COGNITE

Industrial Organizations Require SaaS Environments to Scale Operational Impact

Industrial Organizations Require SaaS Environments to Scale Operational Impact

About Cognite

Cognite is a globally recognized leader in industrial software with a clear vision: to rapidly empower industrial companies with accessible, trustworthy, and contextualized data and drive the full-scale digital transformation of asset-heavy industries.

With its market-leading Industrial DataOps platform, Cognite Data Fusion®, and a comprehensive suite of Industrial generative AI capabilities, Cognite AI, Cognite makes it easy for decision-makers to access and understand complex industrial data. Cognite Data Fusion® is a user-friendly, secure, and scalable platform that enables industrial data and domain users to collaborate quickly and safely to develop, deploy, and scale industrial generative AI solutions that deliver both profitability and sustainability.

↘ Table of contents

WHY COGNITE DATA FUSION® →   BENEFITS FOR YOUR TEAMS →   CONTACT SALES →

⬊ Deployment environments 101

While the number of companies leveraging the speed and scalability of Software-as-a-Service (SaaS) solutions is growing, many industrial organizations still have an aversion to any solution that will not run on-premises or in their private cloud.

In many cases these organizations believe that only on-premises (on-prem) or private cloud environments are acceptable ways to securely store sensitive data, meet regulatory requirements, maintain control of their data, and avoid vendor lock-in. These concerns are valid, and a SaaS solution must fully address each concern while delivering faster time-to-value and solution scalability.

Before delving into how industrial companies can benefit from embracing SaaS, let's begin by defining the commonly accepted benefits and shortcomings of on-prem, private clouds, SaaS, and public clouds.

On premises (on-prem)

In an on-premises (on-prem) environments organizations own and maintain the physical infrastructure (including storage, data servers, networking equipment, etc.) hardware, and software required to host and operate their systems.

Organizations are responsible for costs to manage all aspects of the IT environment, including security, backups, upgrades, installations and more. On-prem is valued most for critical use cases like process control where network interruptions do not impact performance. On-prem infrastructure limits scalability, flexibility, and innovation. Managing and maintaining on-prem solutions require additional investments in infrastructure and resources.

Private cloud

Private cloud is a cloud environment dedicated to a single organization, providing them with exclusive use and control over the cloud infrastructure and resources. It offers enhanced control, security, and customization options but usually requires high upfront costs and ongoing maintenance. Private cloud is valued for the control, reduced hardware management, and increased scalability over on-prem solutions. However, the infrastructure still requires dedicated IT teams to manage and allocate resources, including software and security upgrades and updates.

Public cloud

Public cloud environment infrastructure and services are shared across multiple users/organizations. The services and resources are accessible to the public over the internet, and customers can access and utilize them on-demand. Concerns around public clouds exist around security, especially related to storing of sensitive data outside of an industrial organization's infrastructure, and time-to-value, where initial projects can take 6+ months to deliver business value. Public clouds are valued for their scalability, cost-efficiency, and a wide range of pre-built services. The IT environment, software upgrades and security updates of underlying services are handled by the cloud provider.

Software-as-a-Service (SaaS)

Software-as-a-Service (SaaS) is a software delivery model in which software is accessed through the internet/public cloud, eliminating the need for local installations. Under the SaaS model, the software is stored on remote servers, continuously managed and updated by the service provider, and easily accessed by customers through web browsers, mobile applications, and application programming interfaces (APIs).

SaaS can be deployed in multiple ways (Fig. 1):

- **Multi-tenant SaaS** – Customers share the same resources with their data kept totally separate from each other.

- **Single-tenant SaaS** – Customers do not share resources and have their own dedicated cluster with their own software instance.

- **Virtual private SaaS** – Customers do not share resources and have their own dedicated cluster with a secure connection (private IP address) that is not accessible/findable via public internet.

- **Private SaaS** – Software solutions are installed inside of a customer's private tenant, giving customers increased responsibility and control over deployment, security, maintenance, and data.

