Tableau Pulse Security Overview

tableau from Salesforce

Contents

Introduction

At Salesforce, Trust is our #1 core value, which is why our data-centric strategy supports the company's commitment to run the most secure, trusted, reliable, and available cloud computing service. Customer success is the driving force behind our data center strategy, so delivering the highest standard in availability, performance, and security is our top priority.

Tableau Pulse leverages the capabilities of generative AI to simplify data analytics for its end users by using large language models (LLMs) from a rich, diverse set of partners to create, execute, and deliver concise, analytical, and actionable data insights for customer end users to act upon.

This document describes the security, privacy, and architectural components that make up Tableau AI - a fundamental piece of Tableau Pulse.

Shared Responsibility Model: Partnership in Data Protection

The shared responsibility model is an important concept for Salesforce, as well as for our customers and partners. As a data processor we must implement robust security and privacy measures to protect the confidentiality, integrity, and availability of Customer Data. It's equally important for our customers to review and understand the capabilities available to them, as well as how to secure their deployment of Salesforce to meet and uphold their security, contractual, and regulatory obligations. Customers own their data, and they control what data is submitted to our services. Salesforce has an established information security policy for data classification. All information that customers electronically submit to our services is categorized as Mission Critical, which provides the highest level of security.

For additional information, check out the Security Perspective on the Shared Responsibility Model.

Next, we'll discuss Tableau Pulse and its generative AI security capabilities, as well as the underlying controls that comprise the overall service.

[1] Tableau Pulse encompasses a variety of features–including generative AI–which can be enabled within the Setup page. In this overview, the generative AI features in Tableau Pulse are collectively referred to as Tableau AI.

[2] Customer Data refers to electronic data and information submitted by or for a Customer to the Services, excluding Content and Non-SFDC Applications

Tableau Cloud Architecture

Tableau Pulse is a newly launched product offering, powered by Tableau AI–a suite of predictive and generative AI capabilities–that resides within Tableau Cloud. As a native service, it benefits from the existing security, availability, and confidentiality controls present in the Tableau Cloud service.

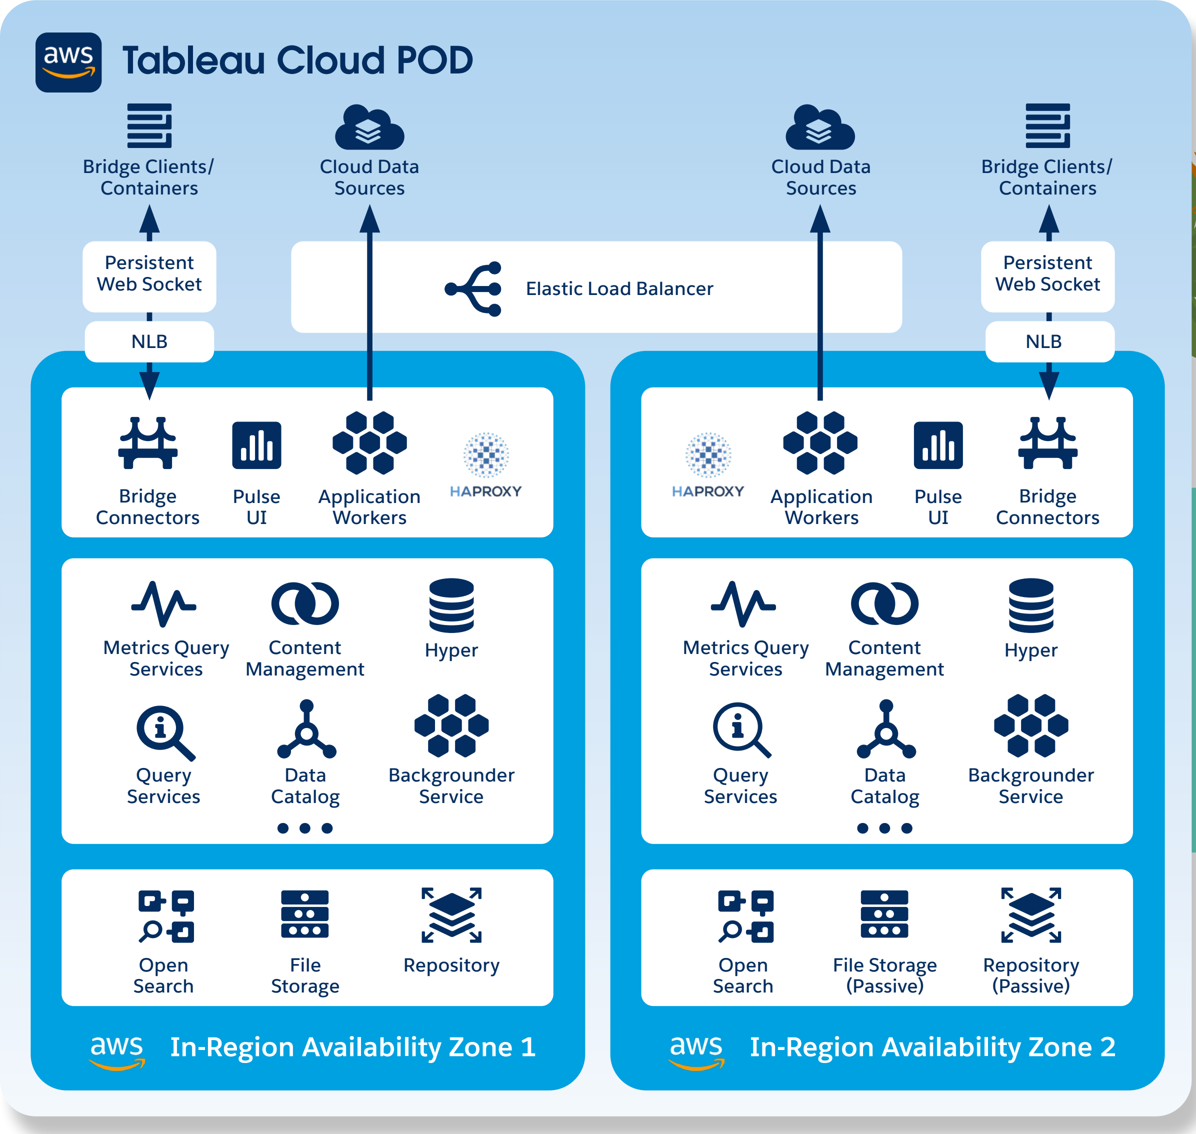

Tableau Pulse is a reimagined data experience. Built on the Tableau platform, it empowers employees with intelligent, personalized, and contextual insights delivered directly in the workflow. Tableau Pulse helps organizations integrate data into daily jobs to help employees make efficient decisions. Tableau Pulse takes users beyond the how and the what to show the why behind their data without users having to learn a new tool or build comprehensive visualizations.

To gain a better understanding of how Tableau Pulse works, let's dive into the overall data flow components. There are multiple layers of integrated capabilities within Tableau Pulse, including a Metrics Layer, an Insights Platform, and Next-Gen Experiences. All of these components combine to offer business end users useful Insights Summaries to aid in decision-making.

