InterSystems IRIS

Data Platform for the Future

InterSystems®
Creative data technology

EXECUTIVE SUMMARY Page 03

CHALLENGES Page 04

SMART DATA FABRIC Page 06

PERFORMANCE Page 08

ANALYTICS Page 10

CONTAINERS Page 14

CUSTOMERS AND PARTNERS Page 16

FUTURE-PROOF DATA MANAGEMENT Page 17

Executive Summary

The **InterSystems IRIS** data platform lets companies take maximum advantage of their own data. It also guarantees modern data management that meets all their requirements – today and in the future. To this end, InterSystems has enhanced a number of the proven strengths its technology has to offer, including the following three: Performance, interoperability, and scalability. And three innovative additions pave the way for the future: **Adaptive Analytics**, the use of artificial intelligence **(AI)** and machine learning **(ML)**, and the combination of **containers, microservices, and API management.**

Inter-Systems IRIS offers professional users real-time insights into their own data, enabling them to make better business decisions at the right time. The first-class performance, the free and flexible scalability, as well as the open system architecture for third-party tools have a positive impact on the work of solution architects in particular. Developers also appreciate the ease of developing, deploying, and updating applications based on the combination of containers, microservices, and API management. In addition, the uniform and open development environment ensures extensive flexibility. InterSystems IRIS enables data bank administrators and data stewards to centrally monitor and manage access to data and also lets them create data models easily and quickly. When it comes to security and data protection, InterSystems IRIS follows proven industry standards. All data is protected from unauthorized access at all times.

InterSystems®
Creative data technology

Meeting the challenge of digitalization with InterSystems IRIS

Meeting the requirements of advancing digitalization – this is the challenge facing companies' development departments as well as software providers, implementation specialists, and system integrators.

Many of them already rely on InterSystems IRIS to optimize their data management so they can lay a robust, powerful foundation for digital transformation. The data platform stands out for first-class performance, interoperability, free and flexible scalability, and security. It also provides integrated capabilities for data analytics, business intelligence (BI), and the use of AI and ML. Overall, InterSystems IRIS lets companies maintain an overview of all their data in real time as well as combining, evaluating, and using it – whether structured or not.

Therefore the solution forms an ideal basis for intelligent, data-driven business models, such as offering up-to-date services in the financial industry. At the same time, it promotes the necessary change processes toward smart factories and networked logistics. Thanks to these new ideas and concepts, companies and software providers alike can maintain their competitive edge and be ready for the future.

"I have yet to experience a situation where I couldn't solve a requirement with InterSystems technology."

Ralf Spielmann, System/Database Architect and Chapter Manager for Database Development, NOVENTI Health SE

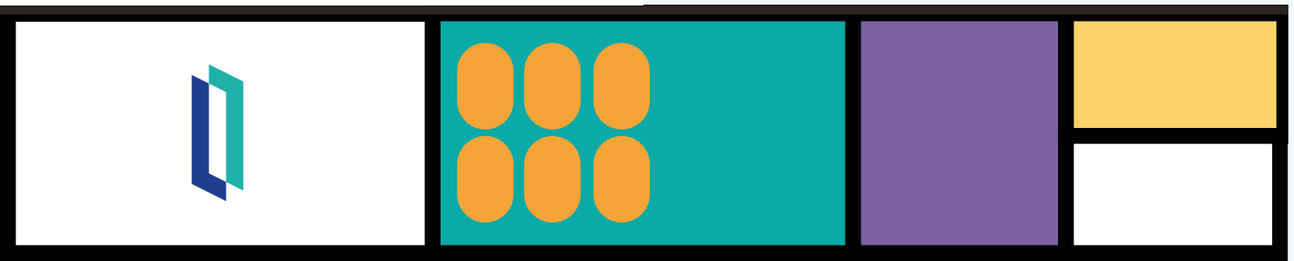

Gaining rapid, usable insights from masses of data

A central feature of digitalization is exponential growth in the volume of data. Companies in every industry need to manage this information overload in order to make informed decisions.