Each SaaS environment has its own tradeoffs, and understanding the differences between these options is imperative when selecting a suitable deployment strategy. Industrial organizations value SaaS solutions for their ability to accelerate deployment, minimize the resources to support the solutions, and scale workflows. Above all else, SaaS solutions must deliver a **delightful, interactive user experience**.
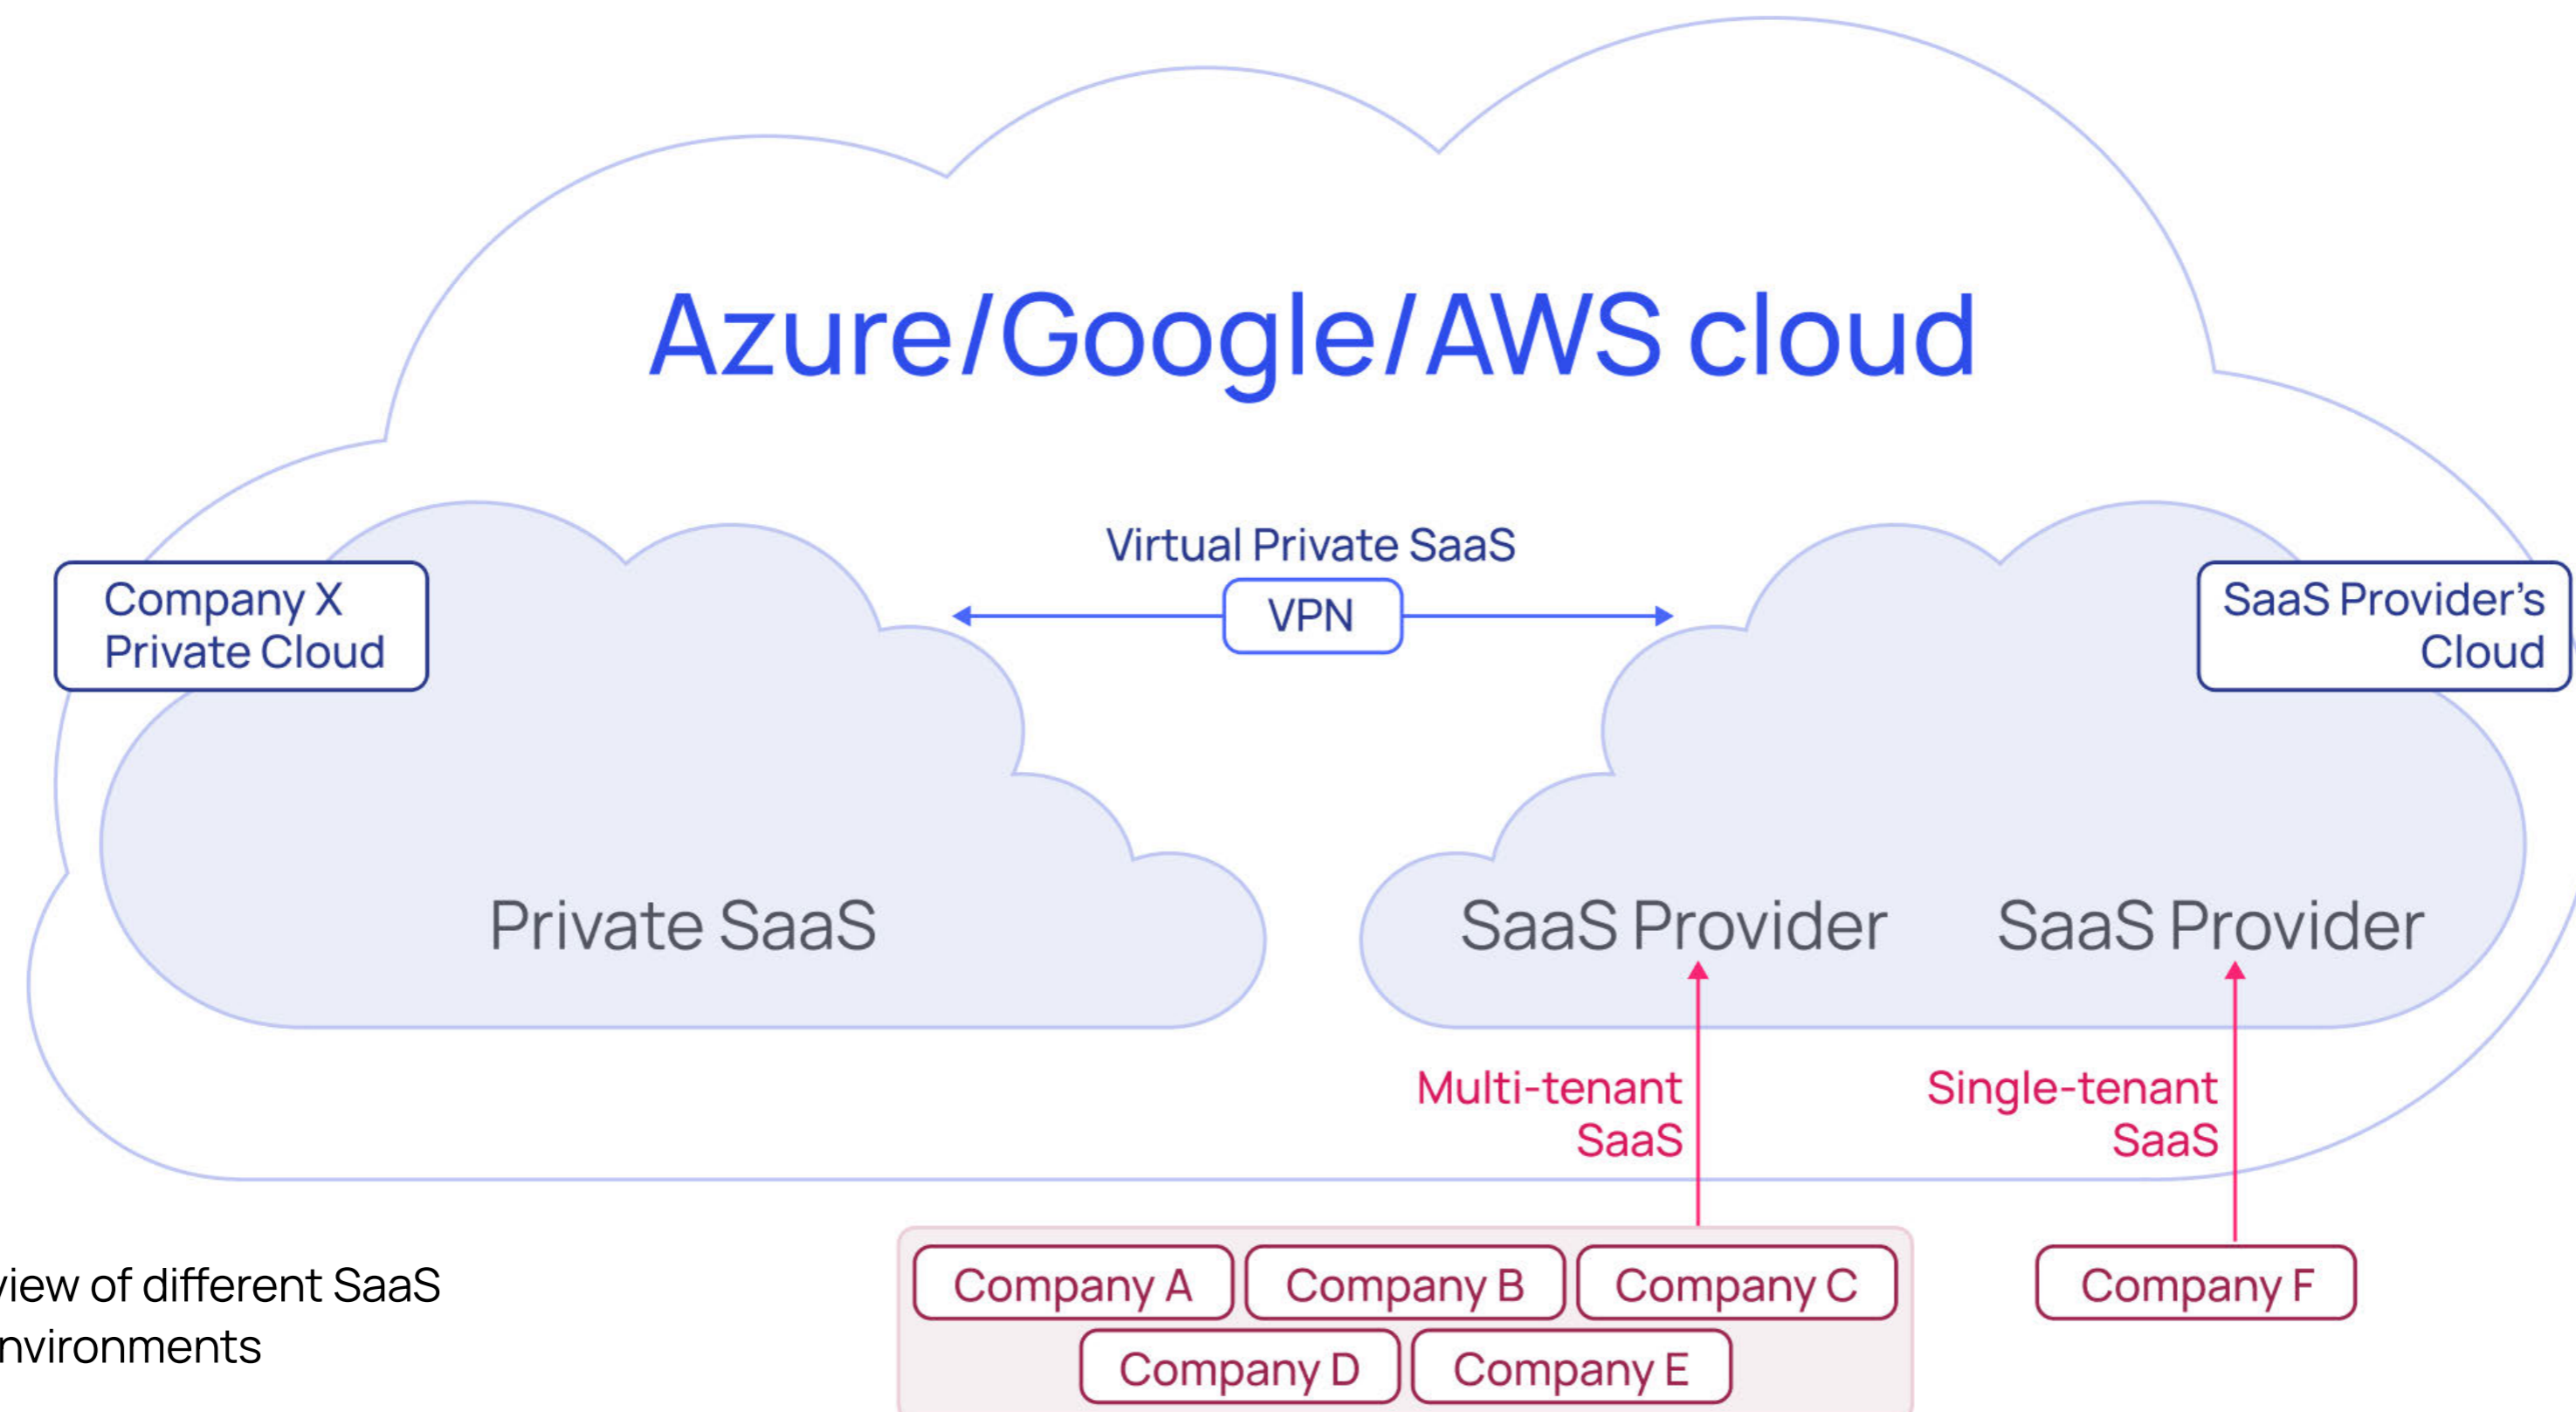
▶ Figure 1: Overview of different SaaS deployment environments

