Six Strategies to Prepare for Trusted Generative AI

Foreword: Looking to the future with generative AI

Generative AI has opened up a tech revolution, offering businesses large and small the opportunity to reshape how they operate, innovate, and engage with customers. AI can enhance efficiency by automating routine tasks, speeding up workflows, and reducing errors. What does that mean for employees? They'll have more time to focus on more strategic tasks that require the magic of human creativity.

At Salesforce, we see a future where customers use generative AI to tackle their biggest challenges. With AI, you're now able to offer service that delights customers by giving agents real-time suggestions, write personalized sales scripts in seconds, and cut through the marketing noise with tailored content that resonates with specific customer needs. Essentially, generative AI helps businesses use their data more effectively to address the most impactful issues to their growth.

The benefits of generative AI expand beyond just business. AI has the potential to help discover new drugs, monitor energy usage, and optimize designs for solar panels and wind turbines. That means big transformations in healthcare, energy, and sustainability that will benefit everyone.

Amid all this excitement, remember that the success of generative AI hinges on trust. Why? Generative AI's output can significantly affect people and businesses. Because of this potential impact, if you're going to adopt generative AI, you must identify and address its potential risks. Then, build a trusted foundation for this transformative technology to ensure it's used safely, reliably, and ethically.

Marla Hay
VP, Product Management, Salesforce

ⓘ This guide will dive into the following:

- Opportunities and impacts of generative AI
- Concerns surrounding generative AI
- Strategies for securing your data in preparation for generative AI

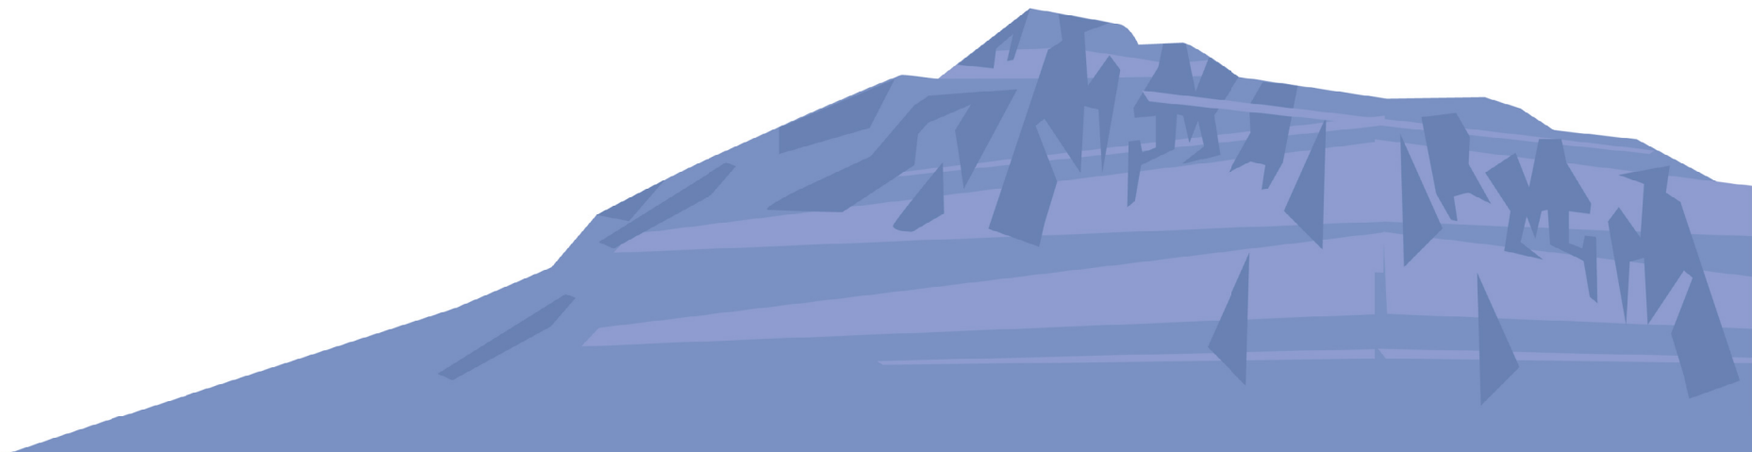

Contents

1

Explore Opportunities With Generative AI

Explore Opportunities With Generative AI

Already, experts are comparing the potential impact of generative AI to major technological milestones, such as the internet and mobile computing. Standing on the cusp of this new era, businesses are looking to understand AI's immediate opportunities and what it could mean for the future.

