




Driving technology transformation with Trust: A Financial Services Perspective

A PwC Australia and Salesforce discussion paper on how to address your technology risk and regulatory requirements with a shared responsibility model for financial services organisations in Australia



Contents

About Salesforce	03
About PwC Australia	04
Disclaimer	05
Scope limitations	05
Overview	06
The Australian markets technology story.	08
How the Australian regulatory environment is evolving to keep pace with changes	09
Shared responsibility model with vendors and regulators	10
Navigating the implementation of a compliance-focused Salesforce platform.	14
AI advancements and the regulatory impact	16
What next?	18
Appendix 1: Mapping of example Salesforce capabilities and practices of representative FS regulatory requirements	20
Appendix 2: Mapping of Salesforce capabilities to Australian Voluntary AI Guidelines	32



As technology rapidly advances, including the rise of AI technologies, the need for secure and responsible innovation has never been greater. At Salesforce, we embed Secure by Design and Privacy by Design principles into everything we build, ensuring that innovation and compliance go hand in hand. Protecting customer data is at the core of our mission, and we are committed to partnering with our customers to help them harness AI-driven advancements while meeting their business and regulatory obligations with confidence.”

CHETAN SANSARE

SENIOR DIRECTOR SECURITY AND REGULATORY
COMPLIANCE APAC

About Salesforce

Salesforce, Inc. is a cloud-based software firm that provides an all-in-one platform of customer relationship management software and applications focused on sales, customer service, marketing automation, analytics, integration, artificial intelligence (AI) and application development. Salesforce’s Software as a Service (SaaS) offerings include:

- 1 **customer relationship management (CRM)** for businesses
- 2 **connecting sales, service, marketing, commerce and IT** on a single platform while personalising experiences
- 3 **data and analytics** to visualise and analyse data
- 4 platforms to **build and run business applications** and deliver services over the internet, through commercially available web connections, browser software, and mobile devices
- 5 **integrated platform with AI and app development**
- 6 **The customer 360 platform** integrates all the offerings above onto a single CRM platform.

[For more information on the products and services refer to the website](#)



Financial services organisations in Australia are at a critical junction where modern technology allows for a focus on customer and employee experience, but also lowers the barriers to disruption. In a regulated industry, leaders need to carefully navigate the transformation journey to stay ahead of competition, deliver the best experiences and uphold stakeholder trust. Having a clear understanding of the accountabilities across your entire ecosystem allows you to safely accelerate transformation.”

NOEL WILLIAMS

PWC AUSTRALIA BANKING AND CAPITAL MARKETS LEADER

About PwC Australia

PwC Australia is a professional services firm with a comprehensive suite of services across assurance, tax and advisory.

At PwC we take a human-led, tech-powered approach, combining diverse perspectives, expertise and relationships with the right technology to solve complex problems and unlock opportunities.

We're part of a global network of firms spanning 149 countries with more than 370,000 people. Our services include artificial intelligence, assurance, digital transformation, deals, tax, consulting, cybersecurity and digital trust.

In Australia, our team of more than 7,000 people has deep expertise in the industries critical to our nation, including: energy, utilities and resources, financial services, healthcare, education and consumer markets.

[For more information on the products and services refer to the website](#)



Disclaimer

This discussion paper is intended for informational purposes only and should not be relied upon as professional advice. It is not tailored to the specific facts and circumstances of any potential client and does not constitute the delivery of professional services.

This document aims to address frequently asked questions and concerns, and help financial services organisations utilising Salesforce services with a suggested approach to managing regulatory compliance and risk responsibilities across their respective organisations. It provides thought leadership on technology transformation in Australia and the interaction with the regulatory environment. It encourages readers to consider how a shared responsibility model can assist in compliance and risk management. It also outlines examples of how Salesforce demonstrates the measures they have implemented to maintain the confidentiality, integrity, availability, and privacy of their customer's data.

Note that this document does not cover all, or all aspects of, regulations and should not be considered exhaustive, nor should it be considered legal advice. PwC and Salesforce strongly encourage customers and prospective customers to consult with their own legal counsel and risk management teams to determine the applicability of relevant laws and regulations. This document was first published in May 2025 and may not reflect regulatory changes after this date.

The information provided in this discussion paper is without any warranty or guarantee of any kind, whether express or implied. Salesforce and PwC shall not be liable for any damages arising from the use of the information in this document. This discussion paper and its contents are protected by intellectual property laws. Unauthorised use or distribution is prohibited.

This discussion paper may include information from third-party sources, and Salesforce and PwC are not responsible for the accuracy or completeness of such information.

Salesforce and PwC reserve the right to update or revise this discussion paper at any time without prior notice.

Throughout this document, Salesforce has highlighted various services that customers can choose to adopt to enhance risk mitigation in Salesforce implementation.

Scope limitations

This document is specific to Salesforce services as defined in the security, privacy and architecture (SPARC) documentation referenced [here](#). Salesforce maintains the SPARC documentation as relevant to their current services.




Overview

Australian organisations find themselves at the precipice of significant change. As well as rapid social, economic and geopolitical changes, Australian leaders find themselves challenged to evolve and re-shape their business strategies to adapt to emerging technology and customer preferences.



62%

of Australian CEOs
believe **macroeconomic
volatility and policy** is a
key threat to their organisation



40%

of all Australian CEOs
expect their **investments in
technology to increase profits**
in the year ahead



56%


of Australia's CEOs
predict **AI will be integrated
into core business strategy**
in the next three years

According to [PwC's 28th annual CEO Survey](#), over half (62%) of all Australia's CEOs view macroeconomic volatility and policy as major threats to their organization. Despite these concerns, there is a strong sense of optimism regarding technology investments, with 40% of CEOs anticipating profitable returns from these investments in the coming year. Further reinforcing this positive outlook, 56% of Australia's CEOs expect artificial intelligence to become a fundamental component of their core strategies within the next three years. Naturally, choosing what to prioritise and where to invest is top of mind for business leaders.

For financial services organisations, this is particularly evident, fragmented legacy technology networks, siloed information across business units, a history of challenging technology migrations, and the pace of disruptive technology innovation make the transformation journey more complex and harder to navigate. Technology-led transformations are further complicated by the rapidly evolving regulatory risk landscape, which is itself trying to keep up with the rapid pace of innovation. Organisations require reliable partners and third-party risk management discipline to ensure technology implementations deliver value and do not come at the risk of customer data security and regulatory scrutiny. Corporate leaders across industries also anticipate that the magnitude of changes ahead will require their organisations to adopt a continuous transformation approach if they are to thrive in the coming decades. With the right strategic support, leaders can convert these obstacles into opportunities and meet the future head-on.

This document outlines the technology and regulatory trends organisations should heed as they look to implement their Salesforce solutions. It focuses on practices that should be considered when determining an organisation's risk-appropriate control environment and when engaging with internal risk teams and regulators. It includes examples of measures taken by Salesforce in relation to: technology infrastructure; cyber security, integrity and availability; and the privacy of their customers' data.

Salesforce has highlighted throughout this document additional capabilities and services designed to enhance the security and control environment of the Salesforce platform. Considering risk management needs and regulatory obligations early will inform an organisation's choice on the need to implement these capabilities and services.



60%

of CEOs expect Gen AI to significantly **change the way their company creates, delivers and captures value**

Source: [PwC's The Reinvention Imperative: Why CEOs need to face into technological disruption](#)



The Australian technology story

Since the start of the decade, there has been increasingly rapid adoption of Software as a Service (SaaS) solutions. SaaS has more recently seen a second wave of demand driven by the need to replace ageing software and infrastructure, and a race to adopt advanced generative and agentic artificial intelligence (AI)—especially from early adopters such as organisations in the financial services, telecommunications, retail and healthcare industries. Growth-minded leaders are recognising the potential for SaaS to leapfrog traditional technology development pipelines and address the longstanding challenge of maintaining technology currency. SaaS not only offers greater potential for differentiation; it can enable business model reinvention through the rapid deployment of digital capabilities at scale. SaaS uses vendors to deliver commoditised functionality—for example, customer service, digital channels and AI-enablement—allowing organisations to focus on differentiating customer offerings and digital experiences. The increasing need for digitisation, the acceleration of technological change, increased productivity demands, and talent shortages are all driving next-generation business models to incorporate SaaS providers as key partners.

Australia has emerged as a leader in the Asia-Pacific region's public cloud market, which is growing by approximately 18% is projected to reach \$15.36bn in 2025, largely driven by the SaaS segment. However, Australia's growth in cloud capabilities trails other developed markets. This could be for a number of reasons. For instance, one in four Australian companies have recently experienced a cyber security incident, up from one in ten in 2020. This has led to concerns around customer data protection and intellectual property loss—and an increased regulatory response. The rise in cyber incidents has left highly regulated industries, such as financial services, cautious. Risk management functions have been tasked with enhancing governance in the use of SaaS, and have become key decision-makers in the selection of technology solutions.

As SaaS adoption continues to increase and scale, we can expect a commensurate uplift in regulatory scrutiny of SaaS, which means risk management requirements of service providers and user organisations will need to continue to mature.





How the Australian regulatory environment is keeping pace with changes

The rapid pace of innovation in technology has kept regulators around the world, including in Australia, alert to ensuring organisations do not compromise customer trust or national infrastructure to meet their business goals.