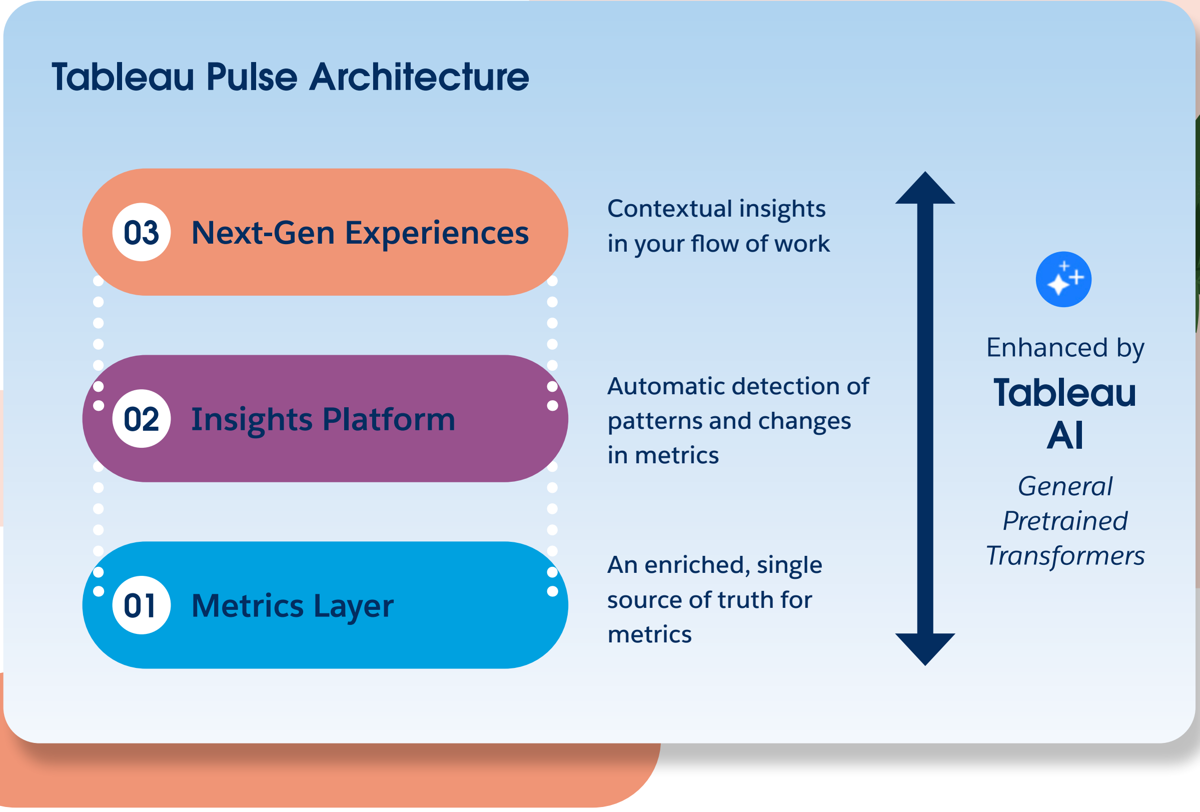

Tableau Pulse Architecture

03 Next-Gen Experiences — Contextual insights in your flow of work

02 Insights Platform — Automatic detection of patterns and changes in metrics

01 Metrics Layer — An enriched, single source of truth for metrics

Enhanced by **Tableau AI**

General Pretrained Transformers

Tableau Pulse Metrics Layer

The Metrics Layer is where analytics professionals create and manage metrics. It's where metric definitions are enriched with real-world business context. In other words, this enriched data model is the "secret sauce" that powers our insights and analysis. Metrics are defined at the data source level, not within a workbook, which means that each metric is a single source of truth that ensures consistency at the consumption layer.

Customers may configure and restrict their Tableau Pulse experience in a limited fashion to use only certain data sources, projects, or workbooks. Some relevant metric examples may include but aren't limited to units of measure, relevant driver dimensions, related metrics, and preferred comparisons.

The Tableau Pulse Metrics Layer natively provides a bounded, secure enclave for insight detection. It provides critical information to help determine which analytics to run and how to present the information in a useful manner.

Permissions for Creating Metric Definitions

- Any user with a site role of Creator, Site Administrator Explorer, or Explorer Can Publish, has the ability to create metric definitions in Tableau Pulse.
- To create a metric definition from a published data source, the user must have the Connect and View permission capabilities for the data source.

Key Permissions Details for Metrics[3]

- Users can access Tableau Pulse from the Tableau Cloud navigation menu.
- The ability to create or view metrics is based on the permissions to connect to and access data within a data source. The data that users see when viewing a metric uses the row-level security (RLS)[4] that's directly applied to the data source.

Permission for Viewing Metrics

- Tableau Pulse doesn't prompt users to sign in to view data; however, one of the following must be true for users to view metric data:
  - The credentials for the data source are embedded.[5]
  - The user's credentials are passed to the data source with single sign-on (SSO).
  - The user's credentials are saved for the data source.[6]
  - The data source doesn't require the user to authenticate to access the data.

Tableau Pulse Insights Platform

The Tableau Pulse Insights Platform provides users with a personalized, AI-generated summary to address the what, why, and so-what metric details to help business users make sense of the data. This platform includes detection and observation about metrics where insights are ranked by usefulness, and it communicates these insights from populated metrics. There are currently several components that comprise the Insights Platform, including:

Insights Generation[7]

- Outlier detection
- Trends
- Sharp increases
- Abnormally high or low data points

Insights Ranking

Tableau Pulse uses insights ranking to issue a rank via an algorithm that focuses on metric impact, relative importance, and personalized user feedback to determine which insights should be presented to end users. Insights are communicated to end users through short, natural language snippets that are accessible in digests via email or on Tableau Pulse Metrics pages. Engagement tracking is also performed to improve future rankings.

[3] permissions-for-metrics
[4] rls-options-overview
[5] publishing-sharing-authentication
[6] manage-stored-credentials
[7] On a periodic basis, Tableau Pulse tests metrics on these insights.

Insights Summary

Tableau Pulse contains a library of important insight types, including:

- An easily-digestible summary of insights and critical points for all metrics.
- A generative AI service that uses prompt generation to review top insights that are detectable when prompts are sent using the service. As soon as Tableau AI processes a prompt to generate an Insights Summary, the prompt and the response are forgotten.
- A generative AI-formulated narrative that's shared in the Tableau Metrics Digest and in the Tableau Pulse UI.

By default, Tableau Pulse generative AI is not enabled for Tableau sites.[8] This also means that features like Pulse Insights Summaries aren't enabled by default. Customers must enable this feature within the Tableau Cloud Setup page.

- If generative AI is turned on, the Digests and Metrics detail pages include generative AI summaries.
- If generative AI is turned off, Salesforce will communicate in a template-based language without generative AI summaries.

Insights Summary Permissions

All insights generated are restricted to and must respect the data security context–including the RLS of the requesting user–to ensure that they access only the data they're authorized to view. RLS in Tableau restricts the rows of data that a user can access within a given workbook or data source at the time they view the data.

Follow-Up Questions

Follow-up questions are presented to users based upon a template of relevant questions that are connected to an insight type. Users can choose from various questions where each subsequent response provides a new list of potential follow-up questions.

[8] pulse-set-up

Next-Gen Experiences

The Tableau Pulse Next-Gen Experiences use insights data to populate relevant summarized insights from captured metrics and permit users to engage with relevant insights. This allows users to act upon everything within the flow of work, such as:

- Email, Slack, mobile, or desktop to deliver meaningful or necessary information.

Insights in Email Digest

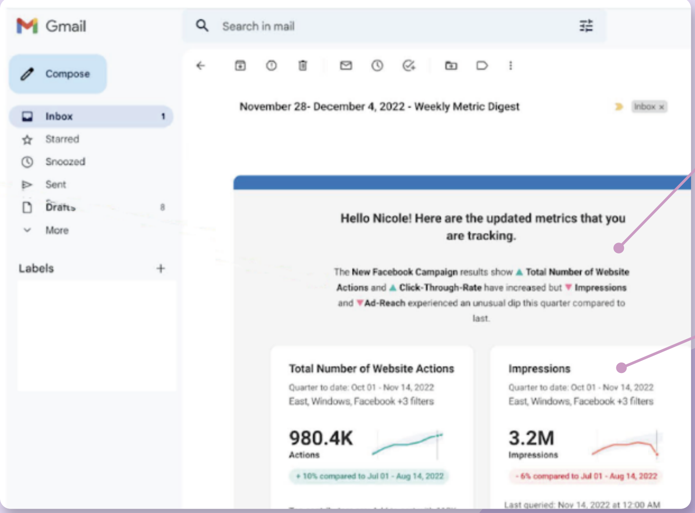

November 28- December 4, 2022 - Weekly Metric Digest

Hello Nicole! Here are the updated metrics that you are tracking.

The **New Facebook Campaign** results show ▲ **Total Number of Website Actions** and ▲ **Click-Through-Rate** have increased but ▼ **Impressions** and ▼**Ad-Reach** experienced an unusual dip this quarter compared to last.

Total Number of Website Actions
Quarter to date: Oct 01 - Nov 14, 2022
East, Windows, Facebook +3 filters
980.4K
Actions
+ 10% compared to Jul 01 - Aug 14, 2022

Impressions
Quarter to date: Oct 01 - Nov 14, 2022
East, Windows, Facebook +3 filters
3.2M
Impressions
- 6% compared to Jul 01 - Aug 14, 2022
Last queried: Nov 14, 2022 at 12:00 AM

Insight Summary containing the most important Insights to pay attention to (GAI)

Top-ranked Insight helps users understand dectected patterns or recent changes

Digests and Alerts insights are sent to users via email digests.

Insights in Slack Digest

Insight Summary containing the most important Insights to pay attention to (GAI)

Top-ranked Insight helps users understand dectected patterns or recent changes

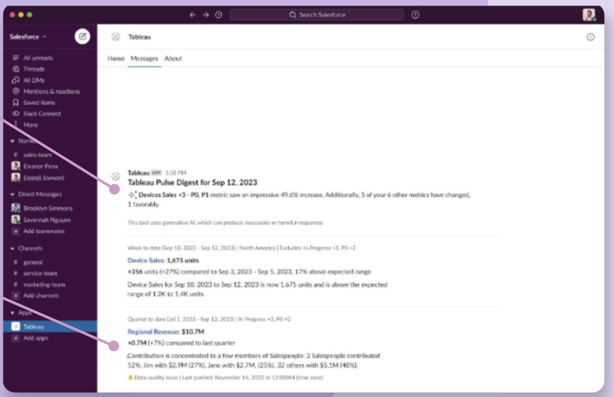

Tableau APP 5:28 PM
Tableau Pulse Digest for Sep 12, 2023
Devices Sales +3 - P0, P1 metric saw an impressive 49.6% increase. Additionally, 5 of your 6 other metrics have changed, 1 favorably.