⬂ A delightful, industrial data user experience

Energy, manufacturing, and power and renewables companies have to deal with enormous amounts of data that includes structured data from enterprise applications, operational technology (OT) data in the form of time series, and unstructured data such as P&IDs, 3D diagrams, technical documentation for equipment reports, work orders and more. The wide diversity and high sensitivity are significant barriers when it comes to accessing, managing, and working with industrial data in a simple, unified experience.

This complexity becomes apparent when working with industrial data in a data warehouse or data lake. Data warehouse environments are inflexible in accommodating new data sources and changing business requirements due to predefined structure. Many organizations started adopting data lakes to overcome the inflexibility of data warehouses only to realize these environments lack structure and governance, which results in challenges related to data organization, data quality, as well as high demand for specialized skills (data scientists) to extract meaningful insights. The need to overcome these challenges, in a simple and scalable way, has led to a rapid growth of SaaS solutions in the field of Industrial DataOps →.

The power of Industrial DataOps lies in its ability to efficiently ingest, contextualize →, and optimize data to deploy applications and solutions to solve operational use cases at scale. Implement-

ing DataOps requires managing large amounts of complex industrial data and that, in itself, requires different types of databases, different clustering, and other optimization techniques to ensure top-notch performance at scale. This optimization is enabled by hundreds of microservices, many API nodes, aggregated controllers, etc. This complexity is invisible to users but is vital in enabling an interactive level of user performance, flexibility and scalability of data (e.g., a Cognite Data Fusion® cluster consists of 1000 Kubernetes → pods representing almost 300 different microservices).

At SaaS companies, like Cognite, there are hundreds of specialized services and people that ensure the environment is stable and highly performant for customers through a SaaS experience. Thus, for a typical industrial company to successfully manage a complex SaaS solution in their private tenant, consistent with a SaaS user experience, it would require a dedicated team solely focused on maintaining the software.

Additionally, scaling similar solutions in a private cloud can be extremely difficult compared to SaaS environments, where scaling is typically handled by the service provider. Private clouds require organizations to manage and scale their infrastructure internally, which is resource-intensive (both human and financial) and slower to adapt to growing users' demands.

Choosing the right SaaS environment

When evaluating different SaaS environments, many have the initial reaction to deploy these solutions within their own private tenant. Deploying inside of one's own tenant passes responsibility and control over deployment, security, maintenance, and data to the industrial organization. While this may be tempting, it is essential to consider the following (Fig. 2):

- **Deployment costs** – While counterintuitive, the more control customers want to own the environment, the higher the cost to the SaaS provider. Depending on the complexity of the solution, a private SaaS solution may be 10-20X more/year than using a multi-tenant.

- **Vendor guarantees their security standards** – SaaS providers practice comprehensive security and privacy measures and conduct thousands of security updates daily. Managing such level of security within a Private SaaS will require a dedicated security team. Moreover, vendors are required to undergo testing by third parties to obtain major security compliance certifications (e.g. ISO 27001, SOC 2 Type 2, etc.). The majority of industrial companies have not received or maintained these certifications.

