

Chapter 2

New Requirements for Enterprise Computing

When thinking about developing a completely new database management system for enterprise computing we first need to clarify whether there are significant opportunities for improvement in the area. The fact that modern companies have changed dramatically towards a more data-driven model gives the strong indication that this might be the case. For a typical enterprise computing use-case like assembly line management it now became viable to use sensors giving instant feedback on various parameters and producing huge amounts of data. Furthermore companies process previously available data at a much larger scale, e.g. competitor behavior, price trends, etc. to support management decisions. All this data volume will continue to increase in the future.

There are two major requirements for a modern database management system:

- Data from various sources need to be combined in a single database management system, and
- This data needs to be analyzed in real-time to support interactive decision-making.

The following sections outline typical use cases for modern enterprises and derive associated requirements for an entirely new enterprise data management system.

2.1 Processing of Event Data

Event data influences enterprises more and more today. It is characterized by the following aspects:

- Each event dataset itself is small (some bytes or kilobytes) compared to the size of traditional enterprise data, such as all data contained in a single sales order, and

- The number of generated events for a specific entity is high compared to the amount of entities, e.g. hundreds or thousand events are generated for a single product.

In the next sections we will give some examples use-cases related to the processing of event data in modern enterprises.

2.1.1 Sensor Data

Sensors are used to supervise the function of more and more systems today. One example is the tracking and tracing of sensitive goods, such as pharmaceuticals, clothes, or spare parts. Hereby packages are equipped with Radio-Frequency Identification (RFID) tags or two-dimensional bar codes, the so-called data matrix. Each product is virtually represented by an Electronic Product Code (EPC), which describes the manufacturer of a product, the product category, and a unique serial number. As a result, each product can be uniquely identified by its EPC code. In contrast, traditional one-dimensional bar codes can only be used for identification of classes of products due to their limited domain set. Once a product passes through a reader gate, a reading event is captured. The reading event consists of the current reading location, timestamp, the current business step, e.g. receiving, unpacking, repacking or shipping, and further related details. All events are stored in decentralized event repositories.

Real-Time Tracking of Pharmaceuticals

For example, approximately 15 billion pharmaceuticals are produced in Europe that are only available on prescription. Tracking any of them results in approx. 8,000 read event notifications per second. These events build the basis for anti-counterfeiting techniques. For example, the route of a specific pharmaceutical can be reconstructed by analyzing all relevant reading events. The in-memory technology enables tracing of 10 billion events in less than 100 ms.

Formula One Racing Cars

Formula one racing cars are also generating excessive sensor data. These sports cars are equipped with up to 600 individual sensors, each recording tens to hundreds of events per second. Capturing sensor data for a 2 h race produces giga- or even terabytes of sensor data depending on their granularity. The challenge is to capture, process, and analyze the acquired data during the race to optimize the car parameters instantly, e.g. to detect part faults, optimize fuel consumption or top speed.

2.1.2 Analysis of Game Events

Personalized content in online games is a success factor for the gaming industry. Browser games can generate a steady stream of tens of thousands events per second, such as player movements, transfer of virtual goods, or general game statistics. Traditional databases do not support processing of these huge amounts of data in an interactive way, e.g. join and full table scans require complex index structures or data warehouse systems optimized to return some selected aspects in a very fast way. However, individual and flexible queries from developers or marketing experts cannot be answered interactively.

Gamers tend to spend money when virtual goods or promotions are provided in a critical game state, e.g. a lost adventure or a long-running level that needs to be passed. In-game trade promotion management needs to analyze the user data, the current in-game events, and external details, e.g. current discount prices.

In-memory database technology is used to conduct in-game trade promotions and, at the same time, conduct A/B testing. To this end, the gamers are divided into two segments. The promotion is applied to one group. Since the feedback of the users is analyzed in real-time, the decision to roll-out a huge promotion can be taken within seconds after the small test group accepted the promotion.

Furthermore, in-memory technology improves discovery of target groups and testing of beta features, real-time prediction, and evaluation of advert placement.

2.2 Combination of Structured and Unstructured Data

First we want to understand structured data as any kind of data that is stored in a format which is automatically processed by computers. Examples for structured data are ERP data stored in relational database tables, tree structures, arrays, etc. Following that we want to understand partially or mostly unstructured data that cannot easily be processed automatically, e.g. all data that is available as raw documents, such as videos or photos. In addition, any kind of unformatted text, such as freely entered text in a text field, document, spreadsheet, or database, is considered as unstructured data unless a data model for its interpretation is available, e.g. a possible semantic ontology.