This tool uses generative AI, which can produce inaccurate or harmful responses

Week to date (Sep 10, 2023 - Sep 12, 2023) | North America | Excludes: In Progress +3, P0 +2
Device Sales: **1,675 units**
+356 units (+27%) compared to Sep 3, 2023 - Sep 5, 2023, 17% above expected range
Device Sales for Sep 10, 2023 to Sep 12, 2023 is now 1,675 units and is above the expected range of 1.2K to 1.4K units.

Quarter to date (Jul 1, 2023 - Sep 12, 2023) | In Progress +3, P0 +2
Regional Revenue: **$10.7M**
+0.7M (+7%) compared to last quarter
Contribution is concentrated to a few members of Salespeople: 2 Salespeople contributed 52%: Jim with $2.9M (27%), Jane with $2.7M, (25%), 32 others with $5.1M (48%).
⚠ Data quality issue | Last queried: November 14, 2022 at 12:00AM (time zone)

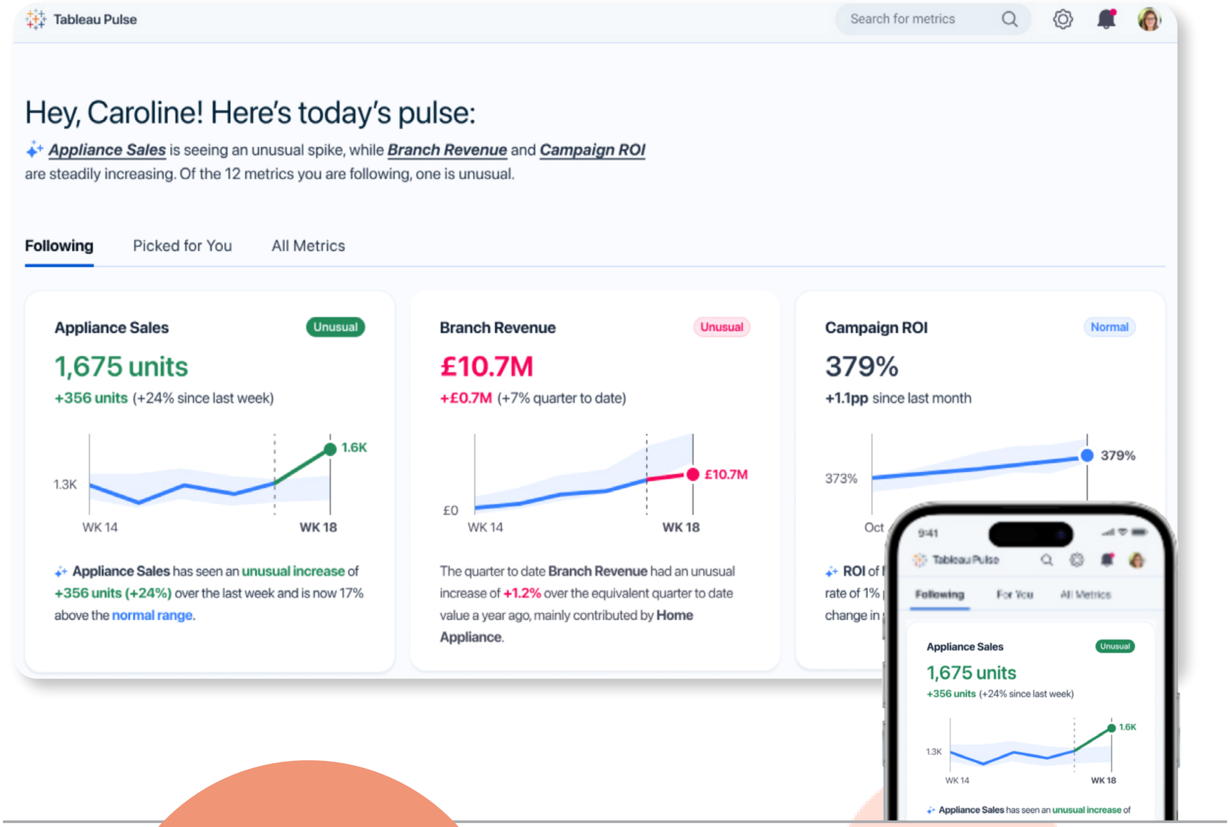

Tableau Pulse

Hey, Caroline! Here's today's pulse:

Appliance Sales is seeing an unusual spike, while *Branch Revenue* and *Campaign ROI* are steadily increasing. Of the 12 metrics you are following, one is unusual.

Appliance Sales Unusual

1,675 units

+356 units (+24% since last week)

Appliance Sales has seen an **unusual increase** of **+356 units (+24%)** over the last week and is now 17% above the **normal range**.

Branch Revenue Unusual

£10.7M

+£0.7M (+7% quarter to date)

The quarter to date **Branch Revenue** had an unusual increase of **+1.2%** over the equivalent quarter to date value a year ago, mainly contributed by **Home Appliance.**

Campaign ROI Normal

379%

+1.1pp since last month

ROI of I

rate of 1%

change in

Appliance Sales Unusual

1,675 units

+356 units (+24% since last week)

Appliance Sales has seen an **unusual increase** of

+

New Consumption User Experience allows users to interact via web UI to explore follow-up questions

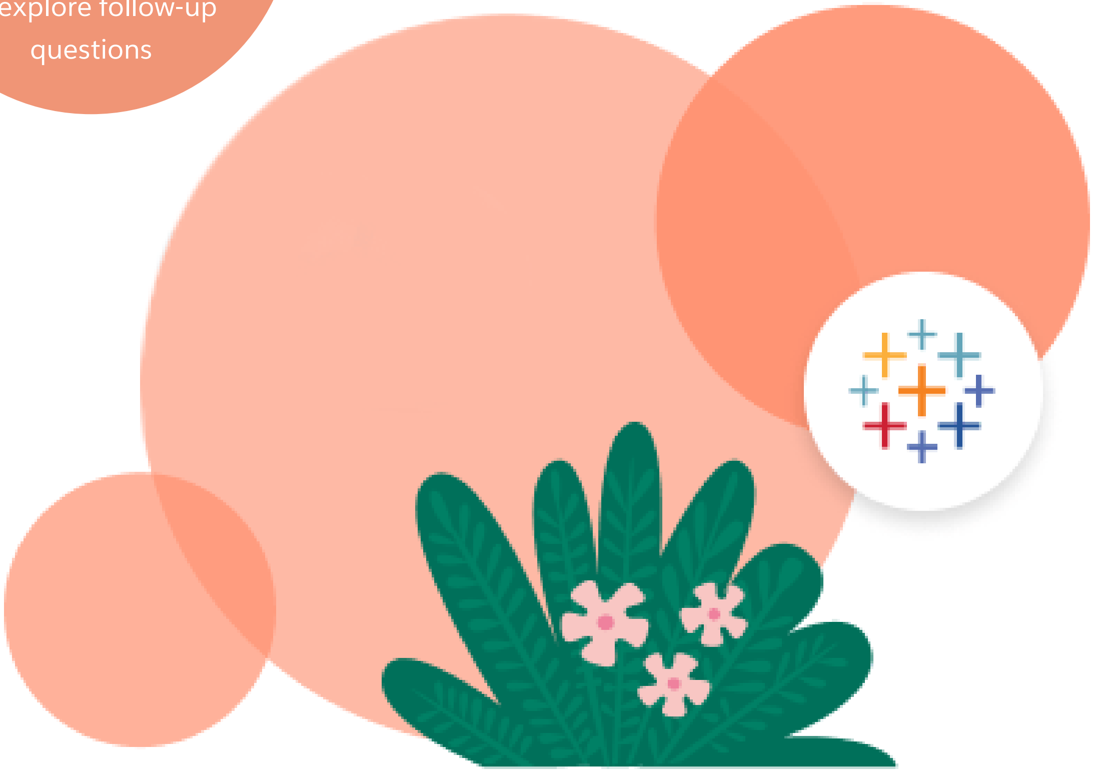

Tableau Pulse Data Flow

The Covered Services use industry-accepted encryption products to protect Customer Data and communications during transmissions between a customer's network and the Covered Services, including Transport Layer Security (TLS) 1.2 or newer. Tableau Pulse uses a native Salesforce multi-cloud to cloud (C2C) framework that enables secure, trust-based authentication using encrypted tokens as the go-between for logical tenants and Salesforce's multi-tenant applications.

Tableau Pulse communications pass through the Salesforce CRM LLM Gateway to the Einstein Trust Layer, as well as the external LLM provider, which maintains authentication over secure communication channels via TLS 1.2.

As part of the shared responsibility model, customer administrators may enable their authorized org users to log in directly or through a defined SSO with multi-factor authentication (MFA) implementation. Tableau Cloud supports System for Cross-domain Identity Management (SCIM), an open standard protocol for automating the exchange of user identity information between identity domains and IT systems, including Azure Active Directory, Okta, and OneLogin.