Many of these regulations, such as the Australian Prudential Regulation Authority's (APRA) CPS 230 Operational Risk Management, are expected to have a transformational impact on business processes. The federal government is also planning a significant overhaul of the Privacy Act, including long-term plans for the introduction of a data subject's right to have all their personal data deleted, and stricter rules around the use of targeted advertising. As at the time of publishing, the first stage of reforms has begun, with draft legislation to require greater disclosure regarding the use of automated decision-making, a tiered penalty regime and children's privacy codes. Finally, the Security of Critical Infrastructure Act aims to strengthen the protection of Australia's critical infrastructure and mandate better cooperation between government and enterprises.

The rising dependence of enterprises on cloud solutions has attracted more attention from regulators seeking to ensure the integrity and sanctity of the financial ecosystem. APRA has laid out a robust framework for regulated entities to identify, assess and effectively manage operational risks, with a special

emphasis on outsourcing to material service providers. APRA's framework generally requires regulated entities to implement a mature governance practice to mitigate the risks of severe, yet plausible, adverse scenarios, and to maintain related tolerance levels for critical processes during potential disruptions. This includes arrangements where processes are delivered across in-house resources and third-party material service providers (e.g. SaaS providers), requiring a close coupling of resilience and information security arrangements between vendors and user organisations.

Due to the volume and nature of today's disruptive technologies, APRA regulations tend to be principles-based, as opposed to prescriptive. In response, enterprises have developed—and continue to evolve—their third-party risk management frameworks and systems to mitigate operational risks (such as cybersecurity, resilience and privacy) within their own risk appetite.

For organisations with global operations, rapidly evolving global regulations are also a key factor for consideration.

In this ever-changing environment, staying on the right side of regulators will require nimbler and highly informed operational processes that appropriately interpret guidance across all facets of the enterprise's technology infrastructure.



Shared responsibility model with vendors and regulators

For most organisations, the model of working with third parties has evolved from a ‘cost-optimised delivery’ partnership to a more strategic partnership model that entrusts the large-scale enablement of core competencies to the SaaS provider. This decentralisation of capability leads to significant challenges in maintaining trust between regulators, user organisations, third parties and end-customers. Organisations are embracing ‘resilience by design’ as a core tenet to deliver critical services reliably. This philosophy is underpinned by a comprehensive understanding of the roles and responsibilities—and dependencies—between business partners.

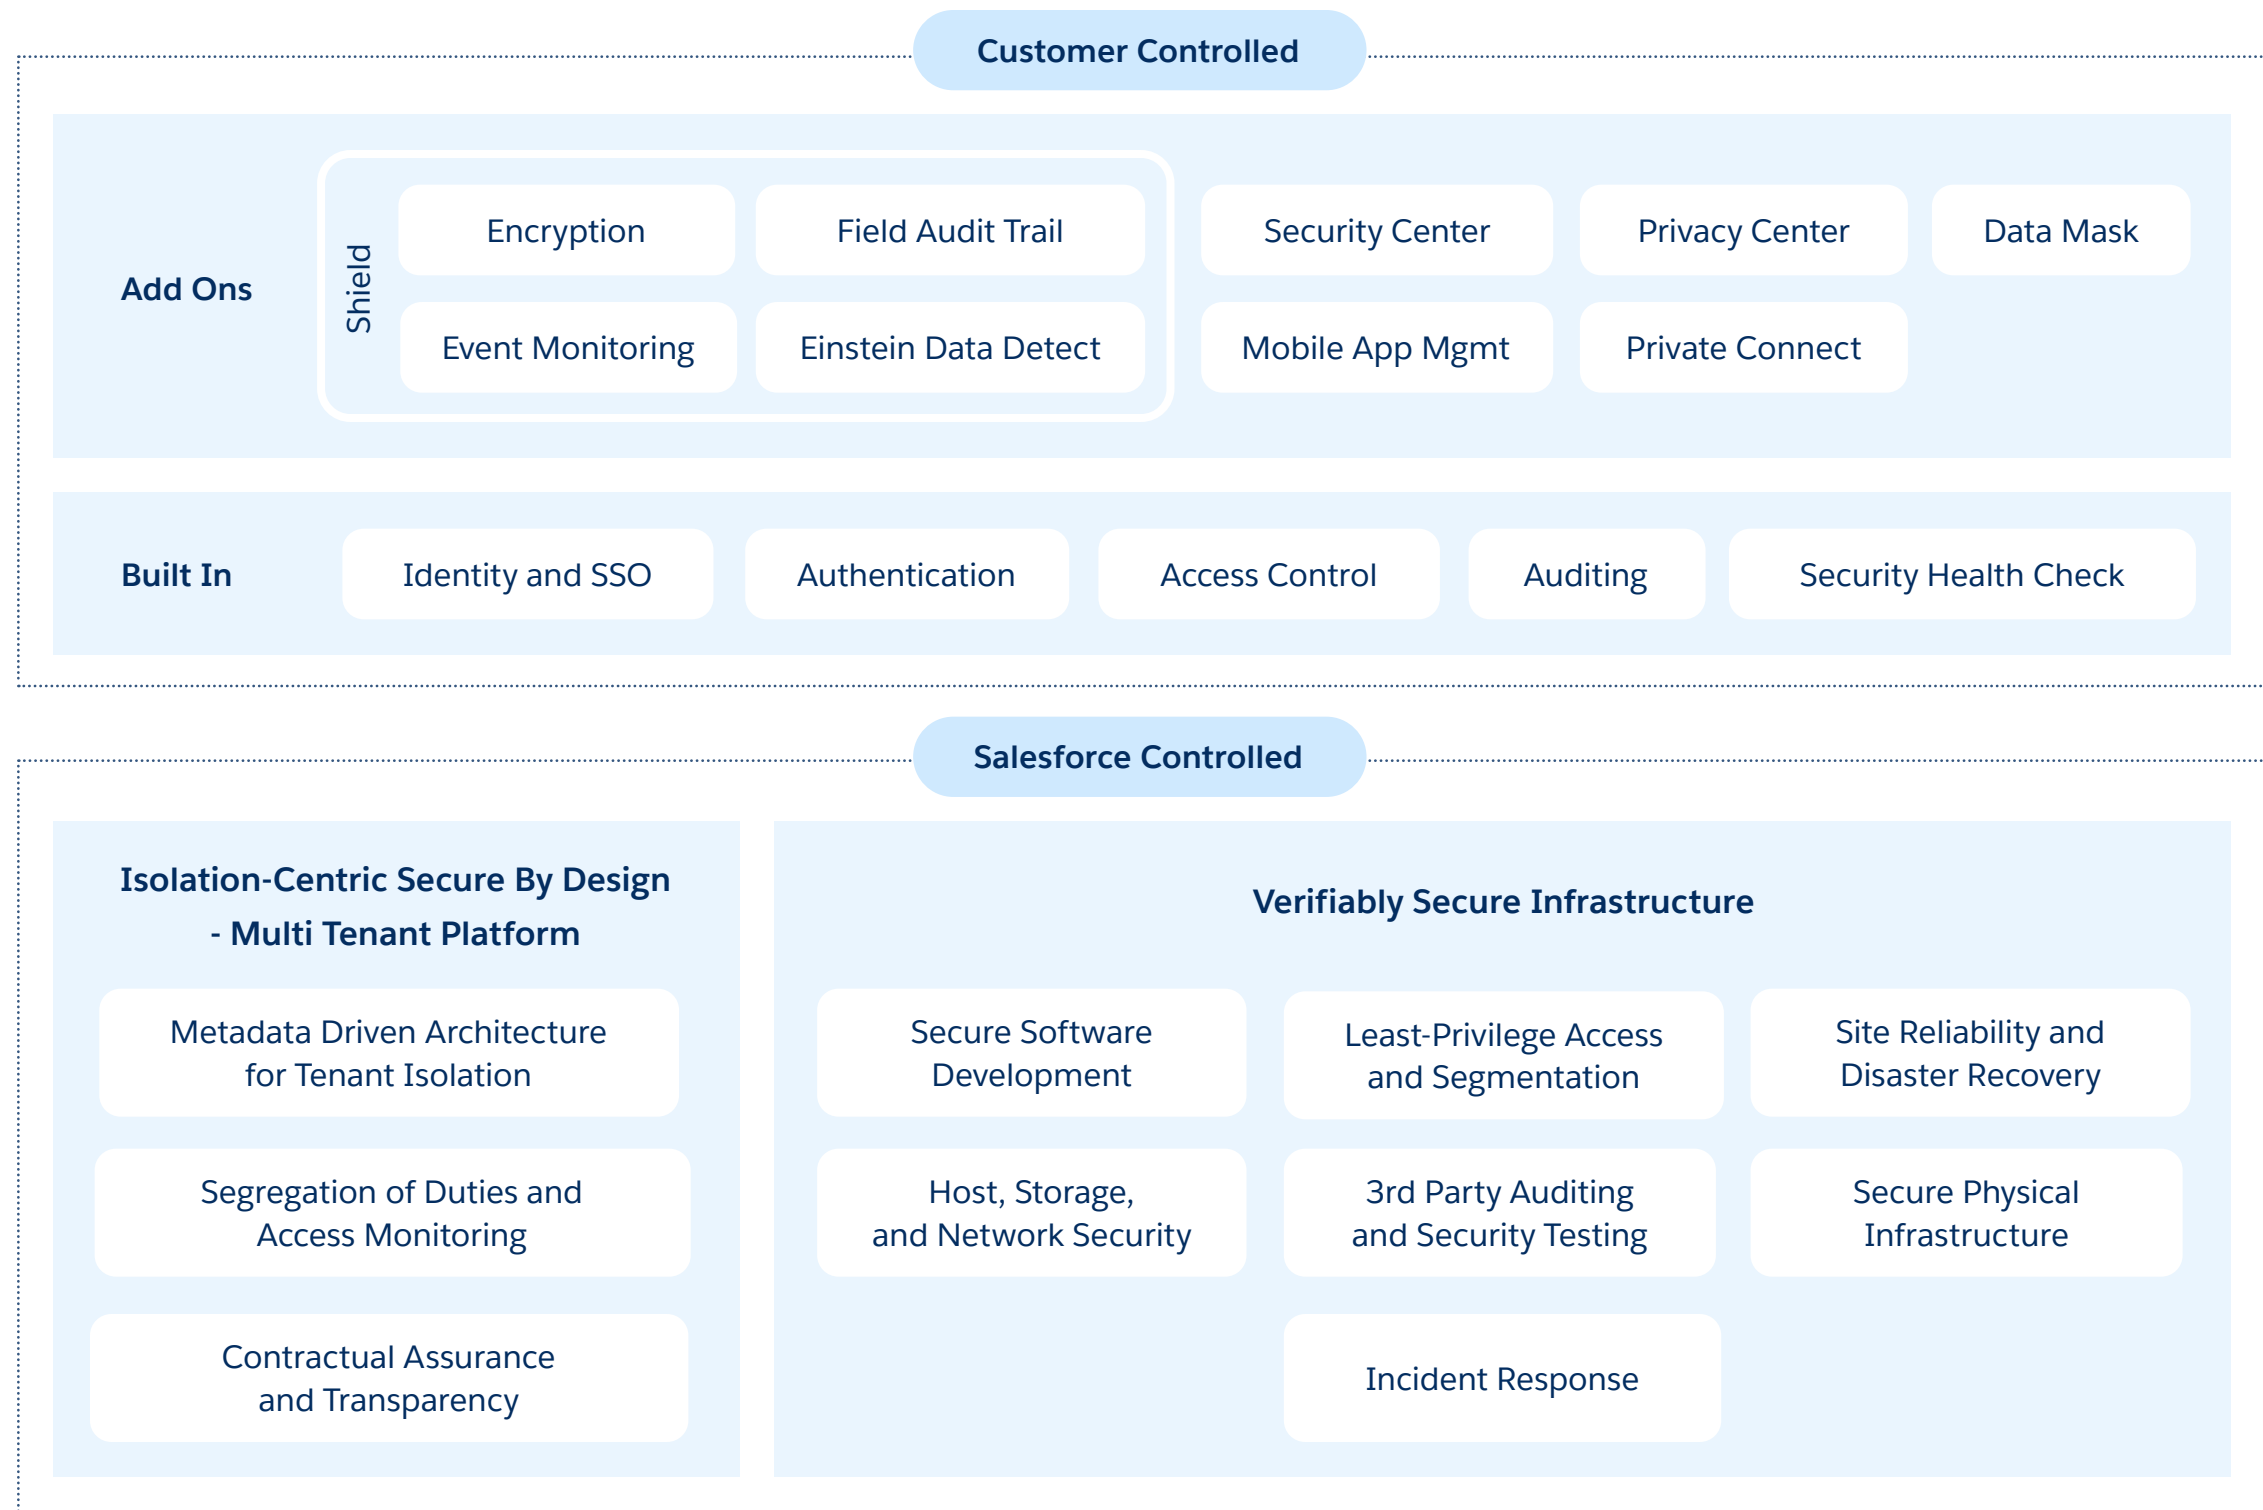
Australian regulators have emphasised that accountability lies with the regulated organisation, regardless of whether aspects of a service are outsourced. Current regulations and guidance, such as CPS 230, have further clarified that while a third-party vendor may be trusted, and significantly scaled, to provide an ongoing service, the risks relating to those functions are still the responsibility of the regulated entity. However, the fact remains that there is some shared responsibility between the parties. This is called a ‘shared responsibility model’. For example, in a shared responsibility model with SaaS vendors, such as Salesforce, enterprises are responsible for elements such as the configuration and utilisation of services, user access, and data protection within their environment.