A company's success increasingly depends on having a direct insight into and immediate access to all relevant data at all times. InterSystems IRIS responds to this challenge by using hybrid transactional/analytical processing (HTAP) that guarantees exceptionally high performance and efficiency for multi-workloads in real time in any use case. This makes it possible to record large amounts of transactional data and perform complex analyses at the same time. Plus, InterSystems IRIS breaks down application and data silos so that the company can make maximum use of its own data. In conjunction with the AI, ML, and BI functions and tools integrated in the data platform, this enables the implementation of a smart data fabric – the evolution of the enterprise data fabric.

"We have been working with InterSystems for decades now and are impressed by the performance, interoperability, and scalability of the company's data management solutions."

Jens Bohl, Chief Information Officer, TRANSWAGGON GmbH

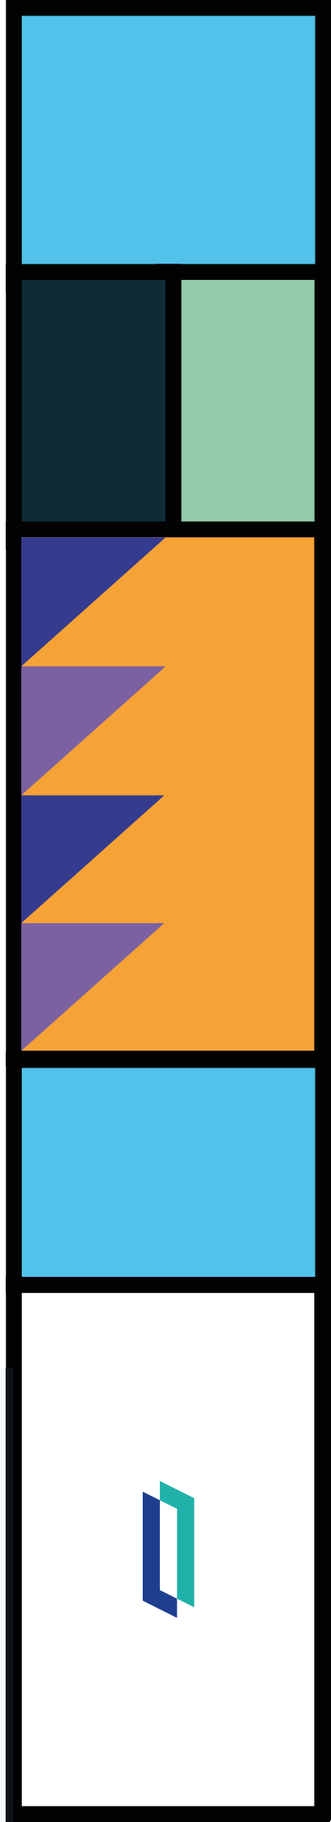

Meeting the challenge of digitalization with InterSystems IRIS

Many companies are under enormous pressure to digitalize their business processes, primarily to meet new data management requirements. They have to impress when it comes to performance, interoperability, scalability, and security and ideally simplify their analyses. Comprehensive interoperability is particularly important because it makes it possible to link all the available data sources, which in turn enables insightful analyses. One concept that meets all the requirements for cutting-edge data management is the **smart data fabric.**

A smart data fabric can be used to merge, cleanse, and harmonize all historical and current data from internal and external sources. All these processes take place regardless of the source of the data, its transmission protocol, or format. The result is comprehensive interoperability between all the company's applications, systems, and services. Given that the smart data fabric is built on top of the existing IT infrastructure, there is no need for a time-consuming, expensive redesign – known as "rip and replace." Instead, existing technologies, applications, and services can still be used, no matter where exactly they are located (they can be operated both locally and in the cloud without any issues). In addition, the smart data fabric significantly reduces the complexity of a company's IT infrastructure, which both simplifies operation and maintenance as well as decreasing the costs involved. Updating the company's data management and automating it to the greatest possible extent will also eliminate any inefficiencies in the workflow.

Integrated functions make the data fabric smart

What distinguishes the smart data fabric from the enterprise data fabric are the built-in capabilities for analyzing data, visually displaying it in BI dashboards, and using AI and ML to gain and leverage insights that are immediately relevant. Companies don't have to use tedious and time-consuming processes to involve other systems, applications, or tools in a task – instead, they can work directly with the data. **IntegratedML**, for example, provides companies with a tool that helps employees with basic SQL knowledge to implement ML scenarios independently. It has a simple user interface that makes it easier to create and execute ML models as well as monitor the result sets. A data scientist then takes care of optimizing the individual models if necessary.

