

Data lake &
analytics factory
by bullion

Big Data at the heart of Customer Transformation Challenges

With digital transformation, analyzing huge sets of data helps businesses identify risks, detect new opportunities and enables fast-decision making for quick action.

Analysts say that organizations integrating coherent information management infrastructure will outperform their industry peers financially by more than 20%.

Hence Big Data is at the heart of Customer Transformation challenges and enables organizations to create value through numerous and very diverse use cases in terms of:

- ▲ **Customer Experience** : anticipate needs, demands and behaviors thanks to a more comprehensive data set
- ▲ **Business Reinvention** : grow core business with differentiated products and services through advanced analytics
- ▲ **Operational Excellence** : gain efficiency and agility with data driven business processes
- ▲ **Trust & Compliance** : unleash the power of analytics to protect your assets and improve compliance with regulations

But to fully benefit from it you need to implement the relevant infrastructure, totally modular, without modification, from the early stage of use case development to fully industrialized platform: a data lake with analytics tools.

The simple revamping of old technologies and processes does not suffice to succeed since analytics reinvention is needed. And adapting your infrastructure for data management is a challenge since this solution must support you in creating value from these data but also fit the needs of different populations in the enterprise:

- ▣ **Business users** – they might be using today Business Intelligence solutions but need to have access to more and better information based on data with different variety and complexity.
- ▣ **IT and operations** – they must react very rapidly to business users requests so the solution must be simple and easy to implement but also flexible since applications must evolve rapidly with business changes. This means also being able to deploy some applications in various environments (on their own systems but also on various cloud providers).
- ▣ **Data scientists** – they are new players who, in connection with business users and IT, are looking for the use cases that will create value for the enterprise through relevant algorithms. Data scientists can spend up to 80% of their time in operational tasks rather than developing advanced analytics. Hence, the solution has to integrate all the analytics tools to be used in a straightforward manner.

The answer is a solution with a new way to collect, store, compute, distribute and analyze data:

- ▣ An industrialized analytics software suite
- ▣ A “state of the art” data lake, vast storage space, based on Hadoop technologies with the tools to deploy and manage it easily
- ▣ Validated and pre-integrated appliance with hadoop distribution
- ▣ One support for all Big Data hardware and software

This is now possible with the Data lake & analytics factory Appliance based on a software suite conceived by Atos and a data lake fully integrated and powered by bullion. This offer is a new expansion to Atos Codex, fully integrated and cross market data analytics end-to-end solution that enables organizations to maximize the value of their data.

<p>CUSTOMER EXPERIENCE</p> <ul style="list-style-type: none"> • Customer 360° • Churn Detection • Dynamic Offering 	<p>OPERATIONAL EXCELLENCE</p> <ul style="list-style-type: none"> • Predictive Maintenance • Supply-chain Optimization • Dynamic Forecasting
<p>BUSINESS REINVENTION</p> <ul style="list-style-type: none"> • Advertising scheduling • Grow core business with new products or service 	<p>TRUST & COMPLIANCE</p> <ul style="list-style-type: none"> • Fraud Prevention • Threat Management • Surveillance

Some Customer Transformation Challenges use cases

The Data lake & analytics factory software suite

Thanks to the selection of different pieces of software and the thoughtful integration, validation and automation, the Data lake & analytics factory software suite dramatically simplifies the jump into the info centric era. This software suite proposes:

- ▣ A state of the art data lake, based on Hadoop technology, with all the tools to deploy it in a few seconds and easily manage all its lifecycle.
- ▣ Combined products to very quickly deploy data refining and analytics capacities
- ▣ A low-level software layer enabling flexibility and scalability to sustain needs and data growth even on various environments (on premise or in the cloud) for temporary capacity add-on.

For business and end-users:

The software suite for the Data lake & analytics factory is the go-to element especially for end-users, as it is the set which delivers the added value and makes data accessible and usable. The software suite will provide:



Data visualization: huge sets of data presented with a very graphic perspective deeply enhance the

brain capacity to get the underlying information. Different visualization possibilities enable different analysis perspective of the very same data set. It is then possible to build advanced dashboards without any code line or filter and seek data to know what happens in real time or at a specific time frame.



