May 17, 2024

Digital Enforcement and Intelligence Branch
Competition Bureau
50 Victoria Street
Gatineau, Quebec
K1A 0C9

Salesforce, Inc. (“Salesforce”) welcomes the opportunity to provide comments responsive to the Competition Bureau’s (hereafter “the Bureau”) calls to gather stakeholder comments about the Bureau’s paper on artificial intelligence and competition. As further described below, Salesforce is committed to building trusted, transparent, open, interoperable, and accountable AI systems that prioritize fairness, accuracy, privacy, and positive societal impact. Salesforce believes that harnessing the power of AI in a trusted way will require governments, businesses, and civil society to work together to advance responsible, safe, risk-based, and globally interoperable AI policy and regulatory frameworks that promote competition and innovation.

Generative AI has the potential to confer significant benefits across business and society and its rapid growth is already disrupting numerous technology markets. Incumbent cloud infrastructure providers have played an important role in supporting this growth as suppliers of a critical input to LLM developers, including through partnerships. As explained below, these partnerships have the potential to promote growth and innovation, but can also serve to position the incumbent cloud infrastructure providers as gatekeepers who can use the partnerships to expand their market power and diminish competition for AI technologies and products that are bundled with or rely on them. Competition authorities should carefully scrutinize these partnerships to ensure that they do not enable markets relevant to AI to become locked up by or tip in favor of these incumbent providers.

Background on Salesforce

Salesforce is a global customer relationship management (CRM) software company, based in San Francisco, California. Salesforce empowers organizations of every size and industry to connect with their customers in a whole new way through the power of AI + data + CRM. Salesforce’s core product offering, marketed as the Salesforce Customer 360, is a cloud-based Software-as-a-Service (SaaS) CRM solution

415 Mission Street
San Francisco,
California 94105
Salesforce, Inc.
/salesforce
@salesforce

that enables companies to manage and improve their relationships with customers through, for example, identification and management of sales opportunities, coordination of marketing and customer engagement initiatives, and provision of customer support. Salesforce's CRM services can be purchased separately or in combination, are typically available in different "editions" to meet differing customer demands for capabilities and cost, and can be supplemented with add-on features that provide additional functionality. Salesforce offers four core CRM services – Sales Cloud, Service Cloud, Marketing Cloud, and Commerce Cloud – as well as CRM products customized for various industries.

Salesforce also offers several solutions that complement its CRM offerings, including:

- Data Cloud: a data platform that organizes and unifies data across Salesforce and other external data sources.
- MuleSoft: an integration solution that enables users to connect any system, application, data or device, whether in the cloud or on-premise, in a unified manner using application networks.
- Tableau: a visual analytics service that helps individuals and organizations better understand their data and solve problems.
- Slack: an enterprise communication and collaboration service.

Salesforce offers its software services globally and has been operating in Canada since 2004. Salesforce maintains corporate offices in two provinces, Ontario and British Columbia, and has more than 3,000 employees working in six different provinces.

Salesforce's Approach to Artificial Intelligence

Salesforce has been active in the research and development of AI technologies for almost a decade. Salesforce introduced its first AI functionalities into its products under the "Salesforce Einstein" brand in 2016. These early AI technologies added predictive AI functionalities to Salesforce's CRM products – for example, predictions about a customer's propensity to buy or the likelihood of a deal closing – and have been steadily augmented in the years since their initial introduction.

To further expand the capabilities of its CRM and related solutions, Salesforce has begun to introduce generative AI functionalities across its product portfolio. These new features enable customers to create and deploy assistive AI experiences natively in Salesforce's services. For example, Salesforce CRM users are now able to use generative AI to draft emails; generate summaries of sales, service, and marketing calls; surface relevant information during customer support interactions; and generate product descriptions for their ecommerce websites. Salesforce also recently introduced "Einstein Copilot," a conversational AI assistant that CRM users can converse with directly to solve issues faster and drive productivity. Salesforce has also introduced generative AI functionalities to Slack and Tableau.