For years enterprise data management focused on structured data only. Structured data is stored in a relational database format using tables with specific attributes. However, many documents, papers, reports, web sites, etc. are only available in an unstructured format, e.g. text documents. Information within these documents is typically identified via the document's meta data. A detailed search within the content of these documents or the extraction of specific facts is however not possible by using the meta data. As a result, there is a need to harvest information buried within unstructured enterprise data. Searching any kind of data—structured or unstructured—needs to be equally flexible and fast.

2.2.1 Patient Data

In the course of the patient treatment process, e.g. in hospitals, structured and unstructured data is generated. Examples of unstructured data are diagnosis reports, histologies, and tumor documentations. Examples of structured data are results of the erythrogram, blood pressure, temperature measurements, or the patient's gender. The in-memory technology enables the combination of both classes of patient data with additional external sources, such as clinical trials, pharmacological combinations or side-effects. As a result, physicians can prove their hypotheses by interactively combing data and reduce necessary manual and time-consuming searches. Physicians are able to access all relevant patient data and to take their decision on latest available patient details.

Due to their high fluctuation of unexpected events, such as emergencies or delayed surgeries, the daily time schedule of physicians is very time-optimized. In addition to certain technical requirements of their tools, they have also very strict response time requirements. For example, the HANA Oncolyzer, an application for physicians and researchers was designed for mobile devices. The mobile application supports the use-as-you-go factor, i.e., the required patient data is available at any location on the hospital campus and the physician is no longer forced to go to a certain desktop computer for checking a certain aspect. In addition, if the required detail is not available in real-time for the physician, she/he will no longer use the application. Thus, all analyses performed by the in-memory database are running on a server landscape in the IT department while the mobile application is the remote user interface for it.

Having the flexibility to request arbitrary analyses and getting the results within milliseconds back to the mobile application makes in-memory technology a perfect technology for the requirements of physicians. Furthermore, the mobility aspect bridges the gap between the IT department where data is stored and the physician that visits multiple work places throughout the hospital every day.

2.2.2 Airplane Maintenance Reports

Airport maintenance logs are documented during exchange of any spare parts. These reports contain structured data, such as date and time of the replacement or order number of the spare part, and unstructured data, e.g. kind of damage, location, and observations in the spacial context of the part. By combining structured and unstructured data, in-memory technology supports the detection of correlations, e.g. how often a specific part was replaced in a specific aircraft or location. As a result, maintenance managers are able to discover risks for damages before a certain risk for human-beings occurs.

2.3 Social Networks and the Web

Social networks are very popular today. Meanwhile, the time when they were only used to update friends about current activities are long gone. Nowadays, they are also used by enterprises for global branding, marketing and recruiting.

Additionally, they generate a huge amount of data, e.g. Twitter deals with one billion new tweets in five days. This data is analyzed, e.g. to detect messages about a new product, competitor activities, or to prevent service abuses. Combining social media data with external details, e.g. sales campaigns or seasonal weather details, market trends for certain products or product classes can be derived. These insights are valuable, e.g. for marketing campaigns or even to control the manufacturing rate.

Another example for extracting business relevant information from the Web is monitoring search terms. The search engine Google analyzes regional and global search trends. For example, searches for “influenza” and flu related terms can be interpreted as a indicator for a spread out of the influenza disease. By combining location data and search terms, Google is able to draw a map of regions that might be affected from an influenza epidemic.

2.4 Operating Cloud Environments

Operating software systems in the cloud require a good data integration strategy. Assume, you process all your company’s human resources (HR) tasks in an on-demand HR system provided by provider A. Consider a change of the provider to cloud provider B. Of course, a standardized data format for HR records can be used to export data from A and import it into B. However, what happens if there is no compatible standard for your application? Then, the data exported from A needs to be migrated, respectively remodeled, before it can be imported into B. Data transformation is a complex and time-consuming task which often has to be done manually due to the required knowledge about source and target formats and many exceptions which have to be solved separately.

In-memory technology provides a transparent view concept. Views describe how input values are transformed to the desired output format. The required transformations are performed automatically when the view is called. For example, consider the attributes `first name` and `last name` that need to be transformed into a single attribute `contact name`. A possible view `contact name` performs the concatenation of both attributes by performing `concat(first name, last name)`.