The vendor is responsible for maintaining core infrastructure, managing the availability and integrity of the system, and for platform security. Having clarity on roles and responsibilities is essential for meeting compliance requirements, fostering collaborative risk mitigation and maintaining stakeholder trust.

The diagram shown on the following page provides a view of how risk management capabilities are enabled between Salesforce and a user organisation. Some of these capabilities are available to all user organisations, while others can be optionally ‘added-on’ for enhanced controls. Similarly, the ongoing operation and effectiveness of these controls is monitored through various mechanisms, for example a controls assurance report provided by Salesforce or by actions taken directly by a user organisation.



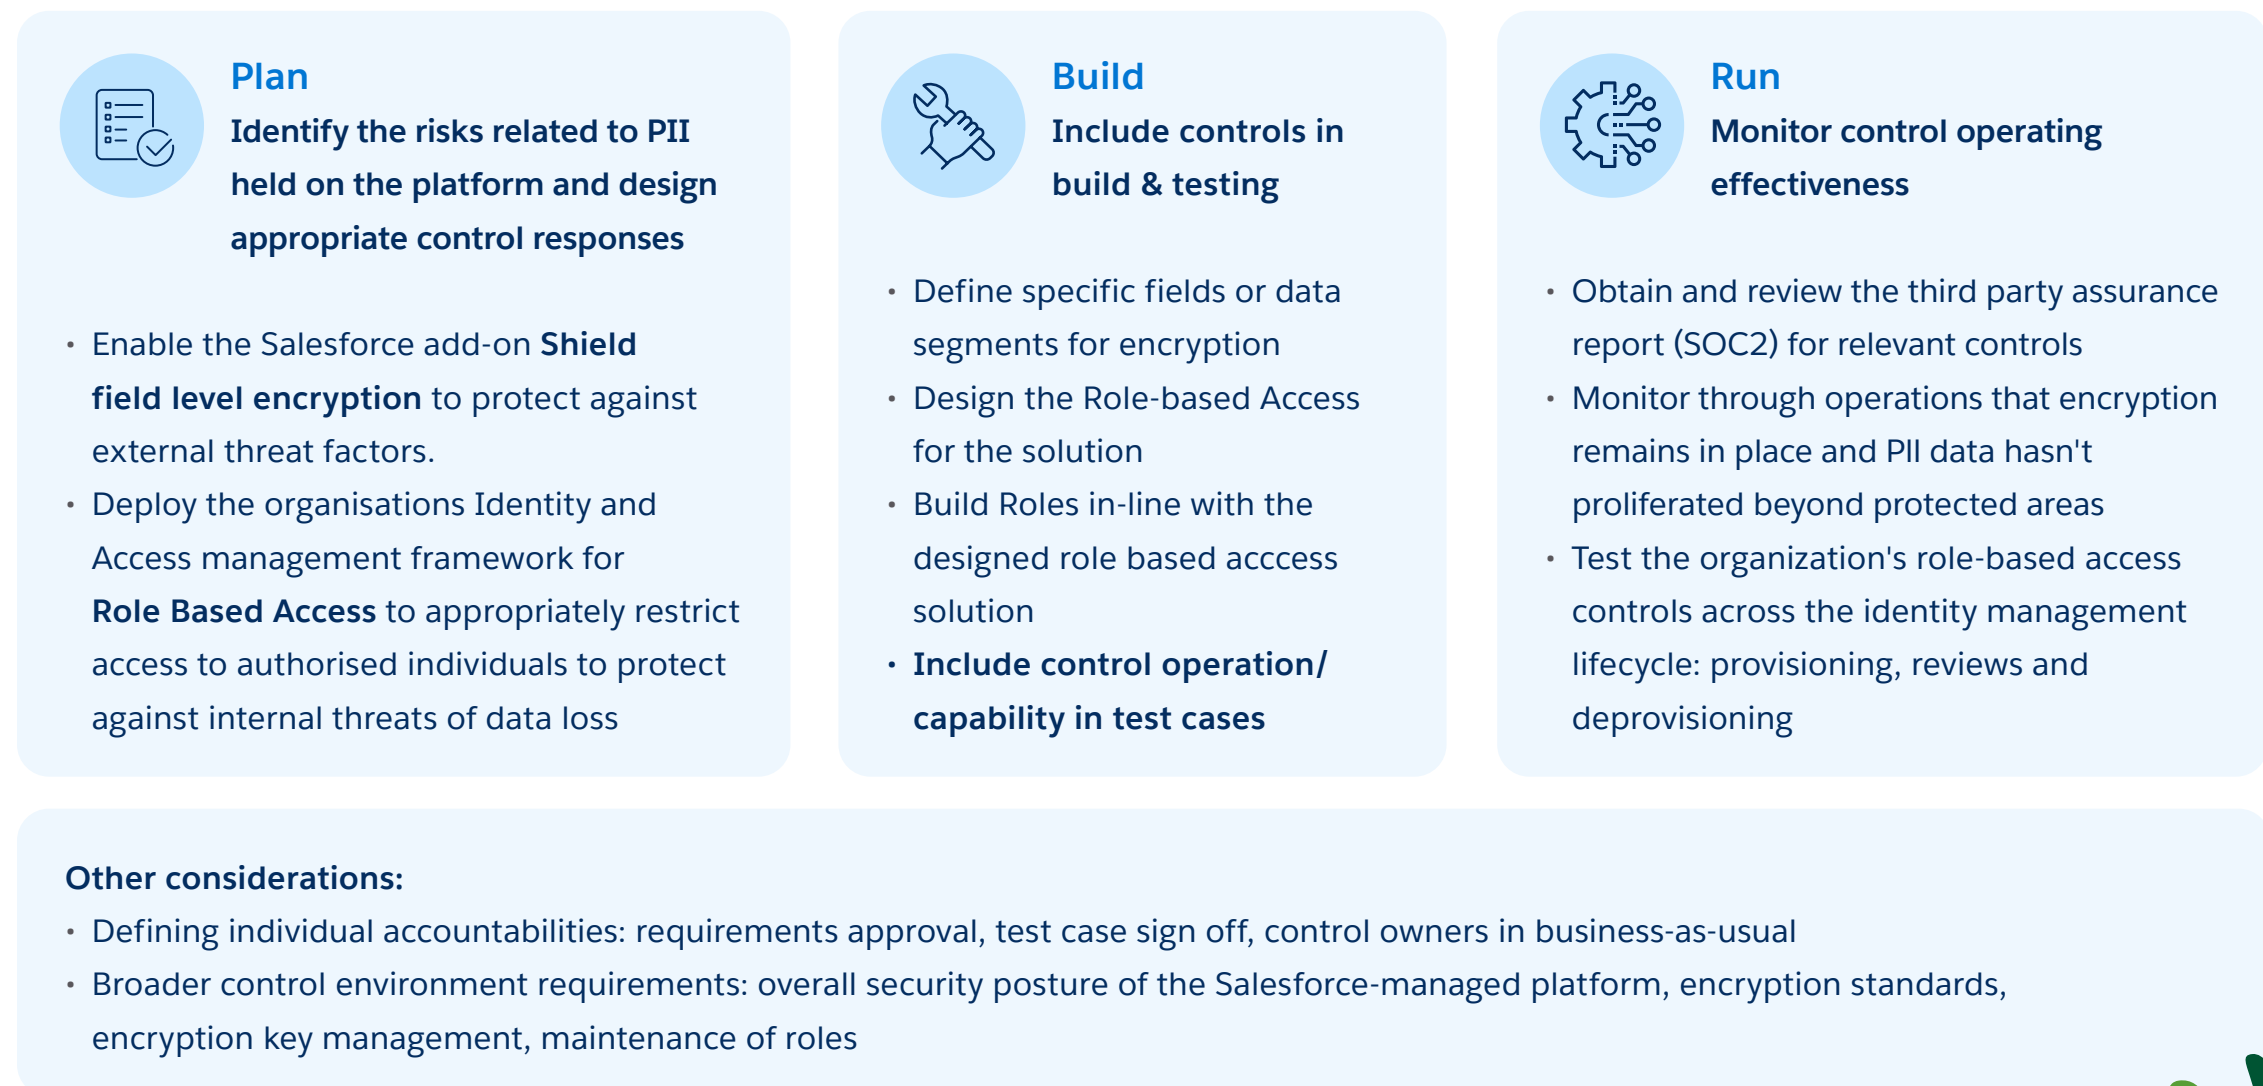
Fig 1: Illustrative shared responsibility model in a Salesforce environment