► Figure 2: SaaS environments comparison table

	SaaS Environments Supported by Cognite			Private SaaS
	Multi-tenant SaaS	Single tenant SaaS	Virtual Private SaaS	
Deployment cost	$	$$	$$$	$$$$
Provider guarantees their security standards	Yes	Yes	Yes	No
Inaccessible from public internet	No	No	Yes	Yes
Software updates maintained by service provider	Yes	Yes	Yes	Yes*
Provider guarantees SLA and user experience	Yes	Yes	Yes	No
Customer retains data ownership	Yes	Yes	Yes	Yes
Data is never stored outside of the customer tenant	No	No	No	Yes

*Requires specialized access to the Private SaaS environment by the vendor

■ **Inaccessible from the public internet** – This can be a requirement for some critical industries or sensitive data. Both Virtual Private and Private SaaS make data inaccessible from the public internet. The Virtual Private SaaS uses a virtual private network (VPN), granting it a local IP address, and the SaaS solution runs in the same datacenter as the customer's tenant, making it indistinguishable from their own network.

■ **Software updates maintained by service provider** – All software maintenance (updates, upgrades, bug fixes, etc.) is handled by SaaS providers. Yet, if an organization chooses Private SaaS deployment, it either needs to manage isolated software installations, upgrades, and necessary infrastructure expansions on its own or rely on the SaaS provider or partner, likely at an additional cost.

■ **Vendor guarantees SLA and user experience** – For providers to guarantee a service level agreement (SLA), they must have access to control the underlying services and infrastructure. Private SaaS deployments can be a blocker for providers to ensure their SLA is met. Additionally, because providers cannot control the resource allocation, they cannot guarantee an interactive user experience as other workloads running in the private cloud may impact performance.

■ **Customer retains data ownership** – This must be true regardless of the chosen SaaS environment. An organization's data is always their own to share, copy, or delete as they see fit.

■ **Data is never stored outside of the customer tenant** – The Private SaaS is the only environment where this is fully guaranteed. Achieving this has significant tradeoffs regarding higher costs, effort to maintain security, SLA, and a delightful user experience.

Cognite Data Fusion® – a SaaS environment trusted by industry

To ensure an interactive user experience for industrial users, Cognite performs continuous optimization of our SaaS offering →. Cognite Data Fusion® takes diverse data types and contextualizes them to build the semantic meaningful relationships that make up the Industrial Knowledge Graph →. It enables users to utilize data at scale and speed,

ensuring a high performance interactive user environment.

Cognite Data Fusion® caches large amounts of data (in many different formats, from many different sources), utilizing many services, and puts it into specialized storage and staging areas that allow users to execute search with a sub-second response time across all data types. This is possible because of the cloud-native deployment architecture, enabling frequent updates for continuous improvement of the user experience (Fig. 3).

In addition, to provide the highest level of availability, security and performance of SaaS environment to customers, Cognite implements and meets the following measures:

- Maintain staging environments that are 100% in sync with the production environment for proper testing and validation of updates. Environments are updated thousands of times per week to provide necessary improvements and optimization.

- Manage a single version of generally available software. All customers benefit from new releases, without having to migrate versions. Managing a single version reduces risk and increases the rate of innovation and new features, such as generative AI-powered search.

- Tenants are clean from other processes, data and jobs with enforced governance to ensure that at any time users access information through Cognite Data Fusion®, they use the proper API channels, enforcing user access rights, security, scalability, etc.

- 24/7 monitoring and support using specialized deployment technology and operational management tools.