Build your understanding of generative AI

How is it different from past technologies?

The magic of large language models

Generative AI involves sophisticated deep learning algorithms that produce various content, including text, images, code, voice, and video. Its foundation rests on large language models (LLMs), and extensive neural networks surpassing traditional AI in complexity and capability.

Unlike traditional neural networks focused on prediction and pattern recognition, LLMs use self-supervised training to generate new content. The shift from simple classification and prediction to the ability to self-train enables a more dynamic understanding of data relationships. Generative AI is faster, more creative, and more adaptable than traditional neural networks.

From chatbots to copilots: The evolution of AI

Initially, chatbots operated on pre-set rules and scripts to handle customer queries. A customer following pre-determined prompts could arrive at the answer they needed as long as it was already in the decision tree, but they couldn't use natural language to ask questions and dig deeper. Traditional chatbots were helpful in simple queries, but it didn't take much to hit the ceiling of their capabilities.

With generative AI, we're now seeing the rise of AI copilots. These advanced AI assistants understand conversational context, learn from interactions, and provide responses that are not only accurate but also contextually relevant and personalized. This shift from basic chatbots to AI copilots marks a big leap toward more intuitive, intelligent, and empathetic interactions – where customers can get answers to complex queries and receive tailored recommendations.

The impact of generative AI on business: How do these opportunities translate into benefits?

Far from being a job eliminator, generative AI empowers teams by helping them use their time more efficiently and providing fast access to deep knowledge and insights. By automating routine tasks, generative AI allows employees to focus on more strategic and creative aspects of their work. Humans must remain at the helm to guide operations and decision-making, but with generative AI, organizations can boost operational efficiency and job satisfaction.

Generative AI is also transformative for customer experiences, where personalizing interactions is a must. For instance, a sales rep today can use generative AI to quickly perform in-depth research on a lead that would otherwise take hours, then instantly craft a personalized email speaking directly to that lead's preferences. Through its capacity to analyze extensive datasets for deeper insights, generative AI helps businesses achieve what they've been seeking – the ability to deliver tailored services that enhance the customer journey and cultivate loyalty.

> "AI is well-poised to absorb today's routine work patterns and free our time for more fulfilling, productive, and profitable work."
>
> – PARAM KAHLON, EVP AND GM, AUTOMATION & INTEGRATION, SALESFORCE

Where businesses stand today in their generative AI journey

Whenever a new tech trend emerges, the question arises: "Is this good for my business, or is it just a fad?"

While most organizations recognize that generative AI isn't a passing trend, some organizations are still hesitant to dip their toes in. They feel paralyzed by the security and accuracy concerns when using generative AI, and don't know how to adequately mitigate those risks.

While this hesitancy is understandable, it also puts organizations at a disadvantage. Businesses that embrace generative AI can produce game-changing results in sales, marketing, and customer service, while those who delay or avoid adopting it risk falling behind their competitors and not meet their customers' expectations.

Opportunities: The transformative power of generative AI

Generative AI can unlock numerous benefits:

Enhanced innovation and creativity

Generative AI can quickly analyze data, trends, and patterns from many perspectives. This allows teams to find insights faster and begin innovating sooner, giving them more time and mental real estate to focus on creative problem-solving.

Personalized customer experiences

AI can process and analyze huge amounts of data, allowing organizations to understand customers' preferences and behaviors more deeply. This enhanced understanding lets them create more targeted, effective customer engagement strategies and more timely, personalized interactions.