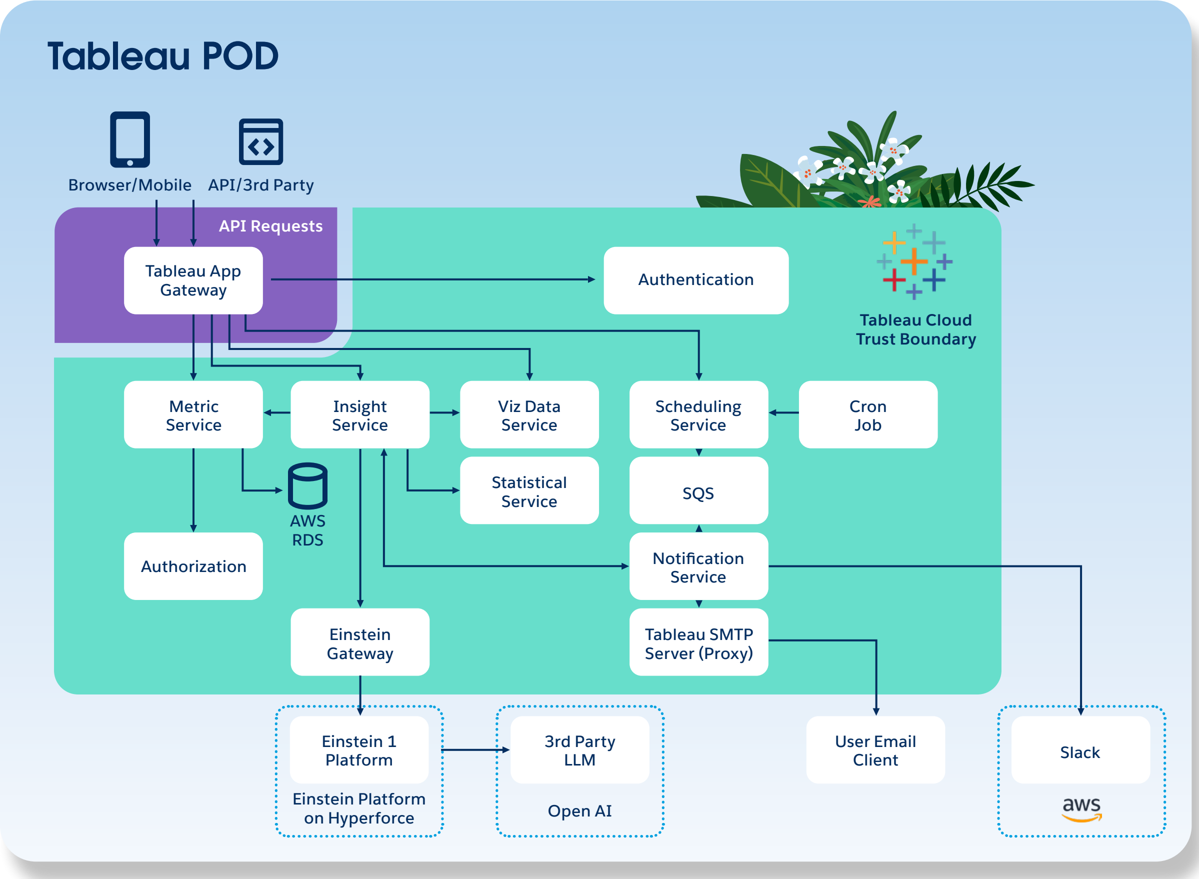

[9] The Einstein Trust Layer is referred to as the Einstein Platform on Hyperforce on the Compliance portal

Authorized users connect to the Tableau Cloud via a web application over a secure channel via TLS 1.2 using a 2048-bit certificate. The browser makes an authenticated request using the workgroup session ID token through the public API Tableau Application Gateway (TAG) to the Tableau Insights Service via a gRPC call over TLS 1.2 to request a Tableau Pulse summary. The Tableau Insights Service performs a cache service check via TLS 1.2 to determine whether a summary for that specific user exists. If the summary is present in the cache, it's returned and displayed to the user. If the summary doesn't exist in the cache, the Tableau UI triggers a new request to generate a summary. If a user doesn't follow any metrics, or no summary has been generated in the past 24 hours, a new request for a summary will begin.

The Insights Service makes a gRPC call over TLS 1.2 to obtain all metric subscriptions for the user in that particular Subscription Service. The Insights Service resolves scoped metric definitions by making a service request through a gRPC call over TLS 1.2 to the Metrics Service. The Insights Service makes a request to Tableau Viz Data Service to generate insights. Once all insights are generated, the service takes the top three highest-ranked insights based upon a usefulness score and forms a prompt to summarize the results of the content and call the Einstein GPT API Gateway.

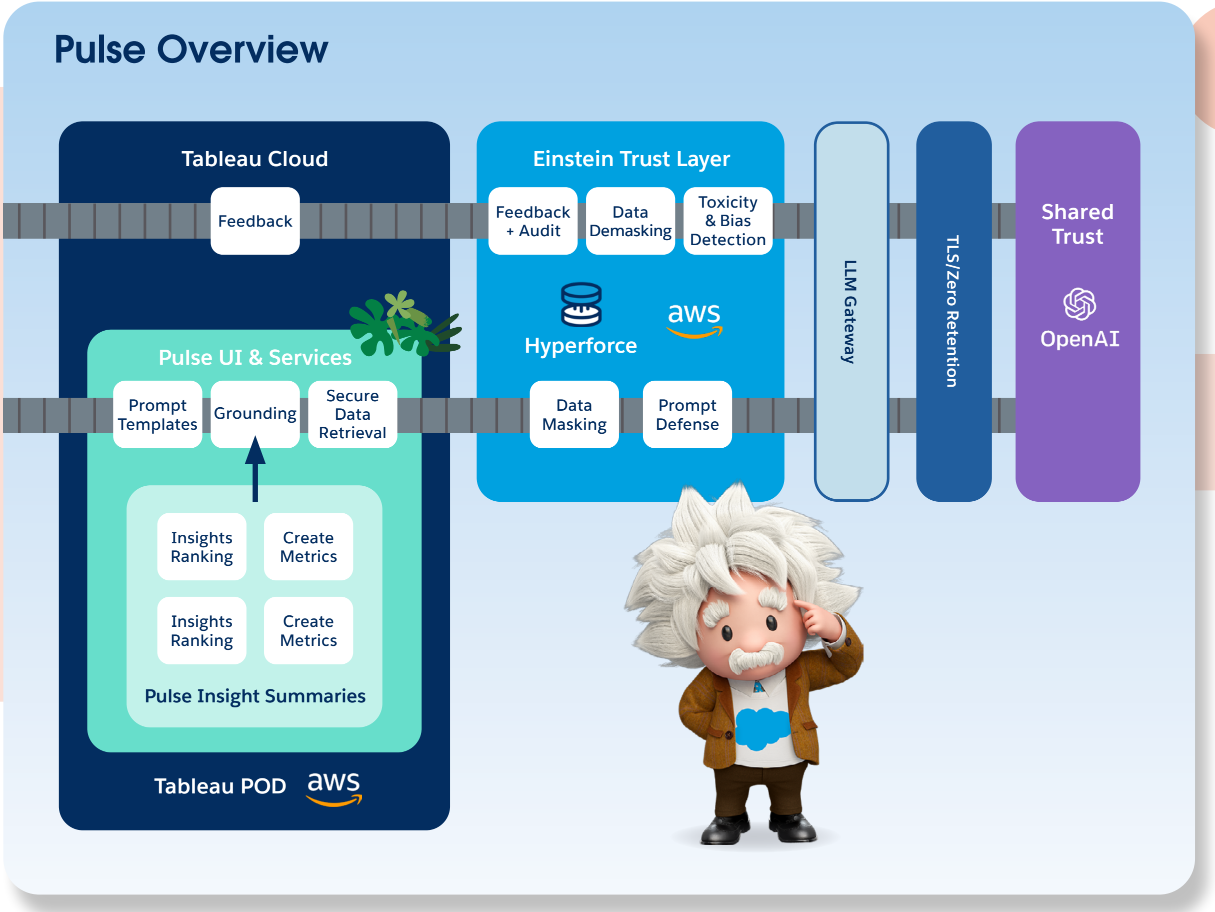

Since Tableau Pulse interacts with the Einstein Trust Layer for generative AI features, we'll walk through several security and privacy guardrails that are embedded in the service to ensure proper data protection and governance controls remain in place for the duration of each request.

Where applicable, these may include prompt templates, grounding, prompt defense, data masking, metering and rate limiting, zero data retention, toxicity detection, feedback framework, and an audit trail, which protect Customer Data for each Tableau AI request.

Prompt Templates and Grounding

Prompt templates provide a structured, versioned way of developing prompts to solve business process challenges with generative AI models. They include instructions that guide the LLM on what to generate, and merge fields to add data to the prompt at runtime. The process of enriching the prompt with relevant data is called grounding; this is the best way to provide the specific context of the business process to the LLM, and reduce the likelihood of undesirable hallucinations.

Grounding context can come from many data sources, including page context, Salesforce objects, SOQL queries and flows, as well as retrieval-augmented generation (RAG), which retrieves unstructured data such as knowledge articles pertinent to the prompt topic. Dynamic grounding pulls relevant data from the org to ensure that the data used remains current, accurate, and complete.