- Security tested and audited by third parties with compliance to ISO 27001, ISO 9001, SOC 2 Type 2, CCC+, World's first DNV certified Digital Twin.
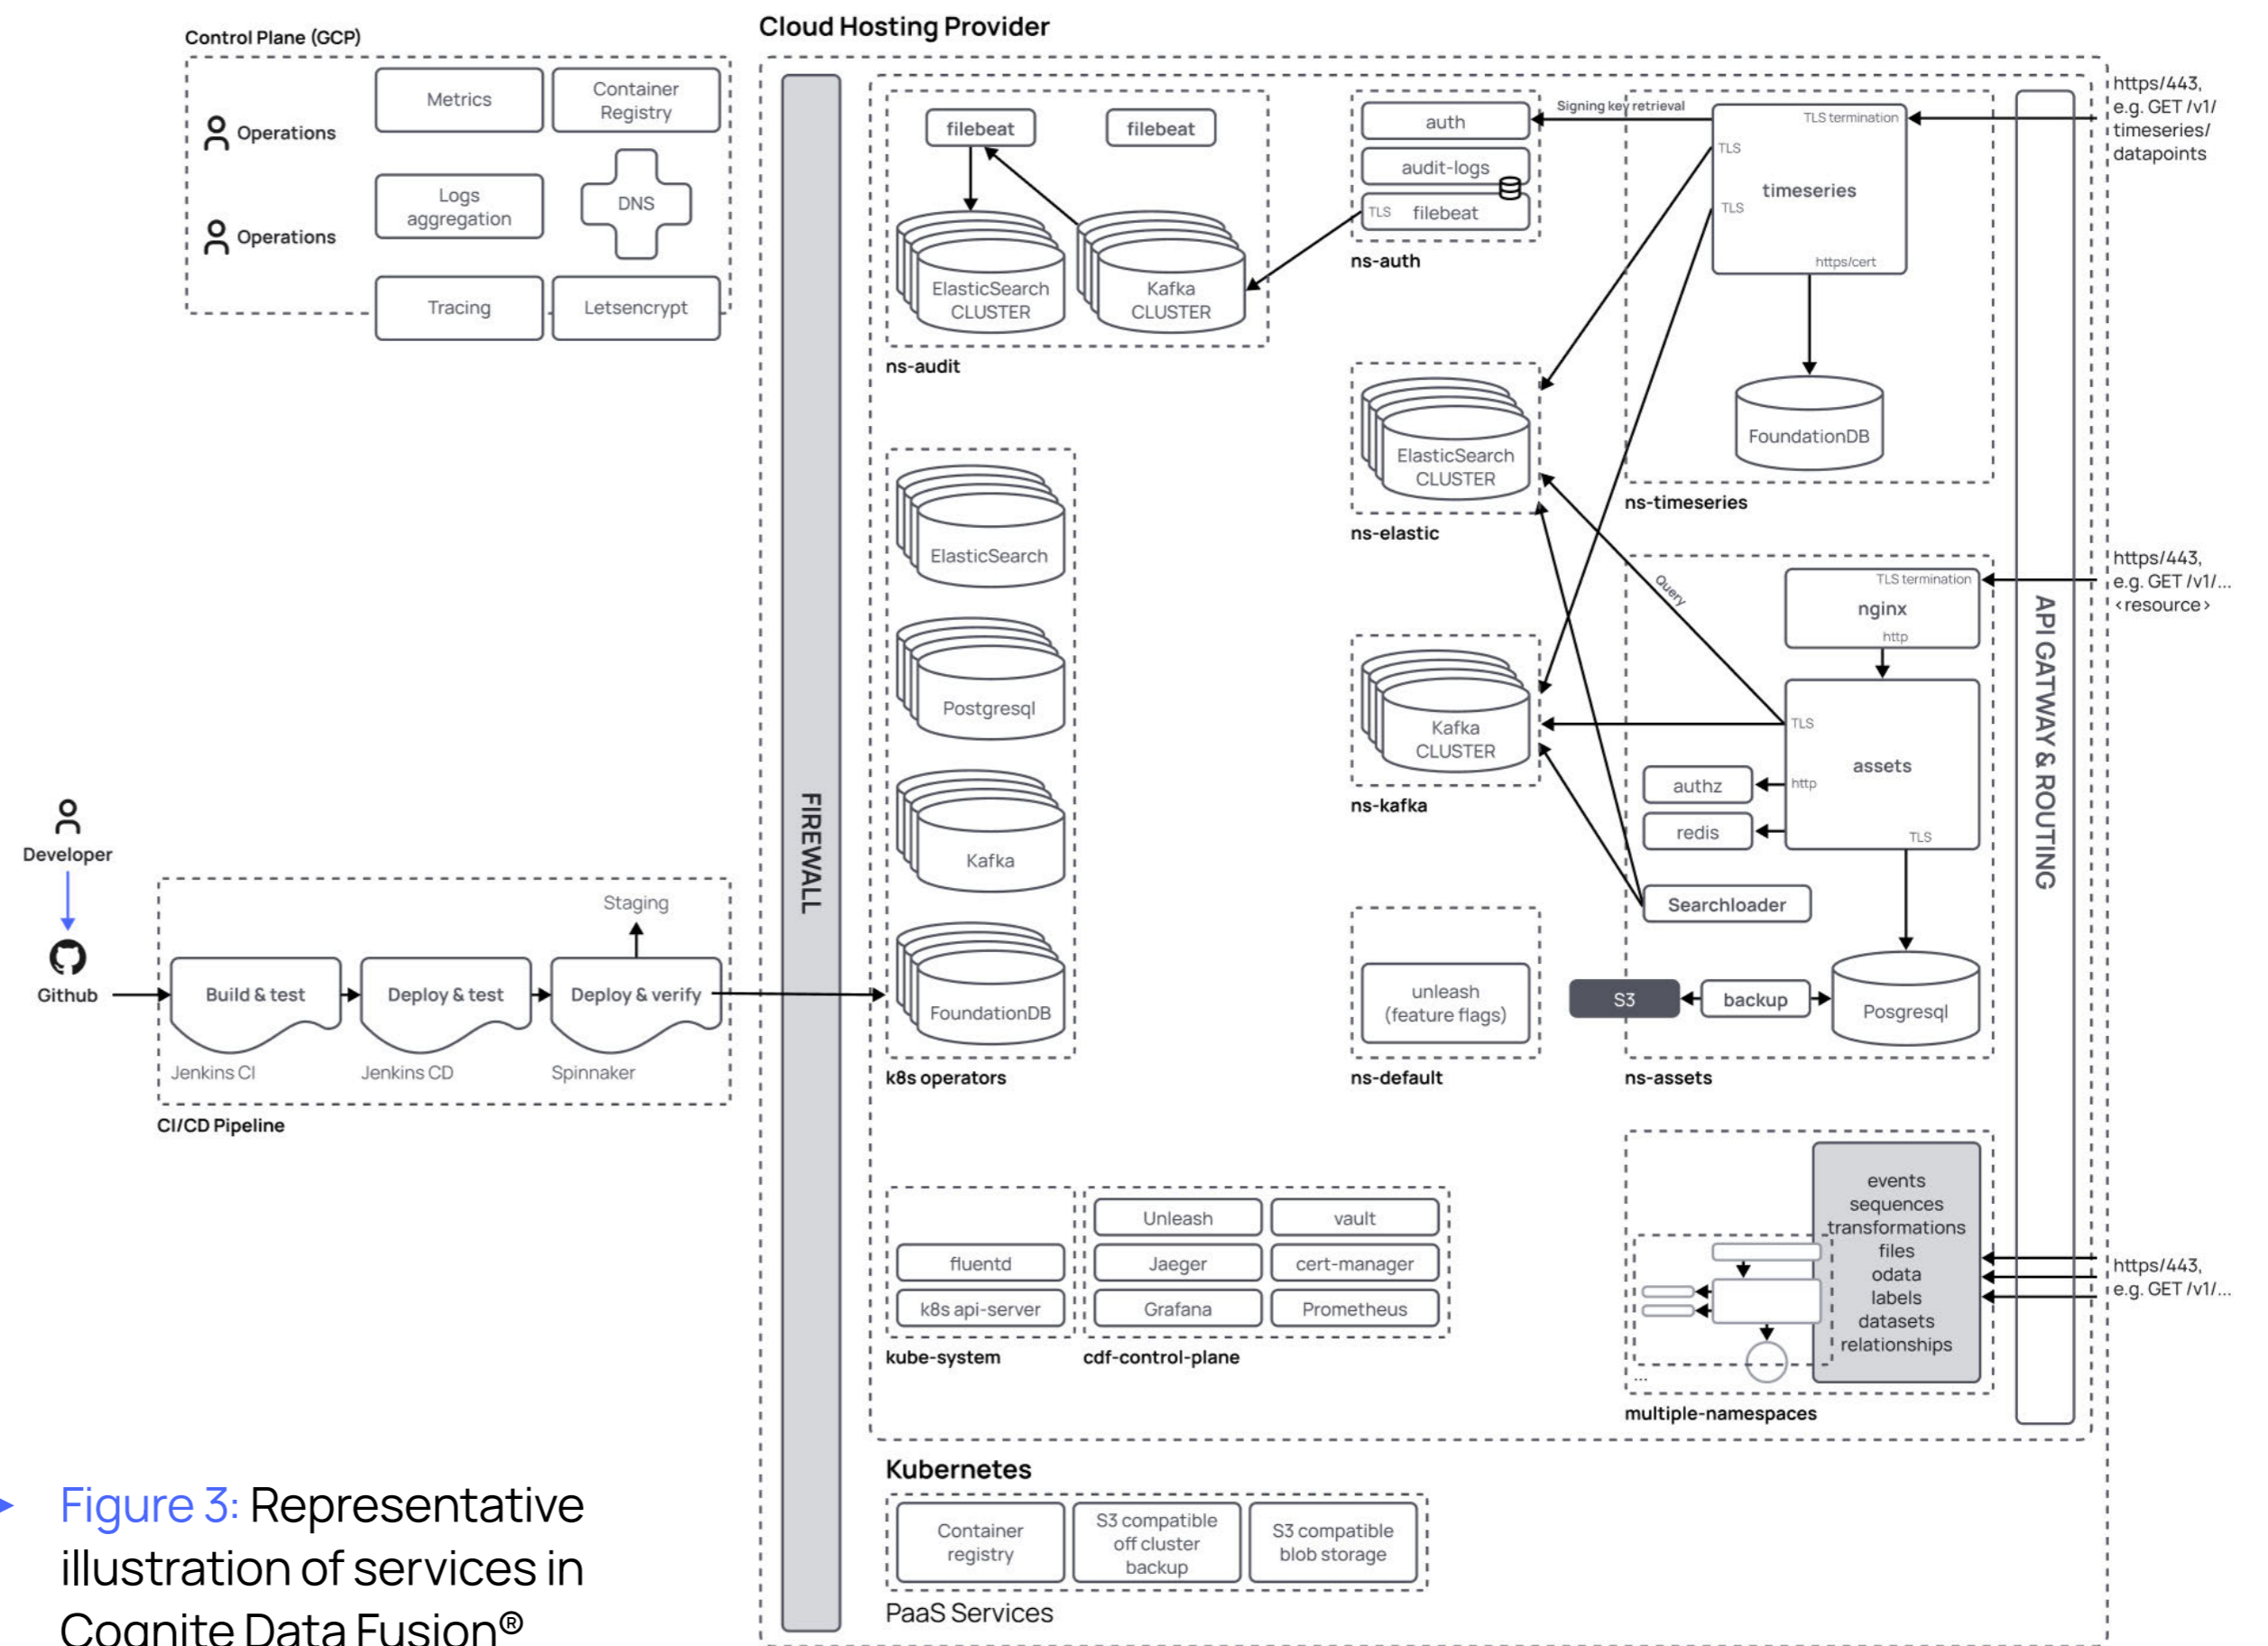
▶ Figure 3: Representative illustration of services in Cognite Data Fusion®