Streamlined operations

Automating routine tasks and processes with generative AI improves efficiency and gives employees the time to focus on higher-value activities. The result is increased productivity and significant cost savings.

Data-driven decision making

Generative AI provides actionable insights from large datasets, allowing for more informed decision-making. By analyzing complex data structures, organizational leaders can easily adapt to market changes and consumer trends.

Risks: The price of inaction

On the other end of the spectrum, organizations that are slow to adopt generative AI might experience challenges that make it harder to keep up in today's market:

Loss of competitive edge

Companies that don't adopt generative AI will struggle to keep up with their competitor's pace of innovation and efficiency. This lag can result in lost market share and shrinking relevance in their industry.

Inefficiency and higher costs

Without the automation and efficiency gains of generative AI, businesses will experience higher operating costs and lower productivity. This can ultimately strain resources and impact the bottom line.

Inability to meet customer expectations

Personalized customer experiences are becoming the expected norm. Using manual processes to understand and predict customer needs can lead to a disconnect from the market, driving down customer satisfaction and loyalty.

Missed insights and opportunities

Generative AI helps process and interpret complex data to spot new market trends while still relevant. Without these advanced analytical capabilities, organizations may miss insights and opportunities to help them grow.

2

Overcome Challenges and Concerns

Overcome Challenges and Concerns

Despite the excitement around generative AI, there are still concerns about using it without putting customers at risk. In a survey of over 500 senior IT leaders, nearly 70% considered generative AI a business priority and planned to fast-track its adoption.[1] However, almost all recognized the need for significant preparation to address integration, security, and accuracy challenges.

Lack of integration could lead to inaccurate results

A significant worry for many business and IT leaders is a lack of data integration. They know that if their data is incomplete or inaccurate, their generative AI tools might give skewed or unreliable results. After all, data-driven decisions and customer interactions are only as good as the data they're drawn from, which means comprehensive data integration is critical.

Security and privacy issues could damage trust

Companies handling sensitive customer information must be sure that their use of generative AI isn't putting that data at risk of exposure. When even one data breach or instance of data misuse could result in near irreparable loss of trust, IT leaders must ensure they have the strictest controls and security protocols before integrating generative AI into their processes.

Uncertainty about the quality of results

Lastly, leaders want to know that they can trust the results they're receiving from generative AI systems. Some are still skeptical of the reliability and quality of AI-generated information. Since generative AI might be used to influence major business decisions, the accuracy and dependability of its outputs are non-negotiable.

Integration, security, and accuracy are vital for building a trusted foundation for AI. Without them, the risks of using AI are far more likely to outweigh the benefits.

Essential needs for implementing AI

Accuracy

59% of IT professionals believe the output of generative AI could be inaccurate.[1]

Security

67% don't believe they're prepared to implement generative AI because of security concerns.[1]

Partnership

83% would collaborate with other businesses to help ensure ethical use of AI.[1]

[1] Salesforce survey shows IT interest in generative AI tempered with technical, ethical concerns - SiliconANGLE

Get to know the tenets of trusted AI

For companies ready to seize AI opportunities, building a foundation of trust is the first essential stop on their journey. At Salesforce, we've pinpointed seven tenets for trusted generative AI.

1. **Your data is not our product:** It is only used for its intended purpose, not as a commodity.
2. **Data residency and compliance:** Data is stored and managed where it legally and ethically belongs.
3. **Customer control and privacy:** Customers are given the tools and options to manage their data according to their unique needs.
4. **Enterprise scale:** Generative AI solutions must handle the growing volume and diversity of data at an enterprise scale.
5. **Built-in security:** Advanced security protocols are embedded into AI systems at every level to protect against threats and vulnerabilities.
6. **Ethical design and practice:** AI systems must be designed to be fair, transparent, and accountable to minimize the negative societal impacts of AI technologies.
7. **Human rights protections:** Customer data must be handled with care and respect to protect the dignity and rights of customers.

