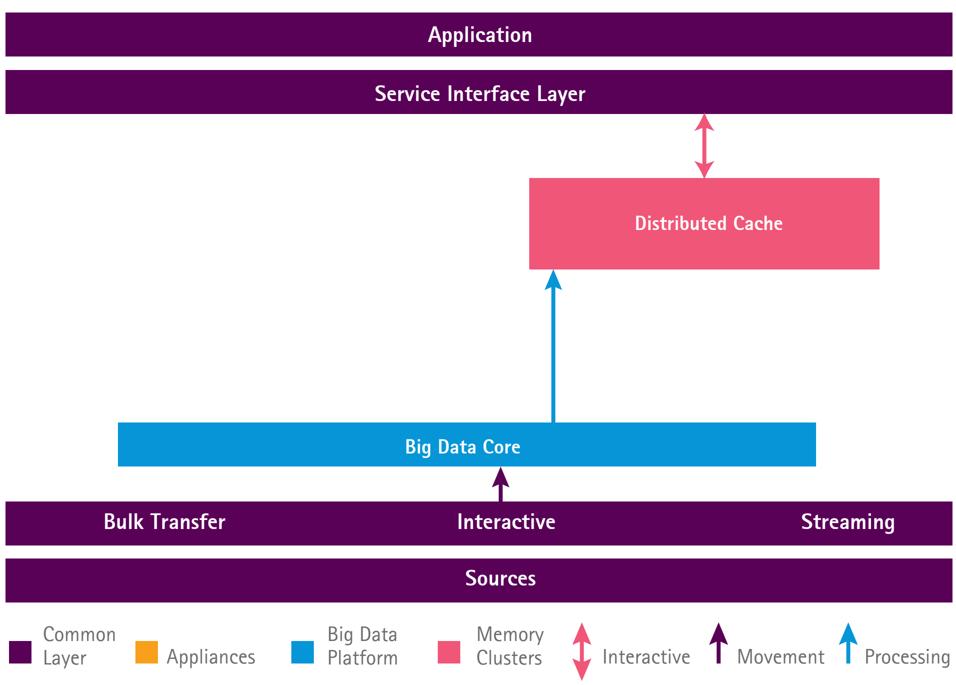

Appliance

Appliance only

Data from the source streams directly into the appliance, which completes processing, analytics, and calculations. The application "talks" directly to the appliance for query requests.

Figure 19: Appliance only

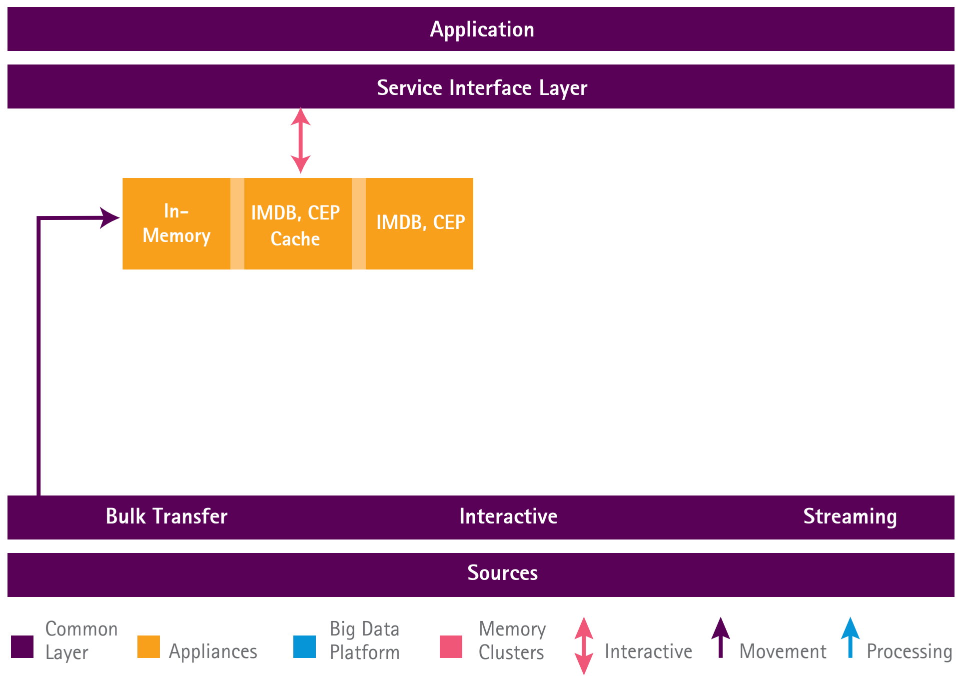

Appliance and big data platform

Data from the source is imported and stored within the platform. The stack may process the data inside the platform before transferring it to the appliance to achieve faster processing speed. The application can also directly talk to the appliance for query requests.