Consistent with Salesforce's commitment to trusted AI, the foundation of Salesforce's generative AI features is Salesforce's "Einstein Trust Layer," which is a secure AI architecture natively built into the Salesforce platform. The Einstein Trust Layer is equipped with security guardrails that allow Salesforce users to benefit from generative AI without compromising their customer data. More specifically, the Einstein Trust Layer's zero-data retention policy means that no customer data inputted into Salesforce's AI features is used for Large Language Model (LLM) training by third-party model providers, or retained outside of a customer's Salesforce instance. Simply put, customer data is not Salesforce's product.

A core guiding value of Salesforce's generative AI strategy is to maximize choice for customers by being open, extensible, and model agnostic in its integrations with the AI ecosystem. Salesforce aims to enable customers to use the LLMs of their choice with Salesforce's services – whether that be a model provided by a third party, Salesforce, or the customer itself. To that end, Salesforce's "Model Builder" will allow customers to use predictive and generative AI models and services from numerous different providers – including Amazon Web Services ("AWS") via Amazon Bedrock and Amazon SageMaker, Anthropic, Microsoft's Azure OpenAI, Cohere, Databricks, Google Cloud's Vertex AI, and OpenAI – or to bring their own models.

Salesforce is also promoting the development of a vibrant LLM ecosystem through its Salesforce Ventures AI Fund, which has invested $500 million to date to support several AI companies, including Anthropic, Cohere, Hugging Face, Hearth.AI, Mistral, and You.com. Salesforce strongly believes that customers succeed when they have the ability to use the solutions – including LLMs and related AI products – that best meet the needs of their business. This is why Salesforce has built its entire CRM platform on open architecture and designed its products, including its AI solutions, to be interoperable and extensible.

Promoting AI Competition

Salesforce recognizes the transformative power of generative AI and the benefits it can bring to consumers and to businesses of all sizes. However, there is a significant risk that these benefits will not fully materialize or be enjoyed widely if a few dominant companies are able to lock up critical inputs either directly or through partnerships. Over the past few years, and most markedly over the past year, the development, evolution, and commercialisation of AI technologies – including generative AI technologies specifically, as well as key components required for the development and deployment of AI systems, such as cloud infrastructure, LLMs, data, and AI chips – have advanced significantly such that they have become critical inputs for a wide range of technology services. Salesforce anticipates that this trend will only accelerate going forward. Regulators and policymakers must endeavor to ensure that the AI space – and, importantly, the wide range of technology services that have started and will continue to rely on AI – remain open, and do not become locked up by or tip in favor of a small number of incumbents who control critical AI components.

Given the important role that AI technologies are so quickly assuming across technology services, Salesforce is concerned in particular that control over key AI components by incumbent cloud infrastructure providers could harm competition, both as it relates to AI technologies specifically and for numerous other technologies that rely on AI and cloud infrastructure. In Salesforce's view, certain technology markets are already impacted by the significant market power enjoyed by companies that have vertically consolidated and strategically distributed numerous integrated technology solutions – such as, for example, public cloud infrastructure, office productivity software, and communication/collaboration software. Salesforce is concerned that such a company's control over key AI components could further expand this market power not only as it relates to AI technologies, but also to services across the technology stack. As the Bureau acknowledged in its Discussion Paper, "the importance of vertical relationships in AI markets may provide an environment where particular exclusionary conduct could arise. Vertically integrated firms could both supply an important input for participation in AI markets, such as data, compute, and AI models, while also competing in the same downstream markets as the firms they supply to. This would be a concern for competition if these firms were to engage in behavior to exclude their downstream competitors from the market."