As you venture into using generative AI, these tenets of trusted AI can act as guiding stars. They'll help you stay grounded in ethical, secure, and responsible practices to innovate without compromising trust.

With clarity around the opportunities and challenges of generative AI, what must organizations do to ensure they build a foundation for trusted generative AI? Let's dive into the six critical strategies.

3

Six Strategies for Building a Foundation for Trusted Generative AI

Six Strategies for Building a Foundation for Trusted Generative AI

Today, many businesses find themselves at a crossroads. Some hesitate to move forward with generative AI because they're cautious of the unknown, while others are eager to dive right in. There are risks to both approaches which must be addressed before taking action. Whichever category your business falls under, you must take strategic steps to ensure trust is built into your generative AI journey.

These six strategies will help you start building a trusted foundation. Once that is in place, you can begin innovating with generative AI to improve experiences, free employees from manual processes, and uncover transformative insights – all with the peace of mind that you've taken the necessary steps to ensure customer trust.

Step 1: Audit your data

Large language models aren't designed to know what data needs extra protection. Conducting an audit of your data will help you understand where protections are needed so you can take the necessary precautions.

Auditing objectives checklist:

- ✓ Understand the data used in your prompts, including confidential or sensitive customer data.
- ✓ Determine which data your large language models can access.
- ✓ Clean and preprocess data by classifying or removing personally identifiable information (PII).

Data Detect

Identify sensitive information

Data Detect lets you quickly find sensitive information in your Salesforce instance that you might not even know exists. This helps you stay compliant and secure before implementing generative AI processes across your business.

Use Data Detect to define policies based on which types of information are considered "sensitive," scan the data in your organization, and identify where you might be storing items like Social Security Numbers and other PII, even if it's buried in long text fields.

- Find and classify sensitive data, such as credit card details and email addresses, and classify data based on your preferences
- Address flagged data and view sensitive data across your records
- Protect customer data by pairing it with other security features that help you comply with protection laws and best practices

Learn more about Data Detect

Step 2: Set up privacy protections

Build and maintain trust by ensuring customer privacy and safeguarding data throughout your AI processes. As the dependence on AI grows for comprehending data and making decisions, it's important to prioritize data protection, especially for PII. Eliminate and obfuscate data that is no longer useful or relevant to your processes.

Privacy protection objectives checklist:

- ✓ Set clear data usage policies that specify how customer data will be handled.
- ✓ Implement preference management to allow customers to opt in or out of having their data used.
- ✓ Eliminate and obfuscate data that is no longer useful or relevant to your processes.

Privacy Center

Set up data privacy protection

Privacy Center helps verify that your AI processes are consented for use in training and prompts. Create retention policies to manage the lifecycle of data used by and generated by AI, including call transcripts, chatbots, and cases automatically logged by AI. Then, remove records that aren't relevant anymore to limit liability in the event of an accidental breach or exposure.

- Set retention policies to anonymize, pseudonymize, and delete data according to preferences
- Honor customer requests, including Right To Be Forgotten (TBF) and Data Subject Access Requests (DSARs)
- Easily capture customer consent with customizable forms

Learn more about Privacy Center capabilities

Step 3: Secure access to your data

This step involves implementing controls to protect customer data from potential misuse. These controls will help ensure your AI integrations stay within the bounds of the data you want to use and protect it from unauthorized access.

Preparing for generative AI checklist:

- ✓ Determine which authorized users and teams should have access to customer data.
- ✓ Implement access controls through user permissions.
- ✓ Continually monitor and audit access on a need-to-know basis to reduce the risk of AI models and unauthorized personnel accessing sensitive data.

Security Center

Manage security controls from a centralized view

Security Center can help you centrally manage user permissions and org configurations for data used in and ingested from AI processes. You can easily grant access on a need-to-know basis for admin, IT teams, contractors, and others.