Having an early understanding of the shared responsibility model ensures risk management and resilience frameworks are embedded in an organisation's operating model when using a SaaS partner, such as Salesforce. Salesforce offers customers options that should be decided on early in the product lifecycle, to enhance security and resilience in line with the enterprise's

risk appetite. The following example illustrates how an organisation enacts the 'shared responsibility model' through the lifecycle, from pre-go-live to release and run. This is just one example risk and example responses - an organisation's risk management framework and team would shape the appropriate response to the organisation's unique circumstances.

Fig 2: Risk: Personally Identifiable Information (PII) breach




As noted in the 'run' state on the previous page, once the system is live, the relevant controls (under the shared responsibility model between Salesforce and its customers), should be regularly monitored as part of the enterprise's risk management framework. Monitoring the ongoing risk posture is often achieved through 'control effectiveness' assessments. In a shared responsibility model, this is accomplished through various approaches and sources. For example, controls operated by a user organisation can be tested internally, while controls fully managed by Salesforce require an upfront understanding of how the control design and effectiveness will be evidenced (refer to Appendix 1 for examples).

When using assurance reports provided by a third party, a unique type of control comes into play: the complementary user entity controls (CUECs). These are controls that a third party defines for their customers to operate, and which complement the vendor's controls to achieve specific control objectives. CUECs must be carefully considered when defining the risk and control framework under the shared responsibility model.

For regulated organisations implementing Salesforce, it is important to consider the relevant regulatory and legal obligations regarding hosting arrangements and to seek independent risk and legal advice before transitioning to the new platform.





Navigating compliance in the implementation of Salesforce

Risk owners who are accountable for regulatory compliance and the overall control environment need to carefully navigate the shared responsibility model and be able to engage with regulators in context of their risk environment. The compliance requirements associated with Salesforce implementations can vary widely depending on the business use cases supported by Salesforce, the data held within the platforms, the specific regulations that apply to the organisation and the organisation's own internal risk and supplier governance processes. Embedding risk subject matter experts (SMEs) at the initiation of the Salesforce program will help you plan for the key activities that need to occur during and after the implementation, and more confidently navigate complex compliance requirements.

Below is an example of key stages in an outsourcing program, and highlights risk and compliance activities at each stage.

Please note that while this paper aims to help customers navigate relevant regulatory requirements, Salesforce advises customers and prospective customers to consult with their own risk team and legal counsel to familiarise themselves with the requirements that govern their specific situations before making a decision to outsource a business activity or work with a third-party vendor.




Fig 3: Example: outsourcing lifecycle

The following example outlines an organisation's SaaS onboarding process from a risk perspective. It includes the key activities that an organisation may need to execute over the lifecycle of the program, based on their regulatory and compliance obligations.



Where required, engage with senior executives or Board at critical milestones for information and approval



AI advancements and the regulatory impact

In 2025, any growth story, regulatory agenda or technology adoption narrative is incomplete without a mention of the disruptive potential of AI. Compared by many to the industrial revolution, AI is expected to impact all facets of an organisation, from operations and customer experience through to strategic outlook. The Australian AI market is growing significantly and is projected to be worth US\$315bn by 2028, according to the Commonwealth Scientific and Industrial Research Organisation (CSIRO).²

Salesforce began using AI in its customer relationship management (CRM) software with the launch of Einstein, in 2014. Today, the natively built predictive AI tools are used to generate two trillion predictions every week, boosting productivity and unlocking more potential than ever. Salesforce continued its investment in AI research and development with the second wave of AI, bringing generative AI to customers with innovations like the CodeGen large language model (LLM).

With the launch of Agentforce in 2024, Salesforce pioneered the third wave of AI in CRM with autonomous agents. Agentforce is a step change in Salesforce's AI technology, including co-pilots. As the name suggests, autonomous agents are autonomous—meaning they don't require conversational prompts to take action. They can anticipate, plan, and reason with minimal, or increasingly no, human help. That means customers can create agents to automate entire workflows or processes, make decisions and adapt to new information, within risk management guardrails and without active human instruction.

2. [CSIRO List of Critical Technologies in the National Interest: AI technologies](#)



Expected regulatory scrutiny

While time and experimentation with AI will help answer many questions, concerns around increased regulation will impact the speed of adoption for many organisations, especially for those that manage customer data and operate in highly regulated markets, such as financial services. Regulatory bodies across many countries have proposed and implemented a plethora of different regulations, and providers of AI technology are also experimenting with self-regulation (e.g. the Frontier Model Forum) to ensure AI is developed safely and responsibly. However, as we have seen with previous emerging technology, regulation will likely fall into one of three major categories:

Continuum of regulation

Laissez-faire:

Governments allow regulation of the technology to take its own course.

Co-regulation:

Governments exist in a dialogic relationship with companies to respond incrementally to new developments and discovered harms arising from the technology.

Command-and-control:

Traditional forms of regulation, where the government enacts laws or binding legal frameworks.

In Australia, a co-regulation approach would likely mean a gradual increase in applicable safeguards and policies that ensure the benefits of AI are leveraged while concerns and future existential risks are mitigated. So far, this includes the voluntary AI Safety Standards set out by the Australian Government to help organisations think about their use of AI in a responsible and trustworthy manner. The standards, released in September 2024, have been followed by a consultation process to gain feedback from organisations, individuals and government agencies on risks including in relation to privacy and consumer law. Appendix 2 of this discussion paper includes a map of Salesforce capabilities and practices against the Voluntary AI Safety Standard for Australia.

It is critical for financial services organisations to engage with regulators and stay informed about upcoming changes. It can also be helpful to consult directly with Salesforce and your advisers, including PwC, on current guidelines, regulations and platform capabilities, should you wish to assess your ability to meet the regulatory requirements.





What's next?

In the highly regulated financial services industry, organisations must be proactive in safely navigating rapid market and regulatory changes. This requires both awareness of, and expertise in, the technological disruptions relevant to the industry. It also requires foresight into evolving risks and regulatory changes to inform timely decision-making.

The right approach for each company will depend on its strategy, operating model, industry context and competitive landscape. According to [PwC Australia's 28th Annual Review](#):

- 1 Nimble resource allocation remains a critical area for CEO attention
- 2 There is value in looking beyond a company's walls to embrace business ecosystems and explore joint ventures or alliances.

Aligning with trusted technology partners, underpinned with a shared accountability model, may mean the difference between being an industry leader or laggard.



How can Salesforce help?

Salesforce's advanced technologies, including AI and machine learning, can help you stay ahead of the technology adoption curve and foster a culture of innovation in your organisation. Highly configurable and extensible Salesforce solutions allow you to adapt the technology to your specific environment and gain operational insights, thereby improving efficiencies and maintaining a robust and sustainable control environment.