Thus, in-memory technology does not change the input data, while offering the required data formats by transparent processing of the view functions. This enables a transparent data integration compared to the traditional Extract Transform and Load (ETL) process used for Business Intelligence (BI) systems.




Fig. 2.1 Inversion of corporate structures

2.5 Mobile Applications

The wide spread of mobile applications fundamentally changed the way enterprises process information. First (BI) systems were designed to provide detailed business insights for CEOs and controllers only. Nowadays, every employee is gaining insights by the use of BI systems. However, for decades information retrieval was bound to stationary desktop computers. With the wide-spread of mobile devices, e.g. PDAs, smartphones, etc., even field workers are able to analyze sales reports or retrieve the latest sales funnel for a certain product or region.

Figure 2.1 depicts the new design of BI systems, which is no longer top-down but bottom-up. Modern BI systems provide all required information to sales representatives directly talking to customers. Thus, customers and sales representatives build the top of the inverted pyramid.

In-memory databases build the foundation for this new corporate structure. On mobile devices, people are eager to get a response within a few seconds [Oul05, OTRK05, RO05]. With the ability to perform complex and freely formulated queries with a sub-second response, in-memory databases can revolutionize the way employees communicate with customers. An example of the radical improvements through in-memory databases is the dunning run. A traditional dunning process took 20 min on an average SAP system, but by rewriting the dunning run on in-memory technology it now takes less than one second.

2.6 Production and Distribution Planning

Two further prominent use cases for in-memory databases are complex and long-running processes such as production planning and availability checking.

2.6.1 *Production Planning*

Production planning identifies the current demand for certain products and consequently adjusts the production rate. It analyzes several indicators, such as the users' historic buying behavior, upcoming promotions, stock levels at manufacturers and whole-sellers. Production planning algorithms are complex due to required calculations, which are comparable to those found in BI systems. With an in-memory database, these calculations are now performed directly on latest transactional data. Thus, algorithms are more accurate with respect to current stock levels or production issues, allowing faster reactions to unexpected incidents.

2.6.2 *Available-to-Promise Check*

The Available-to-Promise (ATP) check validates the availability of certain goods. It analyzes whether the amount of sold and manufactured goods are in balance. With raising numbers of products and sold goods, the complexity of the check increases. In certain situations it can be advantageous to withdraw already agreed goods from certain customers and reschedule them to customers with a higher priority. ATP checks can also take additional data into account, e.g. fees for delayed or canceled deliveries or costs for express delivery if the manufacturer is not able to send out all goods in time.

Due to the long processing time, ATP checks are executed on top of pre-aggregated totals, e.g. stock level aggregates per day. Using in-memory databases enables ATP checks to be performed on the latest data without using pre-aggregated totals. Thus, manufacturing and rescheduling decisions can be taken on real-time data. Furthermore, removing aggregates simplifies the overall system architecture significantly whilst adding flexibility.

2.7 Mathematical and Scientific Applications

Mathematics is the most powerful, omnipresent tool we have. It can be used to put a man on the moon or to calculate the numbers of atoms in a molecule. It can be found in the geometric transformations governing the canons of J.S. Bach or in the catalogue of symmetries in Islamic art. Nowadays the field of applied mathematics and scientific computing is heavily dependent on computer aided data processing due to the ever-growing amounts of data. Researchers and scientists ideally want to process such big amounts of data as fast as possible and see their results in seconds—not hours or days—in order to decide how to proceed. As mentioned throughout the book, these requirements are in the context of business

data processing addressed by using an in-memory database system: in this section we want to demonstrate that the advancements in the business data processing sector are also beneficial for mathematical applications. Although current mathematical and scientific applications are not directly in the focus of enterprise computing, they may pave the way for tomorrow's new business segments.

2.7.1 Proteomics Research

Proteomics research focuses on the study of the proteome, a term referring to the totality of proteins expressed by a genome, cell, tissue or organism at a given time. The proteome is highly dynamic, changing constantly as a response to the needs and state of the organism. Factors such as cancer or other diseases, can also change the composition of the proteome. Thus, analysis of these changes can be used i.e. for cancer diagnosis.

The prominent way to analyse the (human) proteome is by means of mass spectrometry, a technique that measures the masses and concentration of peptides in a biological sample and exports the data in form of a spectrum. Mass spectrometer emit their measurements as raw data then processed in an analysis pipeline consisting of numerous preprocessing (smoothing, filtering) and analysis steps. To identify signals in the proteome for cancer, multiple samples from different cohorts derived in clinical studies have to be compared, which adds selection and join operations to access cohort data from clinical information systems to the pipeline.