Geo positioning: as a complement to the data visualization, the capacity to watch how

things, events and information are linked geographically can provide an unmatched perspective and capacity to

monitor existing elements, detect new opportunities or identify threats.



Advanced search: though searching seems a rather basic feature, it becomes more and more important

when the data set grows. Analysts for example estimate that business workers lose each year hundreds of hours just seeking information, files and data they already have had at hand. The capacity to search in a broader spectrum is a way to leverage data capital.



Analyze: with data sources multiplying exponentially, accurate refining and analysis is more and more essential

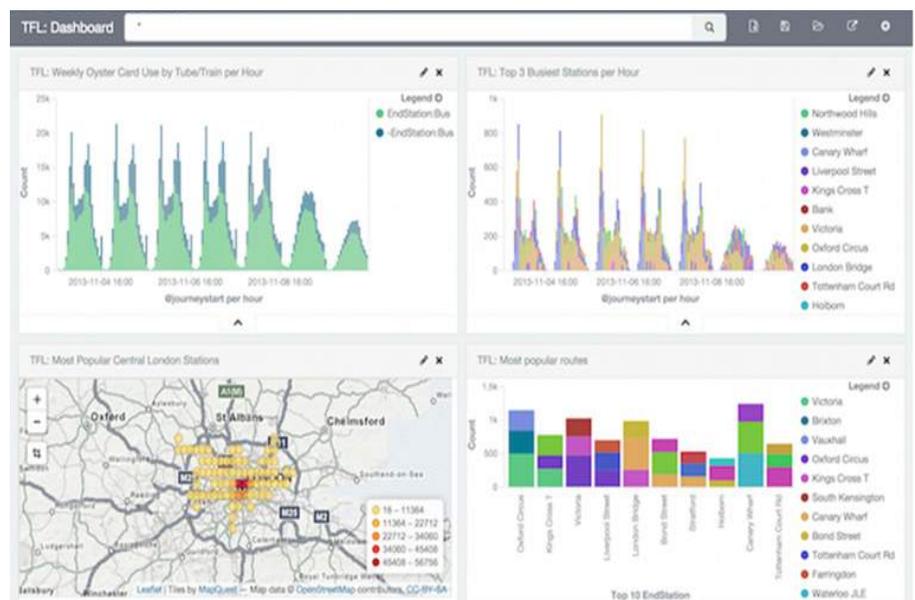
to provide useful information. Making sense of machine data, sensors and such is considered as one of the future next gold mine. One way to extract value from it is to process and analyze all the log files. The software suite proposes the Smart Log Analytics tool to simplify log analysis.

Focus on log analysis

To analyze user behavior and apps performance, help troubleshoot applications, find security audit leads, or perform metrology and capacity planning, the usage of applicative log files is today a vastly underestimated corporate asset.

Analysis and indexation of log files are possible with the "ELK" (Elastic Search - Logstash - Kibana) stack. Elastic search indexes data as they are created, Logstash analyzes data in real time and Kibana provides dashboards to analyze them. This stack is well known, designed for data correlation and easily scalable.

On top of this stack, Bull has integrated additional components such as dedicated blueprints in order to automate deployment, environment set up and life cycle management..



Example of combining various log elements in a single dashboard

The Data lake & analytics factory software suite

For data scientists:

The Data lake & analytics factory delivers very interesting features also for non end-users. Indeed data scientists often spend 70-80% of their time on IT operations. With automation of the life cycle of an Hadoop based environment, providing pre-packaged and integrated software suite, there is no need to install or upgrade specific components, one by one, anymore. This will definitely save time for data scientists and thus leverage their skills to create more value for the organization.

In addition, the software suite enables to deploy Hadoop based environments in self-service mode and proposes catalogs of self-service components. Specialized blueprints enable a quick and easy deployment of consistent software stacks on adapted resources. These capacities are built on a Hadoop platform (based on MapR® or Hortonworks® and later on Cloudera® distribution).