With respect to generative AI specifically, incumbent cloud infrastructure providers have played an important role in the development of and innovation in new AI technologies as they provide large amounts of the compute resources required by LLM and other AI developers. However, the position these incumbent cloud infrastructure companies occupy as suppliers of a critical input for AI technologies creates the risk that these companies could leverage their positions as gatekeepers in ways that expand their market power and diminish competition not just for AI technologies, but also for technologies that rely on and/or are bundled with AI products. Some of the ways that incumbent cloud infrastructure providers might leverage their powerful position that warrant careful scrutiny by regulators and policymakers include:

- ***Agreements that require an LLM to run exclusively on one cloud infrastructure provider's services:*** Exclusive agreements of this type have the potential to harm competition in several ways. At the customer level, they diminish choice by forcing customers who wish to access a particular LLM – either directly or in connection with providing services to their own customers that integrate that LLM – to use the infrastructure provider's services. These customers may have already made substantial investments in another cloud infrastructure provider or may be a competitor to that infrastructure provider, such that a requirement to use the infrastructure provider's services would be contrary to the customer's needs or interests and/or burdensome. In addition, in cases where an LLM provider cannot directly offer solutions that meet customer needs – such as where an enterprise customer has global data residency and compliance requirements – the customer may be forced to contract directly with the cloud infrastructure provider for the LLM bundled with the provider's infrastructure services in order to access the LLM at all. In such a case, the cloud infrastructure provider enjoys an effective monopoly over the provision of the LLM and can dictate the terms upon which that LLM can be accessed and used.

  At the LLM level, exclusivity can diminish innovation by preventing the LLM provider from working with multiple cloud infrastructure providers to develop new models and related solutions and skew the LLM provider's development efforts in favor of the cloud infrastructure provider's own products. At the infrastructure level, such agreements deprive other cloud infrastructure providers of the ability to host the relevant LLM and to make that LLM available to their users, which may diminish competition between infrastructure providers where the LLM occupies a significant position in the market. Finally, exclusive agreements of this type can drive significant additional demand to the cloud infrastructure provider and position the cloud infrastructure provider to leverage this demand to tie its other products, including AI products, to the LLM offering. This may raise competitive concerns depending on the characteristics of the bundled offering and availability of the LLM outside of the bundle.

- ***Agreements between a cloud infrastructure provider and an LLM developer that provide downstream competitors with inferior or delayed access to key LLMs***. Downstream competitors could face significant competitive disadvantages against a vertically integrated cloud infrastructure provider if an agreement between the cloud infrastructure provider and an LLM developer has the effect of forcing the downstream competitors to use less advanced or otherwise inferior models, or, in the alternative, to delay product releases until they are able to make full use of a key LLM.

- ***Exclusionary terms or practices by a cloud infrastructure provider that unfairly advantage the cloud infrastructure provider's own AI applications' use of its infrastructure.*** For example, downstream competitors could be harmed by terms or practices that limit their access to key infrastructure APIs used internally in connection with the cloud infrastructure provider's AI products. Such self-preferencing behavior could impact non-vertically integrated downstream competitors' ability to compete with the cloud infrastructure provider.

- ***Exclusionary terms or practices by a cloud infrastructure provider that limit access to key data or chips necessary to train LLMs.*** For example, downstream competitors could be harmed by agreements that provide a cloud infrastructure provider with exclusive access to key data necessary to train LLMs or that lock up the chips needed to develop and run them.

- ***Restrictions that limit interoperability and/or data portability.*** Restrictions that limit interoperability and/or data portability with a cloud infrastructure provider's services and/or applications could increase entry barriers and lock customers into the cloud infrastructure provider's AI ecosystem.

- ***Partnerships between a cloud infrastructure provider and LLM developer that enable access to rivals' competitively sensitive information.*** Absent robust safeguards, partnerships between a cloud infrastructure provider and an LLM developer could give the cloud infrastructure provider access to competitively sensitive information about competitors that utilize the developer's LLMs, including, for example, information regarding its competitors' relationships and use of the LLM provider, LLM prompts and responses, or product roadmaps (including planned innovations related to AI features).

Thank you for considering Salesforce's feedback. We look forward to continuing to engage with the Bureau on these important matters and are available should the Bureau require further information.

Sincerely,

Zach Carstensen
Director, Canadian Government Affairs
zcarstensen@salesforce.com
Salesforce, Inc.
+1 (206) 518-0985