Depending on the approach, the analysis pipeline for proteomics mass spectrum data represents a data-flow program, where each step applies mathematical transformations on the data flowing through the pipeline. Each operator step can be parametrized and replaced with a different implementation or method to solve the transformation, e.g. different algorithms for minimization problems can be applied, or different model types for classifications can be embedded into the pipeline.

Proteomics research benefits from the integration of relational database operators and arbitrary mathematical operations in SanssouciDB, allowing the proteomics researcher to compose complex analysis pipelines in one environment using clinical patient data as well as raw spectrum data inside one database. Since every parameter and algorithm influences the accuracy of the resulting statistical model, iterative tuning of such pipelines in real-time is a requirement in this field: here the fast traversal of data and algorithm execution is crucial.

2.7.2 Graph Data Processing

As motivated in the previous subsection, graph processing presents a major challenge for modern database management systems: there are two very different access

types that are typically supported by different storage engines. Graphs are a generic representation suitable for almost any information. Even relational databases and the data stored in tables are basically graphs of related nodes with attributes.

The most common operations on graphs are twofold: On the one hand is graph exploration with the goal to traverse and explore singular paths and on the other hand is graph analytics trying to explore and analyze the whole graph or multiple instances of a similar graph.

In native graph databases explorative traversals are directly executed on a graph structure while in in-memory databases multi-way joins are required to follow a single path. However, the same is true for analytical queries. The advantage of a modern in-memory database is that multiple explorative graph queries can be executed at once using joins and aggregations. For analytical queries, the queries can leverage the capabilities of the parallel join execution engine inside SanssouciDB.

Compared to disk-based databases, in-memory column stores have the advantage of fast scanning and data traversal of the necessary attributes which make it possible to execute complex queries. Depending on the application use case it is not always advisable to use relational databases for graph processing. Especially breadth-first and depth-first graph search algorithms can become very expensive due to repetitive joins. On the other hand there are algorithms to resemble graph structures that can be executed very efficiently, like aggregation on hierarchical tree structures used for bill of materials explosion.

2.7.3 SanssouciDB Data Scientist

The SanssouciDB Data Scientist tool supports users who want to implement and interact with complex analysis pipelines that go beyond standard analytical operations consisting of complex mathematical and machine learning operators orchestrated to complex pipelines.

The tool provides a coherent development and modeling environment integrating the graphical data mining modeling paradigm and the source code development paradigm in one framework. The user can graphically design data flow programs to define her analysis pipeline consisting of relational, domain specific native operators, as well as custom operators which e.g. can be directly implemented by the user in the R programming language for statistical computing.

Modeled pipelines are compiled on the fly and executed on SanssouciDB's calculation engine leveraging data locality and transparent parallel execution. Users can interactively manipulate their pipeline, by tuning operator parameters, exchanging complete operators or manipulate code of custom operators within the tool, while the pipeline is compiled to SanssouciDB on the fly, and thereby showing the resulting effects immediately.

2.8 Self Test Questions

1. Data explosion

Consider the formula 1 race car tracking example, with each race car having 512 sensors, each sensor records 32 events per second whereby each event is 64 byte in size.

How much data is produced by a F1 team, if a team has two cars in the race and the race takes two hours?

Please use the following unit conversions: 1,000 byte = 1 KB, 1,000 KB = 1 MB, 1,000 MB = 1 GB.

- (a) 14 GB
- (b) 15.1 GB
- (c) 32 GB
- (d) 7.7 GB

References

- [OTRK05] A. Oulasvirta, S. Tamminen, V. Roto, J. Kuorelahti, Interaction in 4-second bursts: the fragmented nature of attentional resources in mobile hci, in *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI '05* (ACM, New York, 2005), pp. 919–928
- [Oul05] A. Oulasvirta, The fragmentation of attention in mobile interaction, and what to do with it. *Interactions* **12**(6), 16–18 (2005)
- [RO05] V. Roto, A. Oulasvirta, Need for non-visual feedback with long response times in mobile hci, in *Special Interest Tracks and Posters of the 14th International Conference on World Wide Web, WWW '05* (ACM, New York, 2005), pp. 775–781

A Course in In-Memory Data Management
The Inner Mechanics of In-Memory Databases
Plattner, H.
2014, XXIII, 312 p. 125 illus., 62 illus. in color.,
Hardcover
ISBN: 978-3-642-55269-4