Currently, Tableau Pulse natively performs dynamic grounding within the Tableau Cloud environment rather than within the Einstein Trust Layer.

Prompt Defense

Salesforce maintains a secure software development lifecycle (SSDL), which follows the Open Worldwide Application Security Project (OWASP®) Top 10 and the MITRE Common Weakness Enumeration (CWE™) Top 25. When it comes to generative AI, Salesforce has been a strong contributor, following the OWASP Top 10 to ensure a secure LLM rollout. Salesforce has introduced specific protections for prompt defense that are designed and put in place to protect against LLM01:2023 (prompt injections). A variety of guardrails also help mitigate the risk of prompt injection type attacks. Such measures may include but aren't limited to:

- Instruction defense, which provides guidance on how to exercise caution regarding what follows a prompt.
- Post prompting, which positions user input near the beginning of a prompt.
- Prompt enclosure, which isolates user-supplied input and restricts it from referencing other parts of a prompt through encasing user data with various techniques, which may include:
    - Encasing prompts with randomly generated text sequences.
    - Prompt filtering, which identifies and blocks specific words or phrases from the prompt (using denylists) before it reaches the LLM to increase security.
    - Length restrictions, which help mitigate and thwart against Do Anything Now (DAN) type prompts and visualization attack vectors by controlling the length of user input.
    - Wrapping all prompts in data and text to reduce the possibility of harm to an end user or potentially causing hallucinations by returning incorrect responses.

Data Masking

In Tableau Pulse, Insights Summaries reference only the language associated with time-series insights (streaks, dips, trends, and so on). They never reference or express specific dimensional values, which ensures that no sensitive information is sent over the LLM. Data masking ensures that each request is protected through a variety of tools to replace sensitive data–including personally identifiable information (PII) and payment card industry (PCI) information–with placeholders before sending the request to the LLM provider. Further, Salesforce follows safety guidelines as defined in the Generative AI: 5 Guidelines for Responsible Development summary.

Metering and Rate Limiting

Salesforce protects the service against denial-of-service (DoS) type attacks through a variety of controls that aim to keep the service continuously available and prevent any one customer or application from impacting the overall service quality. The Salesforce SSDL adheres to the OWASP Top 10 and the CWE Top 25, and it requires proper input validation and sanitization for each request and interaction with the Einstein Trust Layer where authentication and authorization are required. Additionally, logging and continuous monitoring are in place to ensure proper functionality. Any detection of irregularities or anomalies are fully investigated to ensure that services remain in a consistent operational state.

Zero Data Retention

Communications between Tableau Cloud–Einstein Trust Layer and our shared trust external LLM partners are secured via API calls over TLS 1.2 connections. At no time does data leave the Salesforce secure shared trust boundary.

Salesforce maintains “zero data retention” agreements with its selected external LLM providers that include provisions whereby no Customer Data is retained, sent to downstream storage, or manually reviewed or accessed by third-party providers to train or improve models. At Salesforce, external third-party LLMs don’t maintain access to the underlying dataset.

The Salesforce external LLM provider continuously performs automated safety classifier scans on data sent to them to identify any violations of the respective law or external LLM safety policies. If a safety classifier is triggered, the external LLM provider must notify Salesforce of any repeated violations.

Currently, Salesforce doesn’t use Tableau Pulse Customer Data to train or improve our internal LLM. As demonstrated in the Salesforce Einstein 1 Platform on Hyperforce SOC 2 Report, Salesforce management performs a quarterly review of third-party LLM integrations to verify and ensure that they’re:

[10] The Einstein Trust Layer comprises a set of features/processes to provide trust on the Salesforce Einstein 1 Platform.

- Configured for zero retention of Customer Data.
- Validated to ensure that they don't store any Customer Data.

Toxicity Detection

In addition to the results returned from the external LLM provider, Salesforce also performs toxicity screening prior to sending the results downstream to the end user. Using an internally-developed AI model hosted by Salesforce on the Einstein Trust Layer, a toxicity detection scan is performed to generate an aggregate safety score that's calculated using various criteria from several unsafe toxicity categories that may include hate speech, violence, explicit content, and other forms of harmful material.

Feedback Framework

The Feedback Framework captures user feedback on LLM-generated responses, including explicit thumbs-up/down signals and implicit telemetry data. Feedback data can be used to identify issues that users experience, and to improve the effectiveness of prompting strategies. The agent options include the ability to voluntarily:

- Accept the output and results.
- Edit or modify the results.
- Ignore or reject the results.
- Select a thumbs-up or a thumbs-down.

Audit Trail

Salesforce maintains an audit trail where requests and responses are collected, such as time-stamped metadata that includes the context of the interaction with the LLM. The detailed history of each transaction may include:

- The requesting user, feature, prompt template, model, and other metadata.
- The original prompt, masked prompt, and masked and de-masked response from the LLM.
- The safety score and toxicity category scores.
- The feedback action performed by the end user, including whether the user accepted the output and results, edited or modified the results, ignored or rejected the results, or selected a thumbs-up or a thumbs-down to the response.

When Tableau Pulse customers use the generative AI features, each customer's user feedback and audit logs are saved in their secure, tenant isolated Data Cloud environment, under full control of the customer's Salesforce admins. In addition, Salesforce internally stores audit trail logs in a secure, tenant isolated manner, encrypted at rest via AES-256 bit encryption, for a retention period of 30 days.

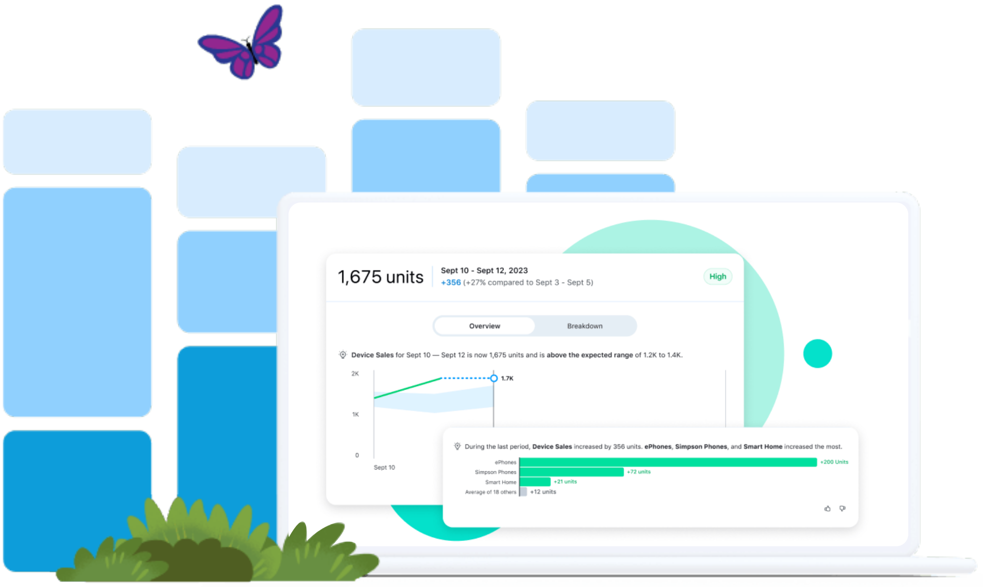

Device Sales for Sept 10 — Sept 12 is now 1,675 units and is **above the expected range** of 1.2K to 1.4K.

During the last period, **Device Sales** increased by 356 units. **ePhones**, **Simpson Phones**, and **Smart Home** increased the most.

End-User Notifications

If the Insights Summary was successful, the service validates the accuracy of the summary. The Insights Service caches the summary using a cache key–consisting of the site ID and the user ID–for a limited time.

After the generation of the insight, the notification service component bifurcates the workflow to dispatch both Slack and email digests based on the customer configuration.

- For email, digests are transmitted through a setup employing TLS for Simple Mail Transfer Protocol (SMTP), ensuring encryption of the digest during both transit and at rest.
- For Slack, digests are dispatched exclusively after a Tableau site administrator establishes integration between the site and the Slack workspace using Open Authorization (OAuth) 2.0. These digests are sent to the end user's Slack in an encrypted manner using the token obtained during the OAuth phase.

Now that we've covered the Einstein Trust Layer protection mechanisms in place, we'll review the underlying multi-cloud infrastructure that enables the secure delivery of the service, including Tableau Cloud, Einstein 1 Platform on Hyperforce, and the Salesforce Data Cloud.

Tableau Cloud

Global Availability to Scale with Customers

Tableau Cloud maintains a global data center hosting model using the Infrastructure as a Service (IaaS) provider Amazon Web Services (AWS). To support a growing customer base with unique data residency requirements, Tableau Cloud maintains a global regional presence across North America; Europe, the Middle East, and Africa (EMEA); and the Asia–Pacific (APAC).

For the most up-to-date information, review the Salesforce Infrastructure and Sub-processors documentation.

Infrastructure Access Control

All physical access to AWS locations is limited in nature and restricted (based on the business unit and role) to a subset of AWS employees. AWS maintains controls that test the design and operating effectiveness on a regular basis as stated in the AWS SOC 2 Report.