As a whole, the smart data fabric enables a central view of all data in the company – a **single source of truth (SSOT)**. In addition to data discovery, data lineage, and data governance, this comprehensive transparency also makes it possible to visualize all processes via a cockpit or dashboard. As a result, business users have quick and easy access to a wide range of clean data, also known as healthy data. This data is always up-to-date, accurate, and trustworthy. Healthy data forms the basis for enabling companies with the corresponding applications to easily perform real-time analyses, to make well-informed decisions much faster, to act more quickly than the competition in the event of problems and crisis situations, and to anticipate future developments using predictive analytics. A larger, unified data base also promotes innovation, as more information triggers new developments around a company's own products and services and optimization opportunities are recognized in a short time and sales potential is quickly identified. This puts companies one step ahead of their competitors and lets them gain the trust of satisfied and loyal customers by offering innovative services.

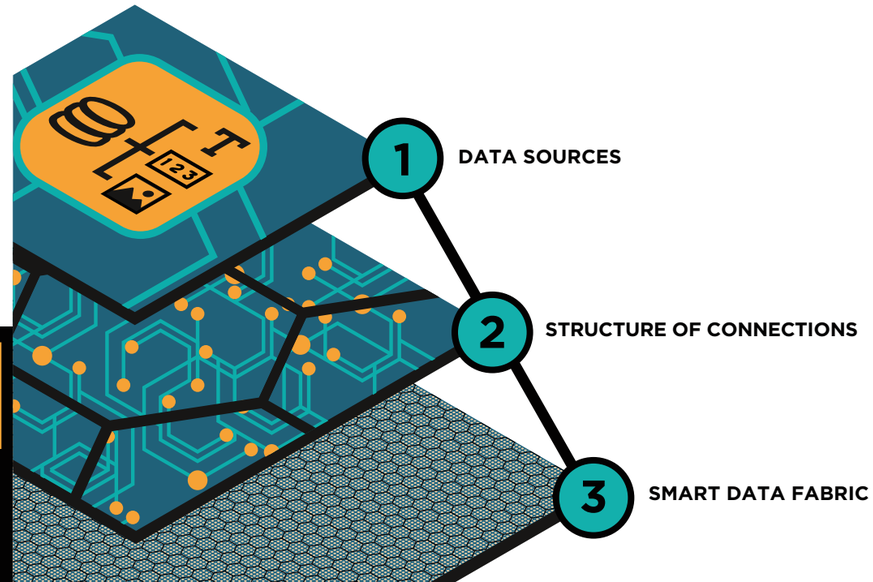

Proven performance

From manufacturing, to financial services, to healthcare – modern applications not only have to process an ever-increasing amount of different data from various sources in real time but ideally analyze it at the same time, too. Only in this way can any production errors be corrected immediately, financial transactions processed in fractions of a second, and life-saving measures initiated without delay. This is why the InterSystems IRIS data platform is designed as a **hybrid transactional/ analytical processing (HTAP)** data platform that guarantees exceptionally high performance and efficiency for multi-workloads in real-time, even with very large data volumes.

InterSystems IRIS achieves this by combining in-memory performance with highly optimized storage utilization and intelligent, data-sensitive distributed caching in which workloads can be distributed and partitioned across multiple nodes in the cluster. Thanks to the decentralized and distributed architecture, there also isn't any need to duplicate data sets completely in the main or fixed storage anymore. In a multi-node setup, InterSystems IRIS can thus be configured so that each node knows exactly what data the others contain. This ensures higher performance, as the overall system is much leaner and access to the individual pieces of information takes place on a selective basis.
The data is stored in so-called **globals**, i.e., compact multi-dimensional arrays. The globals allow applications to simultaneously perform multiple operations at a high level of performance, such as inserting, updating, retrieving, or deleting data. Globals are not limited to the relational model, making InterSystems IRIS a **multi-model data platform**. This gives developers the freedom to create optimal data structures for their applications.

InterSystems IRIS performs significantly better in a direct comparison with other transaction processing and analytical query solutions, such as in-memory and column-store technologies. As one example, InterSystems IRIS queries real-time transactional data thirty times faster than the leading commercial in-memory database.