- Track your security posture and get insights into performance
- Simplify security management to reduce security risks
- Combine metrics into a consolidated view to eliminate blind spots

Learn more about Security Center

Step 4: Test processes

Testing in a sandbox environment serves two key purposes: testing AI processes and training employees on using AI safely and responsibly. It allows them to explore the technology's capabilities and build their confidence before applying them in the real world.

Testing objectives checklist:

- ✓ Assess and refine the performance and behavior of your generative AI models before deploying them in real-world scenarios.
- ✓ Identify and mitigate potential issues, such as biases, errors, or unintended consequences that may arise during a generative AI process.
- ✓ Give employees hands-on experience and training in using generative AI tools and systems.

Data Mask

Test AI systems without compromising sensitive data

Data Mask protects confidential data that is used to build and test AI applications. With flexible access, employees and contractors can test apps without opening your organization to risk. It also lets you easily eliminate or obfuscate data that shouldn't be included in testing.

- Scramble a field's content into unreadable results to make data anonymous
- Create pseudonyms by converting a field into a readable value not connected to the original value
- Delete fields completely to remove them from the data set

Learn more about Data Mask

Step 5: Back up data

The potential for mistakes grows as you feed more data into your AI systems. Backing up your data helps make sure those mistakes don't become irreversible. Once your data is backed up, it can be restored to its state before the AI system impacted it.

Data backup objectives checklist:

- ✓ Safeguard data against unintentional errors to prevent data loss or omissions.
- ✓ Enable quick recovery that restores data to its original form.
- ✓ Ensure data integrity to avoid legal and compliance issues.

Salesforce Backup

Keep data intact and compliant

A single data entry or deletion mishap could drastically affect an AI model's output. Salesforce Backup helps ensure these errors can be undone by restoring the data impacted to its pre-AI status.

- Set backup policies from an admin interface to deploy backups quickly
- Restore data in a few clicks to minimize the disruption to your operations
- Monitor your backup and restore activities in real time

Learn more about Salesforce Backup

Step 6: Monitor processes

Sometimes, AI systems will try to access data they don't have permission for. Having a robust monitoring mechanism helps detect and get alerts when an AI system tries to gain unauthorized access. Regular audits and reviews of AI integration processes and access logs can help identify deviations or potential security risks.

Monitoring objectives checklist:

- ✓ Set up security policies that detect threats to sensitive data.
- ✓ Send alerts and block attempts to prevent users from accessing or exporting data.
- ✓ Perform continuous reviews and audits of your AI processes to ensure sensitive data is protected.

Event Monitoring

Investigate AI activities to mitigate threats to your data

Event Monitoring gives you a granular view of user activities to catch unusual behavior that might threaten your data. This tool gives you insights into who's viewed your data, where it was accessed, when users change records, and more.

- See who, when, and where data is being accessed
- Get alerts on unusual user behavior and define policies to block access
- Continually analyze user behavior to improve app performance and experiences

Learn more about Event Monitoring

Find a reliable partner to help on your journey

Innovative technologies inevitably come with uncertainty, and generative AI is no different. To maximize the effectiveness of generative AI tools, businesses need a reliable partner who can help them establish trusted AI practices, navigating them through uncertainty and onward to confidence and innovation.

Shared trust is a data security model between Salesforce and our customers that helps them secure their data while preparing for generative AI. While our Einstein Trust Layer implements best-in-class security and privacy, our customers are responsible for ensuring that they're leveraging our technologies to prepare their data, protect it, and monitor it for future use.

Using this shared approach, our customers can confidently check off their top needs for successful generative AI integration:

- AI tools produce results that are accurate and relevant to the business
- Data used by AI models is protected, and permissions are enforced
- AI systems offer apparent oversight and transparency

Stay aware of regulatory insights

While setting up your trusted foundation for generative AI, staying current with the regulations that apply to your organization and the customers you serve is critical. As technology advances, the laws will change to help ensure necessary human rights aren't compromised.