Tableau Cloud maintains strict logical access for remote access. MFA is required for all corporate remote access. Tableau Cloud maintains role-based access controls (RBAC) and follows the principle of least privilege (PoLP) and the need-to-know (NTK) principle. Only a limited number of Tableau engineering and security personnel maintain logical access to the production environment. For those users to access the production environment, they must first have the right privileges on various systems, including but not limited to their end-user asset, Tableau virtual private network (VPN), and Okta Advanced Server Access (ASA) (previously ScaleFT). Users must first successfully connect to Tableau VPN via a MFA solution, and then authenticate using Okta ASA to gain access to production via a hardened bastion or jump host that allows only for Secure Shell Protocol (SSH) and Remote Desktop Protocol (RDP) (all other services are removed). The Tableau Cloud bastion host acts as a jump server to permit a secure connection to instances provisioned without a public IP address. Okta ASA transparently enables SSH best practices for traversing bastion hops with a zero trust model.

In compliance with Salesforce information security policies, all Tableau administrator access to the Tableau Cloud production environment is encrypted in transit. Every connection between the SSH client and the target host–including bastion connections–maintains end-to-end encryption (via mutual authentication) authorized with ephemeral client certificates. In addition, only privileged Tableau Cloud infrastructure administrators require and maintain sudo access rights to the bastion host. All access attempts–successful and unsuccessful–are logged and fed into the Salesforce security information and event management (SIEM) solution. Reviews of logs and user access are performed on a periodic basis. Respective controls are tested for design and operating effectiveness within the Tableau Cloud SOC 2 Report.

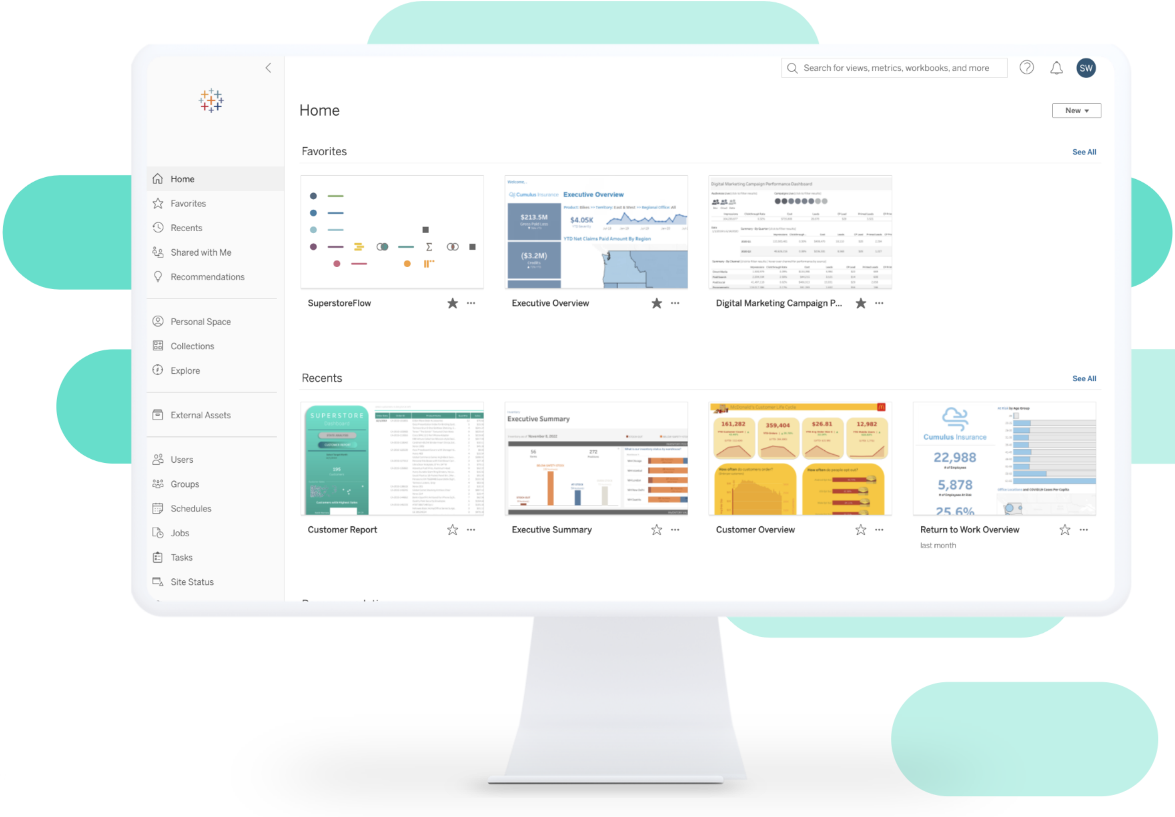

Data Encryption

Tableau Cloud maintains encryption in transit for all communications between a customer's network or web browser and the Tableau Cloud services using TLS 1.2 or higher.

Tableau Pulse data isn't stored outside of Tableau. All Tableau Cloud Customer Data is stored as Hyper extracts on AWS Elastic File System (EFS). Hyper is Tableau's in-memory Data Engine technology that optimizes fast data ingestion and analytical query processing on large or complex datasets. These Hyper extracts are encrypted at rest using the AWS-native EFS encryption that uses AES-256 bit encryption.

The Tableau Pulse feature uses a multi-cloud secure service delivery model designed to protect users' data during its interaction with the Einstein Trust Layer on the Hyperforce platform. Next, we'll explore the infrastructure and security fundamentals of the Einstein 1 Platform on Hyperforce.

Einstein 1 Platform on Hyperforce

Global Availability to Scale with Customers

Salesforce Tableau AI uses Einstein generative AI technology. The Salesforce Einstein 1 Platform on Hyperforce uses AWS as the underlying IaaS provider. The Einstein 1 Platform is currently available in North America, EMEA, and APAC, with hosting in the United States, Germany, India, Australia, and Japan.

For the latest information, check out Where are the Einstein 1 Platform on Hyperforce Instances Located?

Infrastructure Access Control

All physical access to AWS locations is limited in nature and restricted (based on the business unit and role) to a subset of AWS employees. AWS maintains controls that test the design and operating effectiveness on a regular basis as stated in the AWS SOC 2 Report.

The Einstein 1 Platform on Hyperforce environment is physically and logically separate from the Salesforce corporate IT and research & development (R&D) environments. Successful MFA is required to gain access to the Salesforce corporate IT environment, and it's a prerequisite for Einstein 1 Platform engineering and security staff to gain logical access to the production environment. Strict logical access requirements are in place for any remote access to the production environment. Access requirements include but aren't limited to a completed, successful background check, logical access management approval based on role or job function, and business unit alignment with organizationally defined RBAC. Access rights follow the PoLP and NTK principles where just-in-time (JIT) access to production environment resources is in place for all privileged commands within the production environment. Additionally, all access requires a unique MFA device assigned to the user plus adherence and compliance to information security standards, including the Salesforce Global Change Management and Cryptographic Key Management responsibilities.

Data Encryption

The Einstein 1 Platform on Hyperforce follows a zero trust model and the concept of "never trust, always verify." Each required human or service access request–regardless of location or device–remains untrusted until it successfully authenticates and is authorized for requests through contextual demonstration of security configuration and posture. All communication transmissions between Tableau Cloud Pulse and the Einstein Trust Layer services use TLS 1.2 or higher. Within the Einstein 1 Platform on Hyperforce, encryption in transit is achieved via a service mesh architecture that supports mutual TLS (mTLS). Customer Data stored in the Einstein 1 Platform on Hyperforce maintains data stores that are encrypted at rest via AES-256 bit encryption using customer- (Salesforce)-managed keys (CMKs) that are configured for annual rotation. As part of the shared responsibility model between Salesforce and AWS, encryption keys are fully managed via AWS as the underlying IaaS provider, and they aren't visible to Salesforce. AWS maintains the responsibility to manage the rotation of CMKs through the AWS Key Management Service (KMS).

Next, we'll explore the infrastructure and security fundamentals of the Salesforce Data Cloud.

Salesforce Data Cloud

Global Availability to Scale with Customers

The Salesforce Tableau AI uses Einstein generative AI technology, which stores audit trail logs in a tenant isolated manner within Salesforce Data Cloud environment data lake objects. The Salesforce Data Cloud on Hyperforce environment uses AWS as the underlying IaaS provider. Currently, the Customer Data Cloud supports data-residency options in North America, EMEA, and APAC, with locations in the United States, Germany, India, Japan, and Australia.

Infrastructure Access Control

All physical access to AWS locations is limited in nature and restricted (based on business unit and role) to a subset of AWS employees. AWS maintains controls that test the design and operating effectiveness on a regular basis as stated in the AWS SOC 2 Report.

The Salesforce Data Cloud on Hyperforce environment is physically and logically separate from the Salesforce corporate IT and R&D environments. Successful MFA is required to gain access to the Salesforce corporate IT environment, and it's a prerequisite for the Salesforce Data Cloud platform engineering and security staff to gain logical access to the production environment. Strict logical access requirements are in place for any remote access to the production environment. Access requirements include but aren't limited to a completed, successful background check, logical access management approval based on role or job function, and business unit alignment with RBAC. Access rights follow the PoLP and NTK principles where JIT access is in place for all privileged commands within the production environment. Additionally, all access requires a unique MFA device assigned to the user plus adherence and compliance to information security standards, including the Salesforce Global Change Management and Cryptographic Key Management responsibilities.