Figure 20: Appliance and big data platform

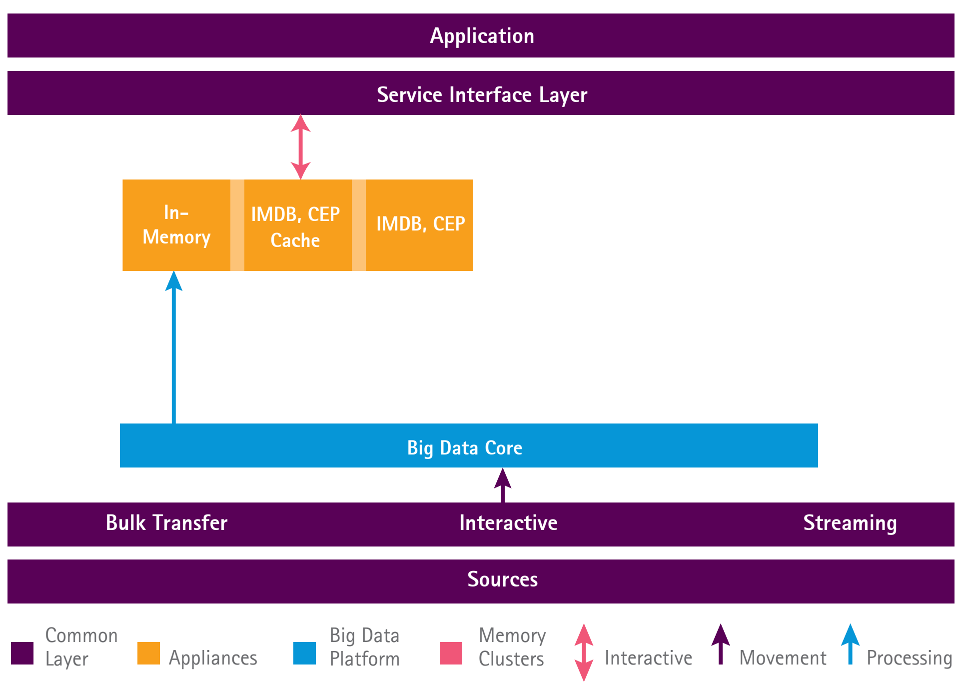

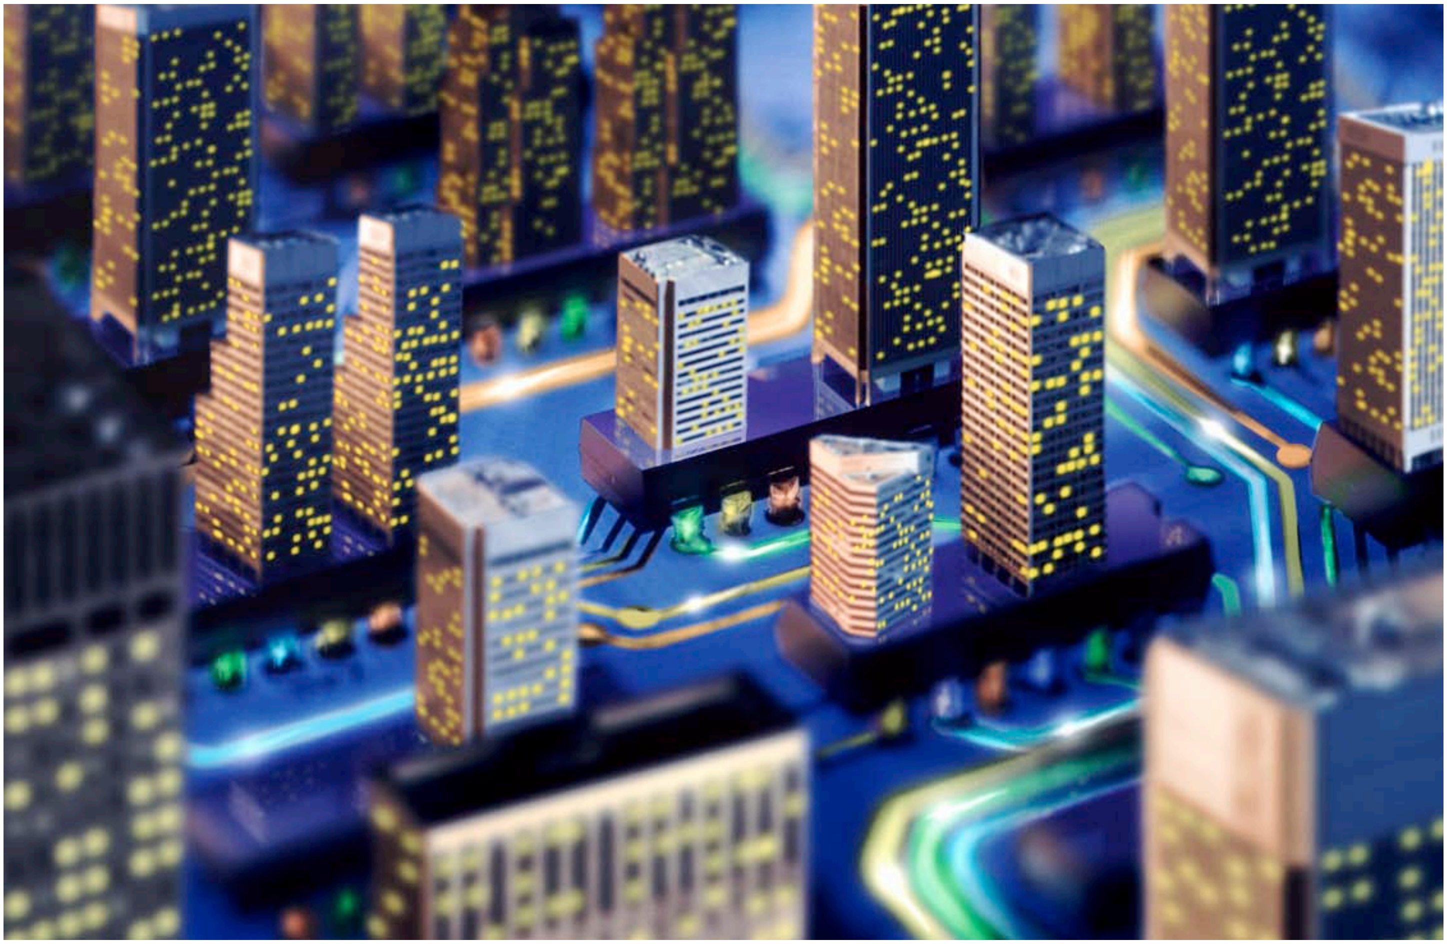

Appliance and streaming

Data from the source is first imported and stored within the big data platform through streaming. The stack may also process the data inside the platform before transferring it to the appliance to achieve optimal processing speed. The application can directory query the appliance for information.