Salesforce has provided an example in Appendix 1 of some practices in selected areas of technology governance and controls and associated guidance available. These practices help you understand Salesforce's control environment, aiding in outsourcing decisions and defining the shared responsibility model.

[Speak to your SF contact for more clarification and how this applies to your org](#)



How can PwC help?

PwC combines unparalleled expertise in regulatory compliance and risk management with deep Salesforce capabilities, acting as a strategic adviser to help you interpret and respond to complex regulations while driving digital transformation. With deep industry knowledge and extensive experience, PwC can guide your organisation through multi-faceted stakeholder engagements, supporting transparent and effective communication with regulators, managing risk and compliance, building resilient solutions, managing project delivery and driving sustainable growth in an ever-evolving economic and regulatory landscape. Some of PwC's key services include:

- ✓ **CRM Strategy & Advisory:** we help align Salesforce capabilities to your business goals, developing CRM strategies that elevate customer engagement and drive sales.
- ✓ **Regulatory strategy and roadmap:** we help you design a plan to navigate the shared responsibility model and engage with your regulators
- ✓ **Risk assessments:** we embed a dedicated Salesforce risk and controls workstream with SMEs to support you with risk assessments, and the design and testing of appropriate controls
- ✓ **Implementation & Delivery:** Our teams enable seamless Salesforce implementation and integration, getting the best of the platform for your unique use cases, delivering value from day one.
- ✓ **Independent assurance** over controls design and documentation.

[Speak to your PwC contact Noel Williams, Banking and Capital Markets Leader, PwC Australia for more clarification and how this applies to your org](#)

Appendix 1: Mapping of example Salesforce capabilities and practices of representative FS regulatory requirements

This appendix provides examples mapping Salesforce's internal risk mitigation and control capabilities within representative categories of financial services regulatory requirements; however, it is not a comprehensive or prescriptive listing of requirements and practices. These are examples only, and to meet the requirements of any of the areas outlined below, it would be necessary to have additional controls across the shared responsibility model that are operated by Salesforce and the user organisation.

Organisations must evaluate their own policies, specific use cases on Salesforce, and regulatory requirements to define their shared responsibility model and assess requirements and which practices are needed to address requirements. The examples provided illustrate the types of information a Salesforce representative can offer to evidence the activities Salesforce undertakes in a typical shared responsibility model.

Supporting information from third parties, including Salesforce, comes in three main categories, which help determine the level of reliance a user organisation may place on this evidence. As such, organisations may need to seek additional information or carry out their own procedures to ensure that their compliance requirements are met. Organisations should consult with their risk team and refer to their operational risk management guidance for whether the Salesforce document reference provides sufficient and appropriate evidence to support controls defined within their shared responsibility model and commensurate with their risk assessment:



Third Party Assurance (T)

Independent auditor opinion under Auditing and Assurance Standards Board standards on the effectiveness of the description, design and (in some cases) operation of controls, e.g. SOC 2 or ASAE 3150 reports. User organisations must carefully review these reports for applicability in terms of the time period covered, scope of services and technology assets included, as well as undertaking a detailed review of the relevant controls and any issues noted.

E.g. SOC 1, SOC 2

Certifications (C)

Independent attestations for alignment with recognised international standards, e.g. PCI / DSS, ISO 27001. The level of control assessment varies by report and therefore user organisations must consider whether these reports provide sufficient evidence as per their own risk mitigation requirements. For example, some certifications represent a point in time view of the design of risk mitigation frameworks, rather than testing the operating effectiveness over a period of time.

E.g. ISO 270001, PCI DSS

Guidance Documents (G)

Documents that outline a service provider's internal policies, procedures and practices, offering insights into their operational standards and risk management approaches. This type of evidence supports understanding, but will not be suitable for cases where control design or operating effectiveness is required to be determined.

E.g. Salesforce best practices/ policies

In addition to the examples of Salesforce practices below, user organisations must be aware of the CUECs (complementary user entity controls) required to complete an effective control framework, as well as the controls the organisation operates wholly independent of their service provider. Further, any custom development or end-user operated functions are controlled by the user organisation. An organisation should consult with their risk management teams, and carefully read any documentation provided by Salesforce, to determine their own controls and CUECs that apply across the shared responsibility model.

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Information Security				
1	Clear roles and responsibilities for information security must be defined and assigned within the organisation	<ul style="list-style-type: none"> • CPS 234 • CPS 231 	<p>Salesforce has established a management framework to initiate and control the implementation and operation of information security within the organisation. Appropriate security roles and responsibilities have been documented and allocated within the security organisation under the Chief Trust Officer.</p> <p>Salesforce Management provides direction and support on a semi-annual basis for the information security program and functions through providing resources, delegating responsibilities, supporting initiatives, driving remediation activities, and holding individuals accountable for assigned roles.</p>	<ul style="list-style-type: none"> • (T) SOC 2 Report - Corporate Services document available from the Security & Compliance portal • (C) ISO 27001
2	Information assets, including those managed by related parties and third parties, should be classified according to criticality and sensitivity.	<ul style="list-style-type: none"> • CPS 234 	<p>Customers themselves control what data is submitted to the Services. Salesforce has an established policy for Data Classification. All information that has been electronically submitted by customers to the Services is categorised as 'Mission Critical' which provides the highest level of security.</p> <p>The Salesforce policy on data classification includes the following points:</p> <ul style="list-style-type: none"> • Salesforce assigns its business information and data according to the following data classification levels: Public, Internal, Confidential, Restricted, and Mission-Critical. • All Salesforce data must be classified and protected based on its classification. • When data of different classifications are mixed, the most restrictive classification applies. • Salesforce handles assets in accordance with the information classification scheme: • Data classification is applied for all assets and systems associated with Salesforce. <p>Additionally Assets must be labelled based on the data classification standard.</p> <p>Customer responsibility</p> <p>Salesforce Data Classification metadata fields allow customers to record the data owner, field usage, data sensitivity, and compliance categorisation for any standard or custom object field. They can also access data classification metadata in the Salesforce API and Apex.</p>	<ul style="list-style-type: none"> • (C) ISO 27001 certification

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Information Security				
3	Existing and emerging information security vulnerabilities and threats caused by insecure configuration of information assets are identified, assessed, and remediated in a timely manner.	• CPG 234	<p>Salesforce has a policy in place in which periodic vulnerability scans are performed on all Salesforce information systems and hosted applications. Frequency and comprehensiveness of scans is defined by security categorisation of the system, data sensitivity and/or specific regulatory requirements. Many of these scans are performed on at least a monthly basis across Salesforce products. Automated mechanisms are employed to compare the results of vulnerability scans over time to determine trends in information system vulnerabilities.</p> <p>Vulnerability scan reports and results from security control assessments are analysed and when new vulnerabilities potentially affect the system/application; they are identified and reported.</p> <p>Vulnerability scanning tools and techniques that facilitate interoperability among tools and automate parts of the vulnerability management process shall be deployed by using standards for:</p> <ul style="list-style-type: none"> • Enumerating platforms, software flaws, and improper configurations; • Formatting checklists and test procedures; and • Measuring vulnerability impact. <p>Identified vulnerabilities are assigned on priority basis with an associated internal service level agreement for remediation based upon risk. Salesforce management reviews vulnerability and patching status on a bi-weekly basis. Patches are deployed for known vulnerabilities at least monthly, or as needed based on the criticality.</p>	<ul style="list-style-type: none"> • (C) Salesforce Services PCI DSS AoC • (T) Salesforce Services SOC 2 report • (C) ISO 27001 Certification • (G) Vulnerability Management and Response Plan Summary • (G) Vulnerability/Penetration Report Summary - Salesforce Services • (G) Salesforce Vulnerability Management Program Overview

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Incident Management				
1	Customers should be notified as soon as possible and, no later than the contracted timeframes, after becoming aware of an information security Incident or an operational risk incident.	<ul style="list-style-type: none"> • CPS 234 • CPS 230 	Salesforce has a formal Incident Management Process that guides the Salesforce Computer Security Incident Response team (CSIRT) in investigation, management, communication, and resolution activities. Salesforce will notify the Customer without undue delay after becoming aware of the accidental or unlawful destruction, loss, alteration, unauthorised disclosure of, or access to Customer Data, including Personal Data, transmitted, stored or otherwise Processed by Salesforce or its Sub-processors of which Salesforce becomes aware. Notification may include phone contact by Salesforce Support, email to the customer's administrator and Security Contact (if submitted by customer), and public posting on trust.salesforce.com. Regular updates are provided to engaged parties until issue resolution. Incident tracking and resolution is documented and managed within an internal ticketing system. In the event that the CSIRT requires additional assistance in responding to a complex, high severity incident, Salesforce can also exercise retainers that are in place with multiple external incident response consulting companies.	<ul style="list-style-type: none"> • (G) Manage Security Contacts for Your Salesforce Organization • Salesforce Data Processing Addendum (Customer Data Incident Management and Notification section)