An equally important aspect is that even if the workload increases significantly, InterSystems IRIS maintains its performance. Plus, the operating costs are lower compared to the competition since fewer IT infrastructure and system resources are required.

Full flexibility

As digitalization continues to progress, the processing and analysis of data is proving to be very dynamic in many companies. Where required, load peaks must be absorbed at short notice and capacities expanded in order to orchestrate all incoming and outgoing data flows. This is no problem with the InterSystems IRIS data platform, because it is freely and flexibly scalable – both vertically and horizontally. As a result, it can manage a large number of simultaneous workloads, data, and users without any loss in performance.

Vertical scalability

InterSystems IRIS is optimized for in-memory processing, but unlike entirely in-memory solutions, it is not limited to the available memory. The data platform can easily make use of additional memory and multiple CPUs at the same time. Flexible caching/tiering enables spillover to fixed storage.
In addition, InterSystems IRIS also provides vertical scalability via parallel SQL processing, with **sharding** providing the basis as a method for partitioning and separating data. In the corresponding cluster, each query is divided into several less extensive queries, which are then executed locally and in parallel on the individual independent file servers – the shards. This not only speeds up the process but also allows a large number of parallel queries in a short space of time. The feedback from the shards is always displayed to users as a combined result.

Horizontal scalability

InterSystems IRIS delivers horizontal scalability by combining sharding with the **unique enterprise cache protocol (ECP)**. In this method, a very large database is divided into several smaller units, which are stored on various independent file servers. A single parent file server acts as the master in the corresponding cluster. The size and number of shards underneath can be changed at any time, allowing free and flexible scaling. It doesn't matter where the shards are physically located – whether on-premise or in the cloud. To ensure the best possible failover, the file servers can be mirrored with InterSystems IRIS. In the event of an instance failure, this mirroring provides automatic failover, ensuring that data remains available for use.

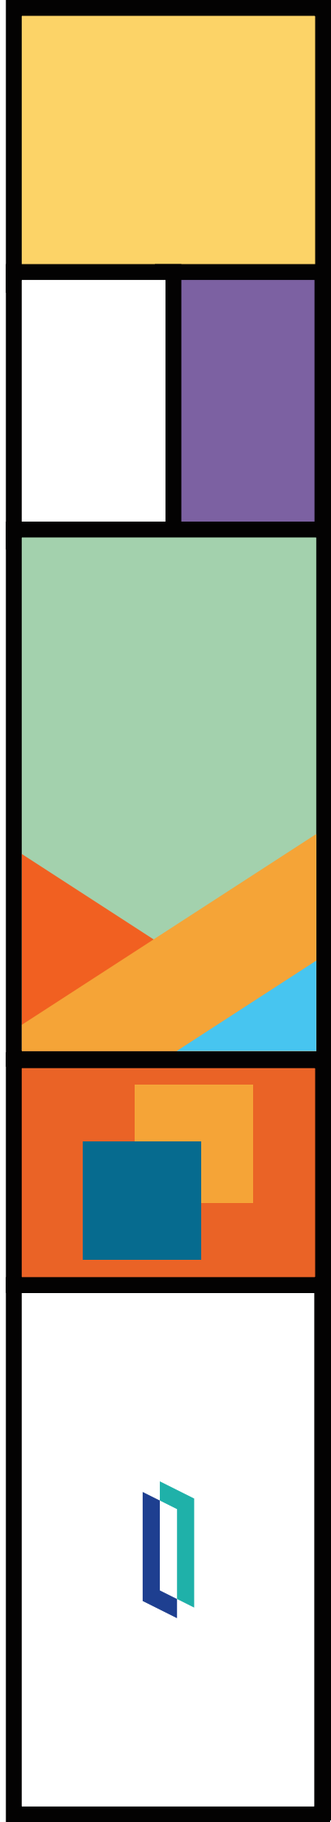

Integrated functions provide insights in real time

InterSystems IRIS has a number of integrated functions for analyzing structured and unstructured data – regardless of its respective format. Patterns, trends, and relationships in the data sets can be identified quickly and easily with the help of predictive modeling, BI, and the corresponding dashboards as well as the complementary use of AI and ML. This enables companies to maintain an overview of the current situation, react quickly to changes, and make accurate predictions. At the same time, a customizable workflow engine is available to issue manual or automatic alerts in case of anomalies and to initiate corrective actions. The same applies to business processes in general: As soon as clearly defined criteria are met, further processes or actions follow, which simplifies and accelerates the workflow significantly.