Data Encryption

When Tableau Pulse customers use the generative AI features, customer user feedback and audit logs are saved in a tenant isolated manner within the Salesforce Data Cloud data lake objects. All communications between the Einstein Trust Layer on Hyperforce and Data Cloud are encrypted in transit using TLS 1.2.

Customer Data in Salesforce Data Cloud data stores are encrypted at rest via AES-256 bit encryption keys using Salesforce CMKs, which are configured for annual rotation. The infrastructure engineering team relies on the unique Salesforce-managed CMK that's assigned to each data store. As part of the shared responsibility model between Salesforce and AWS, encryption keys are fully managed via AWS as the underlying IaaS provider, and aren't visible to Salesforce. AWS maintains the responsibility to manage the rotation of CMKs through the use of the AWS KMS.

Customer-facing audit trail logs are encrypted at rest via AES-256 bit encryption and retained internally within the Salesforce Data Cloud environment for a period of 30 days.

As demonstrated throughout this overview, the protection of Customer Data is paramount to maintaining customer trust. Next, we'll review the ongoing nature of our threat and vulnerability management program.

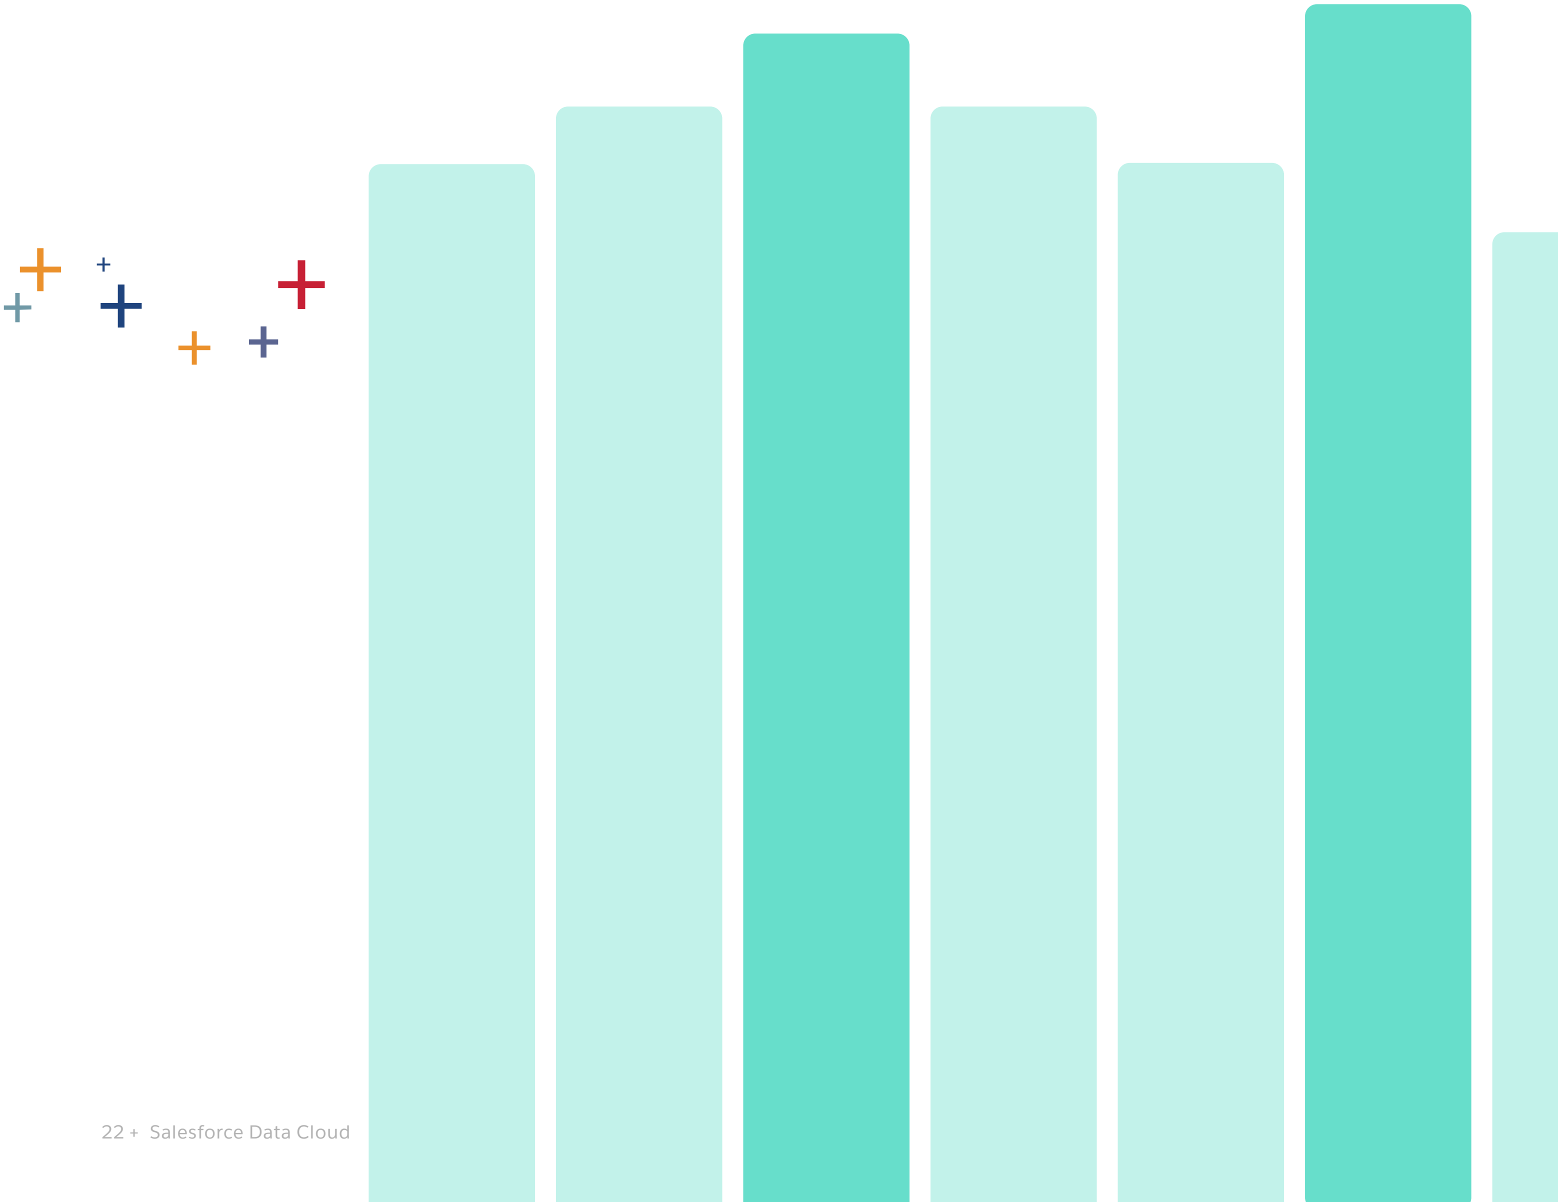

Vulnerability Management

Internal Assessments

Salesforce uses a variety of techniques to perform internal vulnerability management. Salesforce built an SSDL, which integrates security requirements throughout the software development and deployment process following industry-best secure coding practices, including the OWASP Top 10 and the MITRE CWE Top 25. Operating on the concept of Secure by Design, Salesforce applies its SSDL across all of our products and services from initial ideas through feature releases. As part of the SSDL, Salesforce employs threat modeling through security assessments and uses a variety of tools to identify the security risks, threats, and vulnerabilities of the proposed design early in the development process to create security mitigations against our requirements.

Regularly run security tools may include but aren't limited to:

- Credential scanning tools to examine source code for sensitive information, including credentials and secrets.
- Open source software (OSS) vulnerability detection and management scans to help identify and detect vulnerabilities at an early stage in the development process.
- Static application security testing (SAST) to scan uncompiled application source code for vulnerabilities by searching for improper use of dangerous APIs, or tracing flows of tainted data through the application (where applicable).
- Dynamic application security testing (DAST) through continuous 24/7 web application security scans.
- Container scans to perform static security scans at the app or operating system layer of container images.
- Third-party product container scans to look for vulnerabilities in operating systems within containers and third-party components and libraries.
- Weekly authenticated and unauthenticated internal and external vulnerability scans performed across environments to address security misconfiguration components with known vulnerabilities.

The Salesforce Product Security team works closely with our development teams to ensure the security and quality of code that goes into the production environment. Separately within the security organization, a fully staffed internal red team tests our overall security program–which covers advanced persistent threat (APT)-style attacks–including penetration tests on Salesforce production, development, and corporate environments.