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
IT Resilience				
1	A comprehensive assessment of the operational risk profile must be maintained which includes identifying and documenting the processes needed to deliver critical operations including technology, information, and key data and controls	• CPS 230	<p>Salesforce conducts a comprehensive assessment of its operational risk profile through a structured and detailed approach. This process involves identifying and documenting the critical operations necessary to deliver essential services, including technology, information, key data, and controls.</p> <p>Identification and Documentation of Processes</p> <p>1. Technology and Information Systems:</p> <ul style="list-style-type: none"> • Salesforce identifies all technology systems and information repositories that are critical to its operations. This includes databases, servers, and cloud infrastructure. • Each system is documented with details on its functionality, dependencies, and the data it handles. <p>2. Key Data and Controls:</p> <ul style="list-style-type: none"> • Critical data elements are identified, including customer data, transaction records, and operational metrics. • Controls are established to ensure data integrity, confidentiality, and availability. These controls include access management, encryption, and regular audits. <p>3. Risk Assessment Procedures:</p> <ul style="list-style-type: none"> • Regular risk assessments are conducted to evaluate potential vulnerabilities and threats to the operational environment. • These assessments involve both automated tools and manual reviews to identify risks such as data breaches, system failures, and compliance issues. <p>4. Operational Continuity:</p> <ul style="list-style-type: none"> • Processes are documented to enable operational continuity in the event of disruptions. This includes disaster recovery plans, backup procedures, and incident response protocols. • Regular drills and simulations are conducted to test the effectiveness of these plans and ensure readiness. <p>5. Compliance and Governance:</p> <ul style="list-style-type: none"> • Salesforce adheres to a comprehensive governance framework that includes policies and procedures for risk management. • Compliance with applicable regulations and industry standards is supported through continuous monitoring and regular updates to the governance framework. 	<ul style="list-style-type: none"> • (T) SOC 2 Report • (C) ISO 27001 Certification

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
IT Resilience				
2	Sufficient isolation of backups from the production environment should be maintained so that a compromise of the production environment does not compromise backups.	<ul style="list-style-type: none"> • CPG 234 	<p>Salesforce Hyperforce is configured and deployed in a highly available manner. The systems are designed to recover from failure in a minimally disruptive manner. All Customer Data submitted to the Covered Services is written to persistent storage across multiple availability zones. Data backups are encrypted using at least AES-256 via FIPS 140-2 validated encryption.</p> <p>Customer responsibility</p> <p>Salesforce conducts backup of customer data and systems to support Business Continuity and recovery of services in the event of a disaster. In addition to this we also provide our customers with tools and capability to back up their Salesforce data through a number of methods. These options include:</p> <ul style="list-style-type: none"> • Data Export Service: Schedule or manually generate a backup of your org's data into comma-separated values (CSV) files. • Export a Report: Manually generate a backup of your org's data from Salesforce Reports. • Data Loader: Manually generate a backup of up to 5,000,000 records into a CSV file. • Bulk API: Create a custom backup to quickly extract and restore large data volumes. • AppExchange: Choose from several third-party apps 	<ul style="list-style-type: none"> • (G) Security, Privacy, and Architecture document for Salesforce Core Services
Governance				
1	Entitles must conduct a thorough risk assessment and due diligence before entering an outsourcing arrangement. This includes evaluating the third party's ability to meet service levels and compliance requirements	<ul style="list-style-type: none"> • CPS 230 • CPS 231 	<p>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. This includes items such as background screening, training and breach of policy and enforcement. Prospective vendors supporting the production environment for the Salesforce 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 go-live. Deficiencies noted in the review are remediated and/or compensating control(s) identified to address key risks, prior to go-live with potential access to Customer Data. Contracts are in place with all third parties that support the production environment, and these third-party vendors are measured against SLAs and terms within their contract, and Salesforce policies and procedures and information and physical security practices on at least an annual basis.</p>	<ul style="list-style-type: none"> • (G) Salesforce Third Party Risk Management Overview

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Governance				
2	Ongoing operational and strategic oversight mechanisms exist that facilitate assessment of performance against agreed service levels, assessment of the ongoing viability of the cloud provider and the service, timely notification of key changes and a timely response to issues and emerging risks.	• CPS 231	<p>The Salesforce Services are designed with the concept of continuous improvement and Trust (e.g. Availability, Performance and Security) in the infrastructure. Salesforce uses commercially reasonable efforts to make its on-demand services available to its customers 24/7, except for planned downtime, for which Salesforce gives customers prior notice.</p> <p>Salesforce does not typically offer Service Level Agreements as part of the base service offering, but can sometimes be negotiated when contracting in certain circumstances. Our approach is to offer a service with availability and timely resolution of problems. Salesforce has a Site Reliability Operations team that monitors service levels and we provide reporting to customers on these, where required.</p>	<ul style="list-style-type: none"> • (G) Salesforce updates key information on a daily basis on the publicly facing website trust.salesforce.com
Privacy				
1	Organisations must have a clearly expressed and up-to-date privacy policy that explains how the organisation collects, uses, handles, and discloses personal information, and which explains how individuals may raise a complaint or request access or correction of their information	• APP1	<p>Salesforce provides Software-as-a-Service (SAAS) to its customers and, in relation to the personal information that is processed during the provision of those services, Salesforce provides transparency about the manner of that processing via its Data Processing Addendum (DPA, available online here) which is incorporated into our standard Main Services Agreement. Our DPA in turn incorporates comprehensive information about the sub-processors that are used to provide the Services, via Infrastructure & Subprocessor documentation. This high level of transparency assists our customers in meeting their obligation under APP1 to be transparent with their end users about how their personal information is processed.</p> <p>In relation to the personal information Salesforce processes about individual employees of its customers and prospective customers, Salesforce provides information about how we collect, use, handle, and disclose personal information, and how individuals may exercise their rights in relation to that personal information, via our Privacy Statement.</p>	<ul style="list-style-type: none"> • (G) Data Processing Addendum • (G) Infrastructure & Subprocessor documentation • (G) Privacy Statement

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Privacy				
2	Organisation must take reasonable steps to protect the personal, sensitive and health information it stores from misuse, interference, and loss, and from unauthorised access, modification, or disclosure	• APP11	<p>Salesforce has implemented technical as well as administrative controls to protect personal information processed by the online services on behalf of customers. We also provide our customers with the ability to review and enhance the security of the data stored and processed on the platform using a number of out of the box and commercially available add-on capabilities. We also provide contractual commitment that we will not materially decrease the overall security of our services during an order term as part of the MSA that we enter into with our customers.</p> <p>Salesforce undergoes numerous security and privacy audits and assessments to ensure we have the right controls in place to protect the confidentiality, integrity, availability, security and privacy of customer data.</p> <p>Salesforce does not need human access to our customers' data unless this access is required under narrow circumstances, such as part of a customer support request.</p> <p>Best practice in securing data in a SAAS context is achieved by both the service provider (i.e. Salesforce) and its customers both playing their respective roles in maintaining a high level of security. This is known within the industry as the “shared responsibility model”. Under this model, the customers can fulfil their responsibility by implementing several key security measures:</p> <ol style="list-style-type: none"> 1. User Authentication and Authorisation 2. Data Encryption 3. Access Controls 4. Monitoring and Auditing 5. Advanced Security Features 6. Backup and Recovery 	<ul style="list-style-type: none"> • (G) Salesforce Data Processing Addendum • (G) Salesforce Main Services Agreement • (T, C, G) Salesforce Services Compliance documentation

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Cybersecurity				
1	<p>Organisation must establish and maintain a process or system – as far as it is practical to do so:</p> <ul style="list-style-type: none"> • minimise or eliminate any material risk of a cyber and information security hazard occurring; and • mitigate the relevant impact of a cyber and information security hazard on the CI asset. 	<ul style="list-style-type: none"> • Security of Critical Infrastructure – Clause 8 	<p>Salesforce addresses the minimisation or elimination of material risks associated with cyber and information security hazards through a comprehensive set of measures. These measures apply to risks arising on the Salesforce network and infrastructure. The risks arising via an organisation's own network must be managed by the user organisation as part of the 'shared responsibility model'.</p> <p>Salesforce-side mitigations include:</p> <ul style="list-style-type: none"> • Administrative, technical, and physical safeguards to protect the security, confidentiality, and integrity of personal data. • Network protection and authentication protocols to secure access. • Physical access controls at data centres to prevent unauthorised entry. • Compliance with information security policies and standards to ensure adherence to best practices in data protection. <p>As part of our ongoing risk assessment and treatment exercise, and in line with our Vulnerability Management standards and Incident Response plans, Salesforce regularly reviews all assets and associated threats. Suspected and confirmed cybersecurity incidents are investigated and addressed to ensure the protection of critical assets.</p>	<ul style="list-style-type: none"> • (T) SOC 2 report • (G) Salesforce Security (Incident) Response Plan
2	<p>Organisation must establish and maintain a process or system to comply with a framework listed below:</p> <ul style="list-style-type: none"> • Australian Standard AS ISO/IEC 27001:2015 • Essential Eight Maturity Model published by the Australian Signals Directorate • Framework for Improving Critical Infrastructure Cybersecurity published by NIST. • Cybersecurity Capability Maturity Model published by the Department of Energy of the United States of America • The 2020-21 AESCSF Framework Core published by Australian Energy Market Operator Limited (ACN 072 010 327) 	<ul style="list-style-type: none"> • Security of Critical Infrastructure – Clause 8 	<p>Salesforce maintains a formal company-wide information security management system (ISMS) that conforms to the requirements of ISO 27001 standard and NIST Cybersecurity Framework (CSF), including security policies, standards, and procedures. Formal policies and procedures are documented for operational areas including but not limited to: data centre operations, development, program management, production management, infrastructure engineering, release management, operations, hiring, and terminations. The Information Security Policy and Standards have been developed to segregate duties and enforce responsibilities based on job functionality.</p> <p>Salesforce Services are also assessed against Australian Government Information Security Registered Assessors Program (IRAP) requirements and against the Australian Cyber Security Center's Essential 8 requirements</p> <p>Salesforce implements the National Institute of Standards and Technology Cybersecurity Framework (NIST CSF) through a comprehensive set of security features and practices designed to protect customer data and provides the capabilities that may enable our customers to ensure compliance with applicable NIST requirements.</p>	<ul style="list-style-type: none"> • (C) ISO 27001 • (G) Salesforce Services on Hyperforce ACSC Essential Eight Maturity report