InterSystems IRIS is an open data platform, which means it allows seamless integration of best-of-breed solutions. Developers and companies can essentially use any tool to provide for better decision making, e.g. the Apache Spark data analytics framework.

BI for the highest demands

Adaptive Analytics is an optional extension for the InterSystems IRIS data platform that makes it even more powerful, enabling accurate and rapid analyses of data at scale, based on BI, AI, and ML.

For this purpose, a **virtual data model layer** is inserted between InterSystems IRIS and common BI, AI, and ML tools. An intuitive user interface provides the experts in charge with access to this layer and helps them to develop a semantic data model in the form of a virtual cube. This cube can be used to organize data, define uniform identifiers across existing application and data silos, and uniquely identify data fields. A **centralized data model** helps companies to solve the problem posed by different definitions or divergent calculations and to provide their employees with a uniform, consistent overview of relevant information and key figures, which enables better business decisions across all departments in a short time.

Key features of Adaptive Analytics

For best-in-class performance, all data always remains in InterSystems IRIS. Adaptive Analytics accesses the data in real time, which prevents content gaps and retrieval of outdated information. Changes to the data model are immediately available in the semantic layer without causing downtime on the user side.

Based on the incoming queries, Adaptive Analytics learns on a continuous basis. For example, the solution captures frequently asked queries and automatically creates aggregates, which in turn greatly increase the speed and efficiency of the system by minimizing access times. In contrast to conventional caching, the queries are also supplemented with additional data fields that may be relevant for queries in the future, which further reduces access times.

Adaptive Analytics supports **self-service BI** so employees in companies can perform interactive and multi-dimensional analyses and create or modify dashboards independently and without the involvement of the IT department. The users are free to decide which tools they use for this purpose. Adaptive Analytics allows the integration of different third-party BI tools – such as Microsoft Excel, Power BI, Tableau, or Qlik – so the unified online analytical processing model (OLAP) can be accessed via them. This is always done independently of the respective query format of the tools so that specific queries are always answered identically.

Easy implementation of AI and ML

AI and ML offer significant potential by optimizing companies' decision-making processes. Using AI and ML technology accelerates these processes, gives them the capacity to work through larger volumes of data, and thus boosts their precision. Above all, AI and ML make it easy to process and analyze the ever-growing volumes of data. At the same time, companies benefit from the automation that comes with them because algorithms – based on live information – can make decisions in real time and thus perform individual business processes completely autonomously, making them more efficient, less prone to error, and more transparent. The use of AI and ML also offers the opportunity to optimize overall productivity and efficiency, provide high-quality and innovative products or services, reduce operating and maintenance costs, and improve the customer experience. In this way, companies can gain a sustainable competitive advantage, open up new sources of revenue, and thus optimize their value creation.

The data base has to be right

A sufficiently large volume of relevant data of high quality is critical to the success of AI and ML projects. InterSystems IRIS makes it possible to merge, cleanse, and harmonize all of an organization's current and historical data – regardless of the source or format. This clean data, also called **healthy data**, is the only way to obtain high-quality results. When it comes to implementing the projects, there's also InterSystems IRIS hybrid transactional/analytical processing (HTAP). This guarantees exceptionally high performance and efficiency for multi-workloads, paving the way for **KI- and ML-optimized processes in real time**. In addition, the data platform is perfectly suited for easily and conveniently adapting the format and structure of data according to the respective requirements prior to analysis.

Combining all the necessary tools

The InterSystems IRIS data platform has high-performance interfaces to powerful third-party tools for developing a company's own AI and ML models. Among others, these include DataRobot, H2O, Apache Spark, and KNIME. Thanks to the **predictive model markup language (PMML)**, the models of external tools can be seamlessly integrated into applications based on InterSystems IRIS. Many popular tools support the export of models via this industry standard. InterSystems IRIS can execute these models natively so they can be implemented in real time as part of individual business processes. The data always remains in InterSystems IRIS; no extraction is required. As a result, the tools benefit from both the first-class performance of the data platform as well as its free and flexible scalability. Taken together, they result in connections with higher throughput, faster queries, more efficient use of system resources, and an excellent user experience.