Figure 21: Appliance and streaming

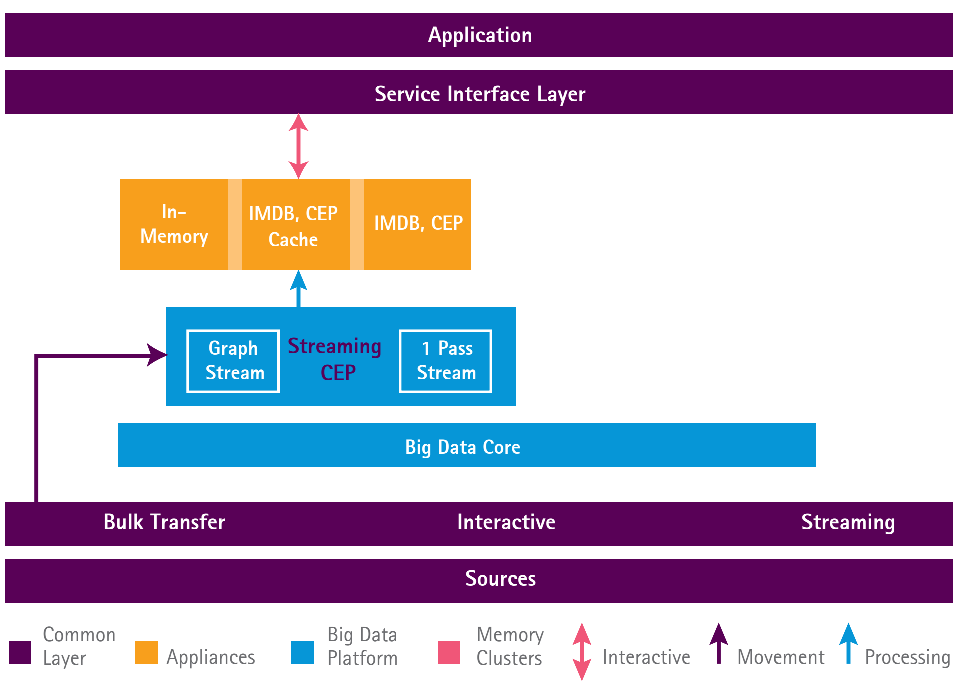

Conclusion

To gain a competitive advantage from its enterprise data, an organization must be able to generate business insights from the data. The barrier of entry is at an all-time low, thanks to the advent of big data. However, big data also brings challenges. To surmount them, organizations need to establish a data supply chain that (among other accomplishments) accelerates data movement, processing, and interactivity—enabling decision makers to more swiftly capture and act on insights from their data as well as achieve returns on their analytics investments.

Yet the landscape of solutions aimed at fostering data acceleration has grown more complex than ever. To build the right data acceleration architecture, executives must first deepen their understanding of the challenges inherent in data movement, processing, and interactivity. Then they have to familiarize themselves with the architectural components now available on the market—each of which supports data acceleration in unique ways.

But even this understanding is not enough: The architectural components deliver maximum value only when they are combined in ways that capitalize on their complementary advantages. By exploring four potential architecture configurations, executives can initiate valuable discussion about which configurations may be best for their organization's needs. Equally important, they can bring a more informed perspective to discussions with vendors about prospective solutions.

This point of view provides an overview that executives can use as a starting point to both understand this evolving landscape and begin familiarizing themselves with appropriate architectural solutions to address their business needs and achieve analytics ROI.

Next steps

To begin building a data supply chain strategy that supports data acceleration in your organization:

- Inventory your data. Start with your most frequently accessed and time-relevant data. This will be given first access to your data platform and accelerated on the platform.
- Identify inefficient processes. Look for any manual, time-consuming data curation processes, such as tagging or cleansing. These may be candidates for replacement with machine learning algorithms.
- Identify data silos. Along with silos, identify corresponding data needs that are currently unmet across the business.
- Simplify data access. Create a strategy for standardizing data access via the data platform. Solutions may be hybrid, combining traditional middleware and API management, or even a platform-as-a-service offering.
- Prioritize individual data supply chains. Prioritizing helps you develop a road map for implementing the data supply chain at scale.
- Consider external data sources. Look outside your organization for external data sources that can be incorporated to complement existing data and help generate more complete insights.
- Choose the data acceleration technology stack for your data and research the optimal deployment methods.

For more information, please contact:

Vincent Dell'Anno
Managing Director, Accenture Analytics - Big Data
vincent.u.dellanno@accenture.com

Carl M. Dukatz
R&D Manager, Accenture Technology Labs - Data Insights
carl.m.dukatz@accenture.com

About Accenture Analytics

Accenture Analytics, part of Accenture Digital, delivers insight-driven outcomes at scale to help organizations improve their performance. With deep industry, functional, business process and technical experience, Accenture Analytics develops innovative consulting and outsourcing services for clients to help ensure they receive returns on their analytics investments. For more information follow us @ISpeakAnalytics and visit www.accenture.com/analytics.

About Accenture Technology Labs

Accenture Technology Labs, the dedicated technology research and development (R&D) organization within Accenture, has been turning technology innovation into business results for more than 20 years. Our R&D team explores new and emerging technologies to create a vision of how technology will shape the future and invent the next wave of cutting-edge business solutions. Working closely with Accenture's global network of specialists, Accenture Technology Labs help clients innovate to achieve high performance. The Labs are located in Silicon Valley, California; Sophia Antipolis, France; Arlington, Virginia; Beijing, China and Bangalore, India. For more information follow us @AccentureLabs and visit www.accenture.com/accenturetechlabs.

About Accenture

Accenture is a global management consulting, technology services and outsourcing company, with approximately 281,000 people serving clients in more than 120 countries. Combining unparalleled experience, comprehensive capabilities across all industries and business functions, and extensive research on the world's most successful companies, Accenture collaborates with clients to help them become high-performance businesses and governments. The company generated net revenues of US$28.6 billion for the fiscal year ended Aug. 31, 2013. Its home page is www.accenture.com.