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
IT General Controls				
1	Maintain controls to manage changes to information assets, including changes to configuration with the aim of maintaining information security	• CPG 234	<p>Salesforce has a formal process for placing a system into production (including the hardware, software and appropriate configuration). This procedure includes a build checklist, server hardening checklist and pre-production testing. Baseline configurations for servers, network devices, and databases are consistent with industry-accepted CIS (Centre for Internet Security) system hardening guidelines that address known security vulnerabilities. Prior to go-live with new infrastructure, management approval is required and the decision is based upon the results of the pre-production testing.</p> <p>Salesforce has internally-developed guidelines for hardening servers, including:</p> <ul style="list-style-type: none"> • All new systems are scrubbed and installed from known baseline images prior to production deployment; • All unused services, ports, and packages are disabled; • A minimal set of user accounts is maintained and no shared password files are allowed; • The environment uses standardised configurations that ensure that the unneeded services are disabled on all machines. <p>The change management process supports a controlled framework as well as proper segregation of duties for the initiation, solution design, test, approval, implementation, and verification of changes. The Salesforce Services change management process is reviewed annually. Additionally, Changes to Salesforce products and data centres, including changes to applications, information systems, network topologies, configurations, and data centre facilities, are managed by a documented change control process.</p>	<ul style="list-style-type: none"> • (C) PCI DSS AoC • (T) Salesforce Services SOC 2 report
2	Maintain complete and accurate records of all privileged accounts	• CPG 234	<p>Salesforce maintains a record of privileged accounts through several mechanisms designed to ensure security and accountability. Here are some key methods:</p> <ul style="list-style-type: none"> • Each user in Salesforce is identified with a unique username and password. This ensures that every action can be traced back to a specific user. • Multi-Factor Authentication (MFA) is required for all logins to Salesforce products, adding an extra layer of security. <p>Technical support and operations staff, who have access to the production environment are granted access based on principles of least privilege and only to functions they need access to. Roles and permissions are defined and documented as part of the access control standards. Salesforce conducts quarterly access reviews to ensure access is granted on a need to know basis. Upon termination, privileged accounts are locked in Kerberos, connections are terminated and tokens are removed. This control is reviewed as part of our SOC 2 assurance audit.</p>	<ul style="list-style-type: none"> • (C) PCI DSS AoC • (T) Salesforce Services SOC 2 report • (C) ISO 27001:2013

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
IT General Controls				
3	Ensure access to information assets is only granted where a valid business need exists, and only for as long as access is required.	• CPG 234 Attachments A and C.	By default staff do not have access to sensitive environments. Access is only granted upon request and post necessary approvals. Access to critical systems/functions is granted on the principles of least privilege. These privileged users must authenticate to the Production Remote Access (PRA) gateway system. The PRA system hosts a secure environment in which the privileged users manage the production systems with the users receiving only a bitmap representation of a virtual screen hosted in the secure environment. Those with privileged access are required to authenticate to a secure server using 2 layers of two factor authentication. Salesforce has also implemented Just-In-Time access to critical systems to ensure access is granted for a predetermined period of time.	<ul style="list-style-type: none"> • (C) PCI DSS AoC • (T) Salesforce Services SOC 2 report • (C) ISO 27001:2013
4	Ensure the strength of identification and authentication is commensurate with the impact should an identity be falsified	• CPG 234 Attachment C.	<p>Salesforce requires all systems and applications to implement a log-in process in accordance with the defined Access Management and Identity Management Standards.</p> <p>As part of the Salesforce Information Security Standards; Salesforce employees are responsible for keeping their passwords secure which include but are not limited to:</p> <ul style="list-style-type: none"> • Not to share passwords or accounts with anyone, only to use individual passwords • Not to write down passwords, nor reuse their Salesforce passwords on non-Salesforce accounts. <p>Internal system privileged accounts on production systems and corporate IT endpoint systems are required to meet the following information security password parameters:</p> <ul style="list-style-type: none"> • Minimum of 16 characters for corporate IT endpoint systems and applications • Minimum of 12 characters for production environment individual user system and application accounts • Passwords must contain a minimum complexity of three out of four: uppercase, lowercase, numeric, symbols based upon available system functionality • Password maximum lifetime is restricted to 365 days for corporate IT endpoint systems and applications • Password maximum lifetime is restricted to 60 days for administrators and production systems • Passwords cannot be reused for at least 10 generations • Account lockout settings are enforced after 5 consecutive invalid login attempts and automatically lock the account after the number of unsuccessful attempts is exceeded 	<ul style="list-style-type: none"> • (T) SOC 2 report • (C) PCI DSS AoC • (C) ISO 27001:2013

Appendix 2: Mapping of Salesforce capabilities to Australian Voluntary AI Guidelines

The AI regulatory landscape is rapidly evolving, necessitating active engagement with relevant authorities to stay abreast of changes. While, at the time of publishing, Australia is yet to implement enforceable regulations, organisations must proactively navigate this dynamic environment to ensure compliance and mitigate risks effectively. The examples provided illustrate the types of information a Salesforce representative can offer to evidence the activities Salesforce undertakes in a typical shared responsibility model.

This appendix provides examples mapping Salesforce's internal risk mitigation and control capabilities within representative categories of voluntary AI Safety Standards; however, it is not a comprehensive or prescriptive listing of requirements and practices. These are examples only, and to meet the requirements of any of the areas outlined below, it would be necessary to have additional controls across the shared responsibility model that are operated by Salesforce and the user organisation.

In addition to the examples of Salesforce practices below, user organisations must be aware of the CUECs (complementary user entity controls) required to complete an effective control framework, as well as the controls the organisation operates wholly independent of their service provider. Further, any custom development or end-user operated functions are controlled by the user organisation. An organisation should consult with their risk management teams, and carefully read any documentation provided by Salesforce, to determine their own controls and CUECs that apply across the shared responsibility model.



S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Artificial Intelligence (AI)				
1	Establish, implement, and publish an accountability process including governance, internal capability, and a strategy for regulatory compliance for the use of AI within the organisation	Guardrail 1, Voluntary AI Safety Standard	<p>Salesforce ensures the establishment, implementation, and publication of an accountability process for AI use through a comprehensive governance framework, internal capabilities, and a strategic approach to regulatory compliance. Here's how Salesforce achieves this:</p> <p>Governance Framework</p> <p>Salesforce has developed a robust AI governance framework that includes policies, processes, and best practices to enable responsible and scalable AI use. This framework aligns with the company's vision and applies to the inventory of current AI solutions used in the business. Key components include:</p> <ul style="list-style-type: none"> • AI Governance Policies: These policies guide the responsible development and use of AI, help to address potential risks and provide mitigation strategies. • Ethical AI Principles: Salesforce adheres to principles such as accuracy, safety, empowerment, and sustainability to guide the development of trusted generative AI. • AI Acceptable Use Policy: This policy provides clarity for customers on how they can use our AI products, aligns with industry standards, and helps provide safe and trusted experiences with AI technologies. <p>Internal Capabilities</p> <p>Salesforce has built internal capabilities to support the ethical and effective use of AI:</p> <ul style="list-style-type: none"> • Bias Detection and Mitigation: Tools like bias detection in Einstein Discovery help identify and remove bias from data and models. • Model Cards: These provide transparency about model inputs, outputs, and ethical considerations, helping users understand and trust AI predictions. • Data Quality Scores: Indicators that help determine the readiness of data for accurate AI predictions, ensuring high-quality input data. <p>Strategy for Regulatory Compliance</p> <p>Salesforce's strategy for regulatory compliance involves several key elements:</p> <ul style="list-style-type: none"> • Secure Data Retrieval and Grounding: Grounding AI prompts only with data that the executing user has access to, respecting existing access controls and permissions. • Data Masking: Sensitive data is detected and masked before being sent to external models • Zero Data Retention Policy: Salesforce's third party LLM providers have committed to ensure that data sent to LLMs is not retained and is deleted after a response is generated. • Prompt Defence: System policies designed to limit hallucinations and decrease the likelihood of unintended or harmful outputs by the LLM. (continued on next page) 	<p>(G) For additional detailed guidance, please review the following Salesforce H&T articles:</p> <ul style="list-style-type: none"> • Ethical AI Model Building • Ethical AI in Einstein Discovery • Einstein Trust Layer • Einstein Generative AI & Trust • Einstein Trust Layer