Embedded Python simplifies and accelerates the development of AI and ML applications

Geared especially toward the quick and easy development of AI- and ML-based applications and services, InterSystems IRIS offers native support for the popular programming language Python since version 2021.2 under the name **Embedded Python**. Embedded Python is an extension to the Python programming language that allows Python code to be executed within the InterSystems IRIS process context, significantly boosting application performance with large volumes of data in particular. An important factor is that Embedded Python can interact natively with objects written in ObjectScript because it uses the same process context as ObjectScript. The reverse is also true, and ObjectScript code can directly access Python modules and libraries. This enables a new level of interoperability between the two languages, allowing objects, methods, and even complete code libraries to be seamlessly exchanged between Python and ObjectScript contexts. Specialists such as data scientists and data architects are thus able to use the numerous freely available Python libraries directly in InterSystems IRIS – which streamlines and simplifies the development process to a significant extent.

IntegratedML enhances in-house projects

IntegratedML also gives companies a tool that helps employees with basic SQL knowledge to implement ML scenarios independently. An in-depth knowledge of ML technologies is not required for this purpose. This is because the tool offers a simple user interface that makes it easier to create and execute ML models as well as monitor the result sets. The individual ML processes are initiated using **SQL-like function calls**. A data scientist then takes care of optimizing the individual models if necessary. The practical significance of a tool like IntegratedML is apparent in light of the acute shortage of skilled workers on the labor market in particular. Besides this, data scientists have more time for more complex tasks when they no longer have to worry about data access and model provisioning.

Natural language processing support

InterSystems IRIS also supports **natural language processing (NLP)**, which refers to functions that go beyond the conventional reading of forms. NLP can be used to analyze the content and sentiment of texts, which in turn allows the extraction of specific information and knowledge – for example, from documents, note fields, or social media feeds. The functions are different from those of other solutions because they take a unique bottom-up approach that automatically recognizes concepts and relationships in text. This is based on a thorough understanding of the language rather than a top-down view in a particular subject area. As a result, there is no need for lengthy, restrictive predefinition of dictionaries or ontologies or for specific domain knowledge. Besides this, InterSystems IRIS also supports the industry standard, Unified Information Management Architecture (UIMA). This allows NLP components to be included in a single pipeline, and the output can be conveniently managed and used in applications that rely on InterSystems IRIS.

An optimal combination for agile application development and flexible deployment

There are always issues that can come up when developing, deploying, and updating applications. This is why developers and companies often work on troubleshooting. But is there a way to get around this and use resources more efficiently? An increasingly popular option is combining containers, microservices, and API management, which is natively supported by the InterSystems IRIS data platform. Thanks to this combination, unwanted surprises are a thing of the past.

Containers simplify all processes

“Runs here, runs everywhere” – working according to this formula, applications in containers are always seamlessly rendered and used. The concept makes it possible to pack an application and all the elements necessary for its operation into a single box, a container. The application functions completely autonomously within this environment. This means that if a container runs smoothly on a developer’s system, easy deployment to the customer or individual employee is guaranteed. Container will behave identically there as well, thereby preventing any incompatibilities with the target system. As a result, new versions and updates of applications can also be rolled out smoothly, allowing easy implementation of approaches around **continuous delivery** and **continuous integration**. In addition, this deployment method provides the freedom to scale up and down dynamically, because containers can be duplicated quickly and easily, for example to cover peaks in demand.

Microservices offer functional autonom

Individual functions of an application can also run as microservices in their own containers. This makes it possible to develop and release the features and their updates separately. Microservices are characterized by functional independence and are therefore protected against the failure of other services and components of the IT infrastructure, resulting in a significantly more stable overall system.