There are also lambda architecture abilities, drill down software stacks,

messaging platform to send the time series capacity for connecting hundreds of data sources. Bull added a studio for application design to easily create, deploy even share or resell home-made applications.

For IT department and operations:

Simplified deployment, administration as well as capacities to manage self-service resources and their lifecycle will not only be an asset for the data scientists but also for the IT people. Big Data products are often based on a very large ecosystem; managing the components selection, upgrade, support and replacement is often a nightmare for IT departments. The software suite solves a big issue thanks to pre-integrated and validated components and enables sound, robust and easy deployment on the datacenter thanks to virtualization technology. The virtualization will indeed greatly enhance Data lake and analytics capacities management, efficiency and

make provisioning or de-provisioning extremely fast and easy.

The IT department security is also enhanced on two different aspects: multi-tenancy and platform availability. The multi-tenancy will likely interest organizations hosting data and IT capacities for third parties and will enable efficient resource sharing while providing dedicated and isolated environments. The platform availability is enhanced by the decorrelation of logical resources and hardware platform like any virtualized projects.

In addition, a low-level software layer enables a complete separation of the software suite and the associated hardware. This software enables deployment inside or outside the enterprise (public cloud) with preserved data security policy. Without any adjustment or development, extra capacities can be added very simply in the resource pool (should they be temporary or not).



Example of REAL TIME IT production dashboard

Data lake & analytics factory by bullion

A Big Data Appliance to make it simple

The Data lake & analytics by bullion pushes further the benefits of the software suite and also provides some appliance-specific benefits such as:

- ▣ Pre-integrated hardware and pre-loaded software
- ▣ One stop shopping and one stop support
- ▣ Predictable performance thanks to validated and benched components
- ▣ Easy upgrades, both at hardware and software level.

The underlying bullion server, delivers an exceptional processing power performance and unmatched modularity, enabling fine tuning of hardware needs thanks to its unique architecture and virtualization technology.

This enables to power both the Data lake and analytics environments on the very same hardware and thus improves dramatically the flexibility, security and TCO.

Flexibility

Bullion servers can be used in a scale-up model when compute or memory needs arise (to deploy more analytics virtual machines for example). They provide from 2 to 16 CPUs, up to 24TB of RAM in a single server and a few hundreds of TB of storage thanks to the attached EBODs. Depending on the organization needs, the appliance can follow a scale up or scale-out track, making the appliance by bullion even more flexible.

The capacity to deploy or re allocate virtual machines in minutes with a whole and consistent stack thanks to dedicated blueprints is a clear advantage especially since needs might evolve fast.



Security

The appliance provides advanced security at three major levels.

The first level that might be of a critical importance (compared to cloud solutions) is that all data are staged on premise and do not transit through public networks.

The second one is linked to the virtual machine motion feature. In case of planned maintenance or unexpected disaster (flood, storm...), being able to move a service/virtual machine across servers, even to a remote data center ensures perfect availability of the system. This feature made possible by virtualization technology is an important asset of the solution.

The third level of security is closely tied to the bullion server. In critical environments, bullion offers unique RAS and warning features that anticipate potential breakdowns. Combined with a wide range of diagnostic resources, this allows you to anticipate problems and react appropriately enabling exceptional **Quality of Service**.

TCO enhancements

The overall infrastructure improves TCO at various levels.

First of all, this appliance is designed to optimize resource allocation thanks to its **unique architecture and virtualization technology**. Bullion hosts both an infrastructure as a service (IAAS) and the virtual machines to run the Data lake & analytics on the very same hardware which improves the TCO.

In addition, the bullion server is powered by high-end processors meaning **better consolidation capabilities** reducing the number of servers, simplifying network topology and reducing licensing costs. Moreover the architecture will allow additional TCO improvements linked to the **upgrade path**. Creating a global resource pool enables to optimize resource allocation since there is a disconnection between CPU, memory and storage capacities. As time goes, data volume will grow and needs will evolve. To address data volume growth, the scalability proposed by bullion architecture enables to add extra compute, memory and/or storage capacity with a very controlled pace. To deploy a new Big Data application, the virtualization technology will enable accelerated and simplified deployment by allocating free resources from the virtualized pool.