These certifications are guaranteed when operating in a Cognite multi-tenant, single tenant, or virtual private tenant. Cognite cannot guarantee security standards if a deployment were to run in a customer's private tenant/cloud.

- Support for specific regulatory and industry requirements including: NIST CSF, IEC 62443.2-4, IEC 62443.3-2, IEC 62443.3-3, IEC 62443.4-1, CMMC, FIPs, NERC CIP v.5, GxP.

Cognite has been providing a secure and stable SaaS environment to industrial customers since its inception. We have done it repeatedly, with global industrial companies, including B.Braun, Hess, Celanese, Statnett, Total Energies and others →. Cognite Data Fusion® is offered as a multi-tenant, single tenant, or virtual private tenant, all depending on the needs and requirements of the end users.

For the most critical industries and most sensitive data, customers use Cognite Data Fusion® via a virtual private tenant.

> **Cognite's Virtual Private SaaS environment combines the security and privacy of a private cloud with scalability and flexibility of single and multi-tenant environments.**

In this environment, Cognite Data Fusion® runs in the same data center as a customer's private tenant, eliminating any regional data residency requirement concerns. Configured within the virtual private local network, it is granted a local IP address, exclusively to your company and unreachable from the public Internet. All data ownership is retained by the customer and only accessible through documented APIs and connectors. The customer always retains ownership of the data.

Our customers trust us not only because we are compliant with the multitenant architecture and follow the best-practices from hyperscalers like Microsoft, SalesForce, Google, Adobe, and others, but also because Cognite has domain expertise in energy, manufacturing and power and renewables. We understand the importance of providing a secure and reliable SaaS environment that is flexible, scalable and quickly adapts to growing business needs. Cognite Data Fusion® → provides simple access to complex industrial data, offering the most interactive user experience for industrial companies to solve problems faster, optimize processes and achieve more operational efficiency.

Want to know more about our product?

Explore more insights from Cognite
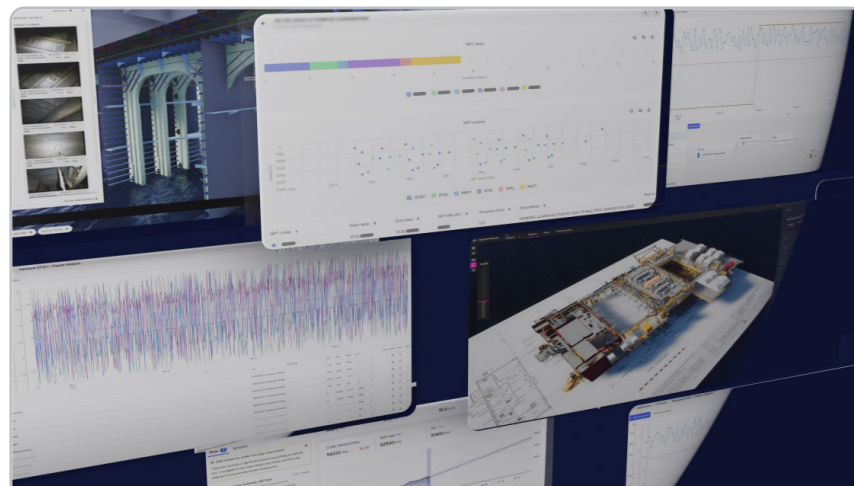
PRODUCT TOUR

Learn from Cognite customers and product managers how Cognite Data Fusion® simplifies and streamlines the data experience of a subject matter expert.

WATCH NOW →

CUSTOMER STORIES

Discover how Cognite Data Fusion® makes data more accessible and meaningful, driving insights that unlock opportunities in real-time, reduce costs, and improve the integrity and sustainability of your operations.

GO TO STORIES →

ANALYST REPORT

Customer interviews and financial analysis reveal an ROI of 400% and total benefits of $21.56M over three years for the Cognite Data Fusion® platform.
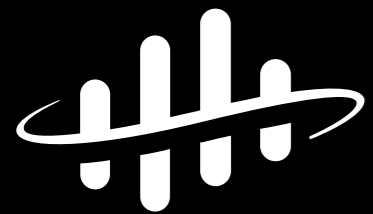
READ THE REPORT →

BLOG

Discover our rich catalog of industry insights and technology deep dives.

READ OUR NEWEST BLOGS →

WHY COGNITE DATA FUSION® →   BENEFITS FOR YOUR TEAMS →   CONTACT SALES →

COGNITE

COGNITE.AI →