API management regulates the exchange of information

Another aspect is API management, which automatically monitors and orchestrates the necessary exchange between microservices. The API Manager integrated in InterSystems IRIS is available for this exact task, providing IT managers with a transparent overview of what is rolled out via containers and how the containers behave during use. This includes an insight into data traffic, which can also be analyzed, managed, throttled, and restricted. The API Manager automatically reacts to errors or temporary resource bottlenecks and initiates the necessary measures. It is also helpful for version changes, as the API Manager can initially make the new version of an application available to only a few users. If it then becomes clear that the code doesn’t have any bugs, the API Manager gradually increases the update rate. It also ensures security and can take care of user authentication, for example via TLS and SSL on the transport layer, OAuth2 on the protocol layer, and by means of role-based access control.

A strong combination for more flexibility and speed

Overall, the combination of containers, microservices, and API management provides more flexibility and speed in developing, deploying, and updating applications and features. IT managers can work far more efficiently, quickly, and with greater focus. At the same time, they are more effective as a team and have a transparent overview of their system at all times. In this context, continuous delivery and continuous integration enhance the product quality of applications because they are regularly modified and tested, which reduces the risk of programming errors in the code.

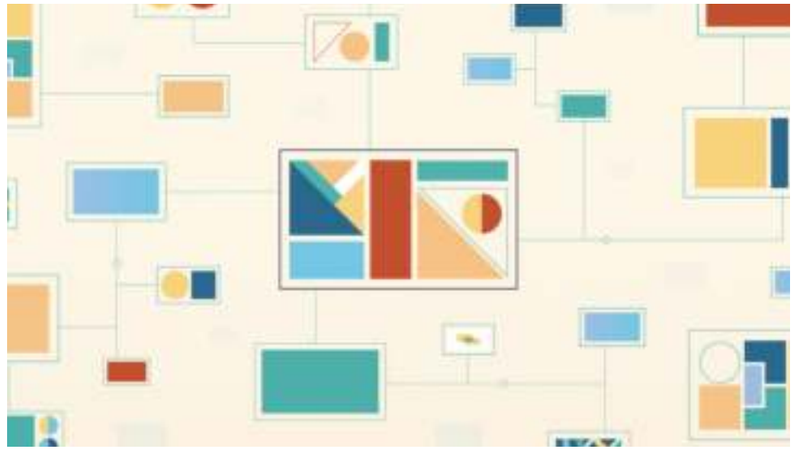

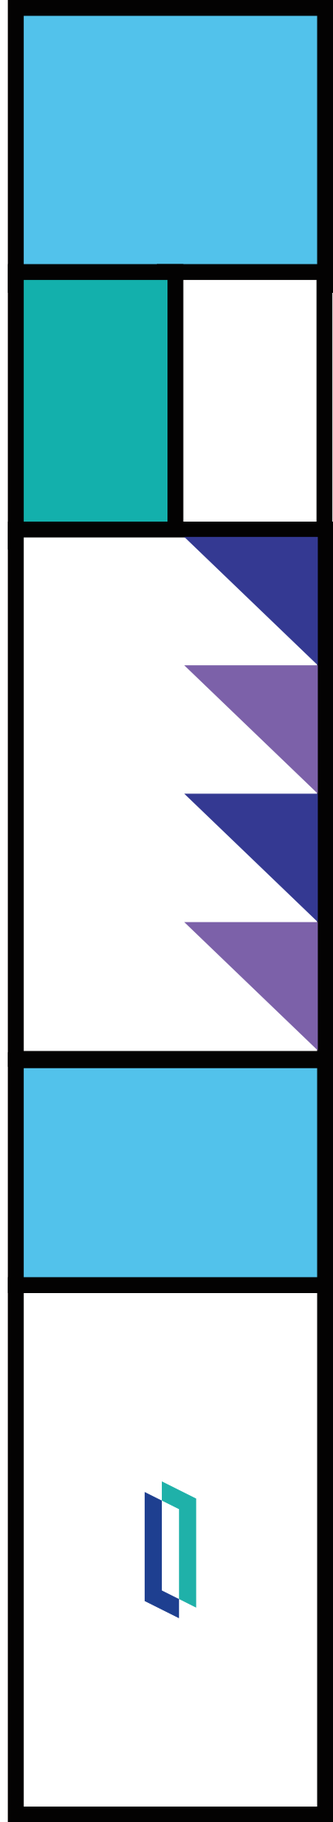

InterSystems: a coach and partner

For over 40 years, InterSystems has provided its customers and partners with outstanding support that has had a decisive impact on our shared success. Among other things, they receive access to extensive in-person and online training to help them exploit the full potential of the InterSystems IRIS data platform. Customers and partners will find a wealth of relevant information summarized in the InterSystems IRIS Experience. A fully functional data platform sandbox is also available here, allowing for practical tests and experiments.Additional expertise is passed on in regular webinars and especially within the lively developer community.