S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Artificial Intelligence (AI)				
1 (continued)			Commitment to Trust At Salesforce, trust is the number one value. The company has agreements with large language model providers to ensure that generative AI capabilities can be used without private data being retained by these providers. Additionally, Salesforce's Einstein Trust Layer is a secure AI architecture designed to safeguard data privacy, enhance AI accuracy, and promote responsible use of AI across the Salesforce ecosystem.	
2	Organisations must establish and implement a risk management process that assesses the AI impact and risk based on how you use the AI system. This should begin with the full range of potential harms with information from a stakeholder impact assessment. Organisation must complete risk assessments on an ongoing basis to ensure the risk mitigations are effective	Guardrail 2, Voluntary AI Safety Standard	Salesforce implements a comprehensive risk management process to help assess the impact and risks associated with AI systems. This process begins with identifying the range of potential harms and involves a stakeholder impact assessment. Here's how Salesforce maintains this process: Identifying Risks and Establishing Guardrails 1. Risk Identification: • Data Leaks: Preventing customer data from being exposed or mishandled. • Regulatory Issues: Taking appropriate steps to comply with relevant laws and regulations. • Reputational Harm: Maintaining trust and avoiding actions that could damage the company's reputation. 2. Risk Mitigation Strategies: • Security Guardrails: • Secure Data Retrieval: Architecting so that prompts use only data that the end user is allowed to access. • Data Masking: Replacing sensitive data with placeholder data to prevent exposure. • Zero Data Retention Policies to be used with external LLM providers: Products are designed so that data isn't stored in the model after the response is generated. • Technical Guardrails: • Prompt Defence: Implementing protections against prompt injection attacks or jailbreaking. • Ethical Guardrails: • Toxicity and Bias Detection: Identifying and mitigating harmful language in prompts and responses. Stakeholder Involvement • Stakeholder Impact Assessment: Engaging with project stakeholders to identify the likelihood and impact of negative effects associated with AI project data. • Building a Risk Profile: Using tools like the Plan Your Trust Strategy unit in Trailhead to create a risk profile and mitigation strategy to share with key stakeholders. (continued on next page)	(G) Additional information regarding DPIA can be found on our website: GDPR and Data Protection Impact Assessments Data Protection Impact Assessments & Salesforce Services

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Artificial Intelligence (AI)				
2 (continued)			<p>Prioritising Risk Mitigation</p> <ul style="list-style-type: none"> • Critical Risk Focus: Prioritising the most critical risks first so that the most significant potential harms are addressed promptly. <p>Ethical AI Practices</p> <ul style="list-style-type: none"> • Ethical AI Model Building: Building AI models ethically, with a focus on removing bias and documenting predictive models. • Transparency and Trust: Providing transparency through model cards and designing AI solutions to be trustworthy and safe. <p>Continuous Improvement</p> <ul style="list-style-type: none"> • Feedback and Audits: Continuously monitoring and improving AI systems through feedback and regular audits to support ongoing compliance and risk mitigation. 	
3	Organisations should test AI systems and AI models before deployment, and then monitor for potential behaviour changes or unintended consequences. These tests should be performed in accordance with the clearly defined acceptance criteria that consider the risk and impact assessment for the relevant use case.	Guardrail 4, Voluntary AI Safety Standard	<p>AI models and applications need a consistent framework to evaluate them during the development time and throughout the model's usage in production. To address the challenges posed by the evolving nature of GenAI, Salesforce Offensive Security teams are committed to testing the core AI systems rigorously.</p> <p>According to OWASP machine learning models have several security risks. The new generation of models such as LLMs bring with them another set of risks as documented here by OWASP. Salesforce follows the secure SDLC approach for any Large Language Models (LLMs) developed by Salesforce.</p> <p>Salesforce uses external LLM providers, such as OpenAI as part of a number of our Generative AI offerings. OpenAI uses rigorous testing methods</p>	<ul style="list-style-type: none"> • (G) Mitigating LLM Risks Across Salesforce's Gen AI

S.No	Principles /Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Artificial Intelligence (AI)				
4	The organisation should enable human control or intervention in an AI system to achieve meaningful human oversight across the life cycle.	Guardrail 5, Voluntary AI Safety Standard	<p>Salesforce ensures human control and intervention in AI systems to help achieve meaningful human oversight across the lifecycle through several key strategies:</p> <ul style="list-style-type: none"> • Human-Centric Design and Ethical AI • Salesforce emphasises ethical AI and human-centric design principles so that its AI systems are developed and deployed with human oversight and control at every stage, across predictive and Generative AI use cases. <p>Here are some key aspects:</p> <p>1. Transparency and Explainability:</p> <ul style="list-style-type: none"> • Model Cards: Salesforce uses model cards to document predictive models. These cards provide critical information about model inputs, outputs, and the conditions under which the models work best. This transparency helps users understand how the models make predictions and recommendations. • Bias Detection: Features like bias detection in Einstein Discovery alert users to potential biases in their data, allowing them to take corrective actions to ensure fair and unbiased outcomes. <p>2. Human-in-the-Loop:</p> <ul style="list-style-type: none"> • Review and Approval: AI-generated responses, especially in generative AI applications, are subject to human review and approval. This ensures that the final output aligns with the organisation's values, voice, and tone. This capability will depend on the specific AI use cases and customers should review these prior to implementing the AI technology for their use case. • Editable Responses: In many cases, users can edit AI-generated responses directly, providing an additional layer of human oversight. 	
5	Organisations must disclose when they use AI, its role and when they are generating content using AI. Disclosure can occur in many ways. It is up to the organisation to identify the most appropriate mechanism based on the use case, stakeholders and technology used.	Guardrail 6, Voluntary AI Safety Standard	<p>Yes, Salesforce does disclose when AI is being used. Transparency is a key principle for Salesforce, especially when it comes to the use of AI technologies. Customers are presented with a solution bill of material (BoM) post a successful discovery exercise. The Bill of Material provides a list of service capabilities that are being proposed to customers. These capabilities are backed by our Trust and Compliance documentation. Customers determine which AI (predictive or generative) capabilities they want to use as part of their implementation and what data these capabilities have access to. Due to the large number of implementation scenarios, based on customer's requirements, it is recommended they review the use cases of the AI system as some of the use cases may not allow the customer to restrict the use of data sources to specific objects/fields. Access controls are implemented and hence the access will continue to be restricted to the user using the service.</p>	<ul style="list-style-type: none"> • (G) Salesforce Trust and Compliance documentation

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Artificial Intelligence (AI)				
6	<p>Organisations must provide information to other organisations across the AI supply chain so they can:</p> <ul style="list-style-type: none"> • understand the components used including data, models, and systems. • understand how it was built. • understand and manage the risk of the use of the AI system. 	Guardrail 8, Voluntary AI Safety Standard	<p>Salesforce provides information to customers on the use of internal or external LLMs. Salesforce Trust and Compliance documentation, specifically the Infrastructure and sub-processor document as well as the Notices and Licences (NLI) documentation provide information necessary to inform our customers on use of third party sub-processors as well as use of software with our services.</p> <p>Salesforce has well defined tools and guidelines for use of AI technologies, especially as it pertains to processing customer data using those. We conduct a risk assessment of the systems prior to these technologies being made generally available. We also undergo audits and assessments, such as SOC 2, ISO 27001 as well as vulnerability assessment and penetration tests. Salesforce may rely on third party LLMs as part of our service. These third party LLM providers are selected based on the tenets of security, privacy, ethical considerations, compliance and availability of these services.</p> <p>Salesforce is deeply committed to the ethical and humane use of artificial intelligence (AI). Here are the core principles that guide the development and implementation of AI at Salesforce:</p> <ol style="list-style-type: none"> 1. Accuracy: AI model responses are backed up with explanations and sources whenever possible. Human oversight is recommended to check model responses before sharing them with end users for most use cases. 2. Safety: Detecting and mitigating bias, toxicity, and harmful responses from AI models through industry-leading detection and mitigation techniques, helping to ensure that AI outputs are safe and appropriate for all users. 3. Transparency: Models and features respect data provenance and are grounded in your data whenever possible. This helps users understand how AI models arrive at their predictions or recommendations. 4. Empowerment: Augmenting human capabilities, making people more efficient and purposeful in their work. The goal is to empower users with AI tools that enhance their productivity and decision-making processes. 5. Sustainability: Building right-sized AI models that prioritise accuracy while also reducing the carbon footprint. This principle underscores the company's commitment to sustainable and responsible AI development. (continued on next page) 	<ul style="list-style-type: none"> • (G) Salesforce Trust and Compliance documentation

S.No	Principles / Requirements	Regulatory Consideration	Relevant Salesforce Practices	Salesforce document reference
Artificial Intelligence (AI)				
6 (continued)			<p>Ethical AI Practices</p> <p>Salesforce has implemented several features and guidelines relating to the ethical use of AI:</p> <ul style="list-style-type: none"> • Bias Detection: Alerts users to potential biases in their data, allowing them to remove distorting effects on analysis and predictions. • Model Cards: Provide documentation for predictive models, communicating important usage information and ethical considerations. • Data Quality Scores: Indicate how ready your data is for AI to make accurate predictions, highlighting areas for improvement. • Bias Flags: Indicate when selected inputs can introduce bias into the model, helping users evaluate and address potential biases. <p>Trusted Generative AI</p> <p>Salesforce's Einstein generative AI solutions are designed with trust as the #1 value. The company has agreements with LLM providers, such as OpenAI, which provide that private data is not retained by these providers. The five principles guiding trusted generative AI are:</p> <ul style="list-style-type: none"> • Accuracy: Backing up model responses with explanations and sources. • Safety: Detecting and mitigating bias and harmful responses. • Transparency: Respecting data provenance and grounding models in user data. • Empowerment: Augmenting human capabilities. • Sustainability: Building accurate and environmentally responsible models. 	
7	Organisations must maintain records to show that they have adopted and are complying with the guardrails. This includes maintaining an AI inventory and consistent AI system documentation.	Guardrail 9, Voluntary AI Safety Standard	Salesforce has a formal Asset Management policy and standard and this includes AI systems being used. Salesforce policies on use of AI systems, risk assessments, configuration guidelines, testing methods and results are all formally documented and reviewed.	



The information provided in this report is strictly for the convenience of our customers and is for general informational purposes only. Publication by Salesforce does not constitute an endorsement. Salesforce does not warrant the accuracy or completeness of any information, text, graphics, links, or other items contained within this guide. Salesforce does not guarantee you will achieve any specific results if you follow any advice in the report. It may be advisable for you to consult with a professional such as a lawyer, accountant, architect, business advisor, or professional engineer to get specific advice that applies to your specific situation.