For the most recent threat and vulnerability management control testing, check out the Corporate Services SOC 2 Report.

Third-Party Assessments

A variety of third-party assessments are conducted against the Salesforce security program and products on an annual basis. These may include but aren't limited to:

- Continuously reviewing and testing the Salesforce security control framework.
- Additional penetration tests designed and executed to replicate real world attack scenarios for certain high-risk features.
- Third-party application security assessments performed on an annual basis and when deemed necessary by qualified assessment firms that may hold a variety of certifications, including but not limited to:
    - PCI Qualified Security Assessor (QSA)
    - PCI Approved Scanning Vendor (ASV)
    - Member of the Council of Registered Ethical Security Testers (CREST)
    - Registered under the U.K. National Cyber Security Centre (NCSC)'s Certified Cyber Security Consultancy (CCSC) scheme
- Quarterly network vulnerability assessments performed (where applicable) to meet security compliance obligations.

For the most recent third-party application security assessments, check out the External Security Assessments report.

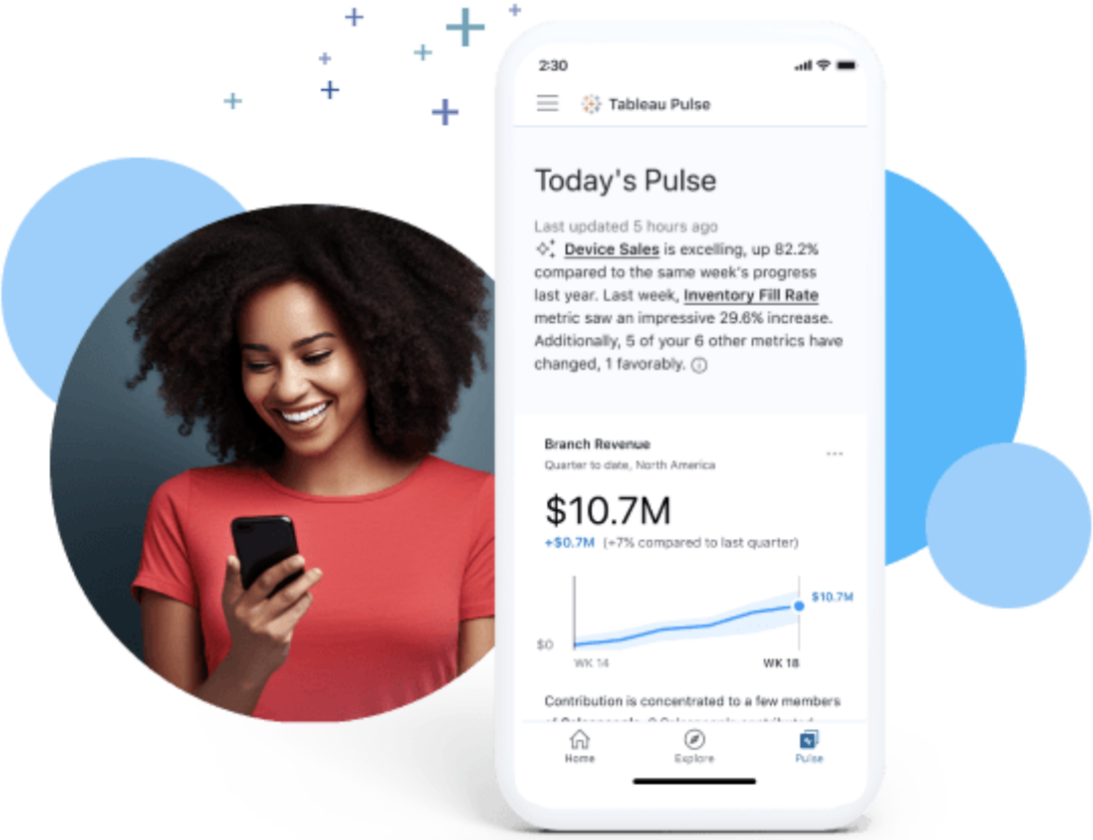

2:30

Tableau Pulse

Today's Pulse

Last updated 5 hours ago

Device Sales is excelling, up 82.2% compared to the same week's progress last year. Last week, Inventory Fill Rate metric saw an impressive 29.6% increase. Additionally, 5 of your 6 other metrics have changed, 1 favorably.

Branch Revenue

Quarter to date, North America

$10.7M

+$0.7M (+7% compared to last quarter)

Contribution is concentrated to a few members

Third-Party Risk Management (TPRM)

Third parties contracted by Salesforce are required to commit to confidentiality agreements covering Customer Data. All third parties are subject to Salesforce policies and procedures as defined in the company Third-Party Suppliers standard and other policies. These include background screening, training, breach of policy, and enforcement. Prospective vendors supporting production services are assessed for their security, compliance, and privacy practices prior to signing contracts for services. These third-party vendors are also evaluated by the Salesforce Compliance team prior to launch.

Deficiencies with potential access to Customer Data (where applicable) are noted in the review and remediated or compensated using control(s) identified to address key risks prior to launch. Once a third party is evaluated and onboarded, periodic assessments are performed on an ongoing basis, depending on the criticality of the third party.

The TPRM team performs security and compliance assessments, which consist of an in-depth review of the vendor's security posture in line with the services provided to Salesforce. This process is driven by internal security standards and best practices, as well as regulatory requirements that include but aren't limited to the System and Organization Controls (SOC), International Organization for Standardization (ISO), Payment Card Industry Data Security Standard (PCI DSS), Federal Risk and Authorization Management Program (FedRAMP™) certifications, Health Information Trust Alliance (HITRUST), Health Insurance Portability and Accountability Act (HIPAA), and General Data Protection Regulation (GDPR).

The third-party assessment determines whether the security controls are effective or ineffective. Even if the security controls are determined to be effective and they meet Salesforce requirements, additional observations, opportunities for improvement, or discussion topics may still be identified and mitigation may be required. If security controls are determined to be ineffective and don't meet the Salesforce requirements, they're identified as findings and mitigation is required. For ineffective security controls, findings are ranked against the likelihood of a risk occurring and the impact to the business. All findings are assigned a risk ranking from very low (opportunity for improvement) to critical and monitored through mitigation. Mitigation timelines are defined based on risk ranking.

Annual reviews are performed for all critical vendors, including data centers and sub-processors. TPRM teams continuously monitor vendors for changes, which may include:

- Security control implementation.
- Third-party infrastructure or applications.
- Regulatory compliance environments and updated requirements.
- Services provided to Salesforce.

Should any of these changes occur, a reevaluation of the third party's security and compliance risks may be required. Additional procedures are also performed. Contracts may be reviewed as required to ensure appropriate agreements are included for privacy certifications and audit rights.

Salesforce security teams perform a thorough internal security assessment of the external LLM provider services prior to permitting them to be used in or with Salesforce products.

Salesforce third-party vendor audit program methodologies and processes are regularly audited by Salesforce third-party auditors. For more information, review the vendor audit program controls in the Corporate Services SOC 2 Report and the Salesforce Third-Party Risk Management Overview.

The TPRM team performs security and compliance assessments, which consist of an in-depth review of the vendor's security posture in line with the services provided to Salesforce.

Conclusion

At Salesforce, we build security into every aspect of our business. The Tableau Pulse product is built upon our #1 core value, Trust, because our customers depend on us to safeguard and protect their Customer Data and to ensure that our platform and services consistently meet our performance and security requirements. For additional details and information, review the following Salesforce Security Compliance Documentation section or contact your Salesforce account executive.

Thank You!

Salesforce Security Compliance Documentation

Tableau Cloud:

Compliance
AWS SOC 2 Report
Corporate Services SOC 2 Report
SOC 2 Report
Vulnerability/Penetration Report Summary

Einstein 1 Platform on Hyperforce:

Compliance
Einstein GPT Security White Paper (EN)
SOC 2 Report
Where are the Einstein Platform on Hyperforce Instances Located?

Einstein Trust Layer

Help Documentation
Trailhead: Meet the Einstein Trust Layer

Salesforce Security Documentation

External Security Assessments
Salesforce Secure Development Lifecycle Overview
Salesforce Third-Party Risk Management Overview
Salesforce Vulnerability Management Program Overview
Security Perspective on the Shared Responsibility Model

Salesforce AI Articles

How Salesforce Develops Ethical Generative AI from the Start
Meet Salesforce's Trusted AI Principles
Generative AI: 5 Guidelines for Responsible Development
Why Salesforce Aims to Build Products That Are 'Ethical by Design'

Salesforce AI Acceptable Use Policy

PDF

Trust and Compliance

Documentation
Salesforce Infrastructure and Sub-processors

Data Processing Addendum

PDF

Salesforce has made a good faith effort to provide you with responses to your request for information that are accurate as of the date of the response and within our knowledge. Because our procedures and policies change from time to time, we cannot guarantee that the answers to the questions you have asked will remain the same over time. The information provided here is for informational purposes only, and the rights and responsibilities of the parties with regard to your use of Salesforce's online software services to which this request for information relates shall be solely as set forth in the applicable mutually executed subscription agreement(s) or online Salesforce terms of service between the parties.