Worldwide Response Center for technical support

Customers and partners can contact the Worldwide Response Center for technical support, which is staffed around the clock with highly qualified software experts. During local office hours, InterSystems customers and partners can always reach an expert who speaks the respective language in the region. InterSystems doesn't outsource services to call centers, so anyone seeking support always communicates directly with InterSystems employees.

These professionals provide immediate support and take as much time as needed to resolve the problem. In doing so, they attach importance to understanding all processes in the company and not just the immediate problem. Customers and partners define the priority of their problem themselves and decide when the support case is closed.

High customer satisfaction

In 2020, Gartner once again named InterSystems "Customer's Choice" in the Gartner Peer Insights Program for Operational Database Management Systems (ODBMS). Gartner Peer Insights is based on reviews from verified end users and their experiences with purchasing, implementing, and operating the InterSystems IRIS data platform. The categories evaluated were scalability, speed, deployment options, and customer service. InterSystems IRIS achieved the best result in the Operational Database Management Systems category with an average rating of 4.7 out of 5.0 stars.

> ***„I told the team at InterSystems: 'In my opinion, "IRIS" really stands for "It Really Is Simple".'"***
>
> **Dagmar Causley, Managing Director, LIB-IT DMS GmbHTRANSWAGGON GmbH**

Future-proof data management thanks to InterSystems IRIS

The use of InterSystems IRIS ensures high-performance, innovative, and future-proof data management that demonstrates its strengths in all industries and use cases. Developers and business users have fast and easy access to all relevant data. In addition, data management is largely automated, which greatly simplifies and accelerates daily work with the ever-growing volume of data. This gives database administrators and data stewards more time to work on other complex tasks and to develop additional data sets.

In terms of its implementation, the data platform offers a number of options: It enables **deployment in the cloud (public, private, hybrid) as well as locally (on-premise) or in virtual machines**. To this end, InterSystems IRIS is available on all leading cloud infrastructures, including AWS, Google Cloud Platform, and Microsoft Azure. In addition, the data platform seamlessly integrates with existing IT infrastructure regardless of the delivery model, allowing companies to continue gaining value from their existing environment.

Solution architects thus have extensive scope to set up their system according to individual requirements. This includes better business process orchestration that can be performed directly on the data platform, making it easier and faster.

> ***„The entire process of converting to IRIS went extremely smoothly, which is remarkable for a project of this scale and complexity."***
>
> **Alexander Marx, Technical Manager and Director IT, Pagero GmbH**

Security and flexible development

A high level of security and data protection is also ensured, as the unified, powerful, and flexible security architecture makes minimal demands on system resources. The architecture is based on authentication, authorization, auditing, and database encryption according to proven industry standards. Authentication and authorization of users is done using passwords and mechanisms such as OAuth2 or two-factor authentication. As part of the auditing process, a log of predefined events relating to a system or application is automatically created so the exact progression of events can be tracked precisely later on. The database encryption includes both data at rest and data in motion. In addition, InterSystems IRIS supports the use of TLS and SSL for data transmission and provides tools for a public key infrastructure (PKI).

InterSystems IRIS facilitates and accelerates the development of innovative applications by providing a uniform and open development environment. Compared to the usual development process, which often involves coordinating different tools, products, and open source projects, working with InterSystems IRIS requires much less time and effort. Among other things, InterSystems IRIS supports not only all common programming languages today – such as Python, C++, and Java

– but also a wide variety of data models. Developers can thus choose their preferred language, and there is a suitable data model for each task which can be used relationally, in an object-oriented way, or via direct/native access.

> ***With InterSystems IRIS, anyone can implement data management for tomorrow today. That's why we want to bring more partners on board and implement innovative projects together. We are always available for advice and support."***

Jochen Boldt, Director of Sales, InterSystems

Notes

InterSystems®
Creative data technology

www.intersystems.com/se info@intersystems.se