Atos, your partner to grow in the Big Data

Experts at your side

Creating the technological requirements is essential but the best technology delivers the desired benefits with the relevant expertise.

At Atos, secure Data Management is at the center of our offering and of our “Journey 2018” strategic vision. As a result, Atos makes unique investment in Analytics & Big Data with 4.500 Big Data experts & 86.000 business technologist to unleash the value of data. We are today the only European player to master the whole data value chain, from infrastructure to software and services.

Atos has developed a methodology and recently launched its new Codex solution to support customers all along the evolution of their project and at any or all stages of the transformation for business analytics.

The services range from digital transformation strategy & consulting, use case business modeling, data science expertise (define use cases, choose algorithms), agile analytics deployment and ongoing evolution management.

Two Global Delivery Centers have been installed to offer competitive pricing and share competencies and skills across Atos. Hundreds of solution architects and data scientists work in the Atos Advanced Analytics Competence Centers on data centric applications and algorithms to deliver tangible business value from client data through Atos Codex solutions.

Launch a successful Big Data POC with us

Whatever the status of your Big Data strategy or programs, we can therefore support you on your project. Participate to one of our Innovation workshops, consult our experts for an opportunity scan or maturity assessment and once the project is identified, an interesting approach is to experiment the solution through a Proof of Concept with the support of our experts.

About Atos & Bull

Atos SE (Societas Europaea) is a leader in digital services with pro forma annual revenue of circa € 12 billion and circa 100,000 employees in 72 countries. Serving a global client base, the Group provides Consulting & Systems Integration services, Managed Services & BPO, Cloud operations, Big Data & Cyber-security solutions, as well as transactional services through Worldline, the European leader in the payments and transactional services industry. With its deep technology expertise and industry knowledge, the Group works with clients across different business sectors: Defense, Financial Services, Health, Manufacturing, Media, Utilities, Public sector, Retail, Telecommunications, and Transportation.

Atos is focused on business technology that powers progress and helps organizations to create their firm of the future. The Group is the Worldwide Information Technology Partner for the Olympic & Paralympic Games and is listed on the Euronext Paris market. Atos operates under the brands Atos, Atos Consulting, Atos Worldgrid, Bull, Canopy, Unify and Worldline.

For more information

atos.net

fr.directionmarketing@atos.net

Bull, the Atos technologies for the digital transformation

Bull is the Atos brand for its technology products and software, which are today distributed in over 50 countries worldwide. With a rich heritage of over 80 years of technological innovation, 2000 patents and a 700 strong R&D team supported by the Atos Scientific Community, it offers products and value-added software to assist clients in their digital transformation, specifically in the areas of Big Data and Cybersecurity.

Bull is the European leader in HPC and its products include bullx, the energy-efficient supercomputer; bullion, one of the most powerful x86 servers in the world developed to meet the challenges of Big Data; Evidian, the software security solutions for identity and access management; Trustway, the hardware security module and Hoox, the ultra-secure smartphone. Bull is part of Atos.

For more information

bull.com/bullion

Atos, the Atos logo, Atos Consulting, Atos Worldgrid, Bull, Canopy, Unify, Worldline registered trademarks of the Atos group. All trademarks are the property of their respective owners. Atos reserves the right to modify this document at any time without notice. Some offerings or parts of offerings described in this document may not be available locally. Please contact your local Atos office for information regarding the offerings available in your country. This document does not represent a contractual commitment. April 2016 © 2016 Atos



THIS BROCHURE IS PRINTED ON PAPER COMBINING 40% ECO-CERTIFIED
FIBERS FROM SUSTAINABLE FORESTS MANAGEMENT AND 60% RECYCLED
FIBERS IN LINE WITH CURRENT ENVIRONMENT STANDARDS (ISO 14001).