Europe: The General Data Protection Regulation law

Formed in 2016, the GDPR protects individuals' privacy and data security in the European Union. It applies to any organization that processes people's data in the EU, no matter where the organization is located. The GDPR operates from seven core principles that set a framework for how data controllers and processors should handle personal data:

1. Lawfulness, transparency, and fairness
2. Purpose limitation
3. Data minimization
4. Accuracy
5. Storage limitation
6. Confidentiality and integrity
7. Accountability

United States of America

The USA doesn't have a comprehensive federal law that regulates data protection and privacy. Instead, different federal and state laws cover various aspects of data privacy, such as health and financial data, but data collection remains largely unregulated in most states.

- **Privacy Act of 1974:** A federal law that regulates how the United States government can collect, maintain, use, and disclose PII about individuals in its records.
- **Health Insurance Portability and Accountability Act (HIPAA):** A federal law that protects the privacy and security of individuals' health information.
- **The Gramm-Leach-Bliley Act (GLBA):** A federal law requiring that financial institutions explain to customers how they share and protect their private information.
- **Children's Online Privacy Protection Act (COPPA):** A federal law limiting how online businesses can collect, share, and use personal information from children younger than 13 years old.
- **California Consumer Privacy Act (CCPA):** A state law that applies to businesses operating in California granting California residents protections of their personal information.

4

Success Stories

Success Stories

Find out how using the Salesforce Platform helped these companies integrate trusted generative AI into their processes to improve customer engagement.

Crexi

Before Agentforce: Admin tasks steal precious hours from customer relationships

Crexi sales reps rely on conversations to make business deals. Unfortunately, critical admin tasks like drafting emails, writing call notes, and performing customer research often monopolized their time. These tasks stole precious hours away from customer-centric conversations, preventing sales reps from making human connections that helped close deals.

After Agentforce: Generative AI provides an always-on assistant to handle admin tasks

With Agentforce, Crexi was able to optimize its sales process and give time back to sales reps so they could make more human connections with customers. Generative AI and agents help handle tasks that don't necessarily
need a human touch, performing research, generating key insights, and producing summaries of conversations and meetings. With Agentforce handling much of the admin work, **80%** of the sales reps' day has been freed up to focus on customer engagement.

Read the full story

Iron Mountain

Before Agentforce: Manual processes slow down support

The service agents at Iron Mountain had a wealth of data, but their manual ordering process was slowing down customer support. Agents had to switch between multiple apps and databases to get the context they needed for each customer order and inquiry. The result was long lead times, a growing backlog of inquiries, and a mountain of customer data that agents couldn't use efficiently.

After Agentforce: AI-generated responses ease the pressure on service agents

With Agentforce, agents automatically get personalized case replies based on their knowledge base and past cases. Agents review the generated response in case it needs editing, but **76%** of the time, the AI-generated reply is good to go without requiring a human's touch. Overall, **85%** of service agents said the AI-generated replies were contextual and accurate and helped reduce repeat calls by **8%**.

Read the full story

Conclusion

Infuse Generative AI With a Layer of Trust

Infuse Generative AI With a Layer of Trust

The Salesforce Platform helps you create a trusted foundation for generative AI in customer service, sales, marketing, and commerce – the entire customer 360-degree experience. Salesforce unifies your data, CRM, AI, and security in a single platform so you can keep your data secure and private while implementing AI systems that let you refocus human attention where it matters most.

With Salesforce, security is a shared commitment. We partner with our customers to build a trusted foundation for generative AI. At Salesforce, when it comes to AI, we start with trust and security in mind. The Einstein Trust Layer handles tasks like secure data retrieval, dynamic grounding, data masking, audit trailing, toxicity detection, and zero retention. This lets you use LLMs to create more efficient and personalized interactions without compromising the privacy and security of your CRM data.

Building trust is central to a successful transformation with generative AI. With that foundation in place, your business will be freer to explore the full potential of generative AI while fostering a culture of confidence and responsibility, which are essential for long-term success in today's data-driven world.

Explore the Salesforce Platform

Contact an expert to learn more