



GenAI for next-gen governments

January 2024

Executive summary

Generative AI (GenAI) has the power to transform governments by facilitating inclusion and improving productivity and accessibility. GenAI, powered by data and driven by sophisticated algorithms, can create novel solutions for governments through informed decision making and hyperpersonalised citizen interaction.

GenAI has the potential to radically transform both the public and private sectors. It has the potential to significantly boost individual productivity while also bridging the digital divide by promoting equitable information access and processes. GenAI, like any other technology, has its unique set of challenges. Issues such as algorithmic bias, accountability, privacy and fairness become central when deploying GenAI for public use. This paper advocates the use of GenAI in public service delivery and governance, and emphasises the necessity of developing guardrails for the safe and responsible use of GenAI in the public realm.

Next-gen governments would have hyperpersonalised public service delivery, transparent functioning and improved public trust. The biggest contribution of GenAI in governance would be in removing information asymmetry for citizens. However, unlocking the full range of GenAI's capabilities would require continuous pursuit of research, collaborative efforts and supportive regulatory measures.

The 'Generative AI for next-gen governments' initiative is a significant leap forward towards helping governments deliver efficient and inclusive governance using GenAI. This paper suggests ways to leverage GenAI's transformative impact on governance while managing the associated risks and challenges.



Table of contents

1.	Generative AI – an overview	4
2.	GenAI for next-gen governments	10
3.	Adopting responsible GenAI	15
4.	Next-gen governance through GenAI for India	24
5.	Conclusion	26



1. Generative AI – an overview

1.1. The GenAI era

2023 was the year of generative AI (GenAI). GenAI, which began with a rule-based approach in the 1950s, has evolved rapidly with breakthroughs like generative adversarial networks (GANs) and transformers. Over the last few years, some ground-breaking work has happened in foundational large language models (LLMs). The first two quarters of 2023 saw an exponential rise in the use of GenAI tools triggered by the launch of ChatGPT, a general-purpose AI chatbot that took human-machine interaction to the next level.

The popularity of ChatGPT also led to the launch and rediscovery of several other GenAI tools. Today, we have AI tools with a wide range of capabilities – from creating websites to composing music, from editing images to generating 3D models from scratch, and from creating presentation slides to reviewing and writing codes. GenAI is used not only by developers and technologists, but also at an individual level to augment human intelligence.

Artists and creative professionals are liberally using GenAI tools to generate new ideas and impactful, rich content. These tools have been used for creating website templates, original images, 3D models and animated videos by providing just a simple text prompt. This has led to a significant boost in productivity.

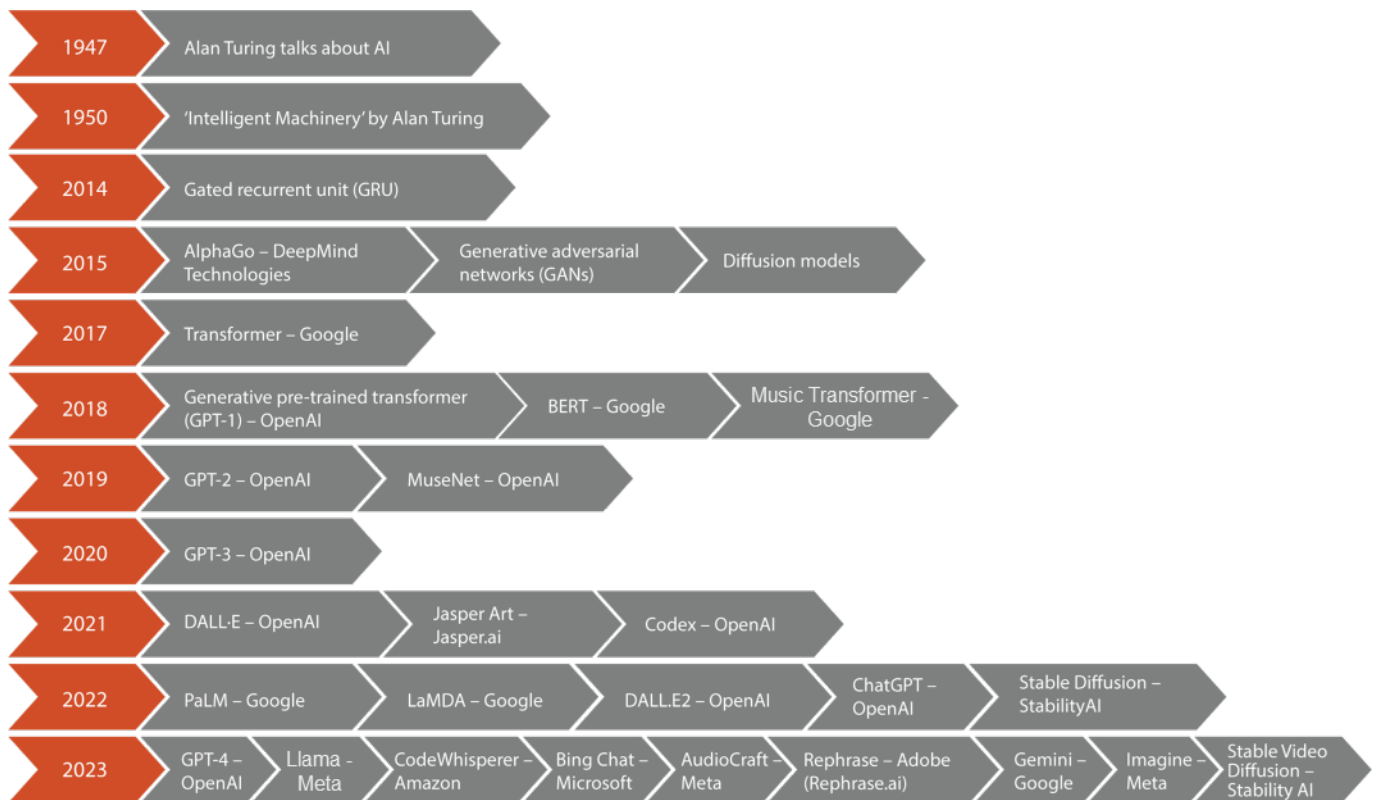


Figure 1: Evolution of GenAI

1.2. GenAI for social good

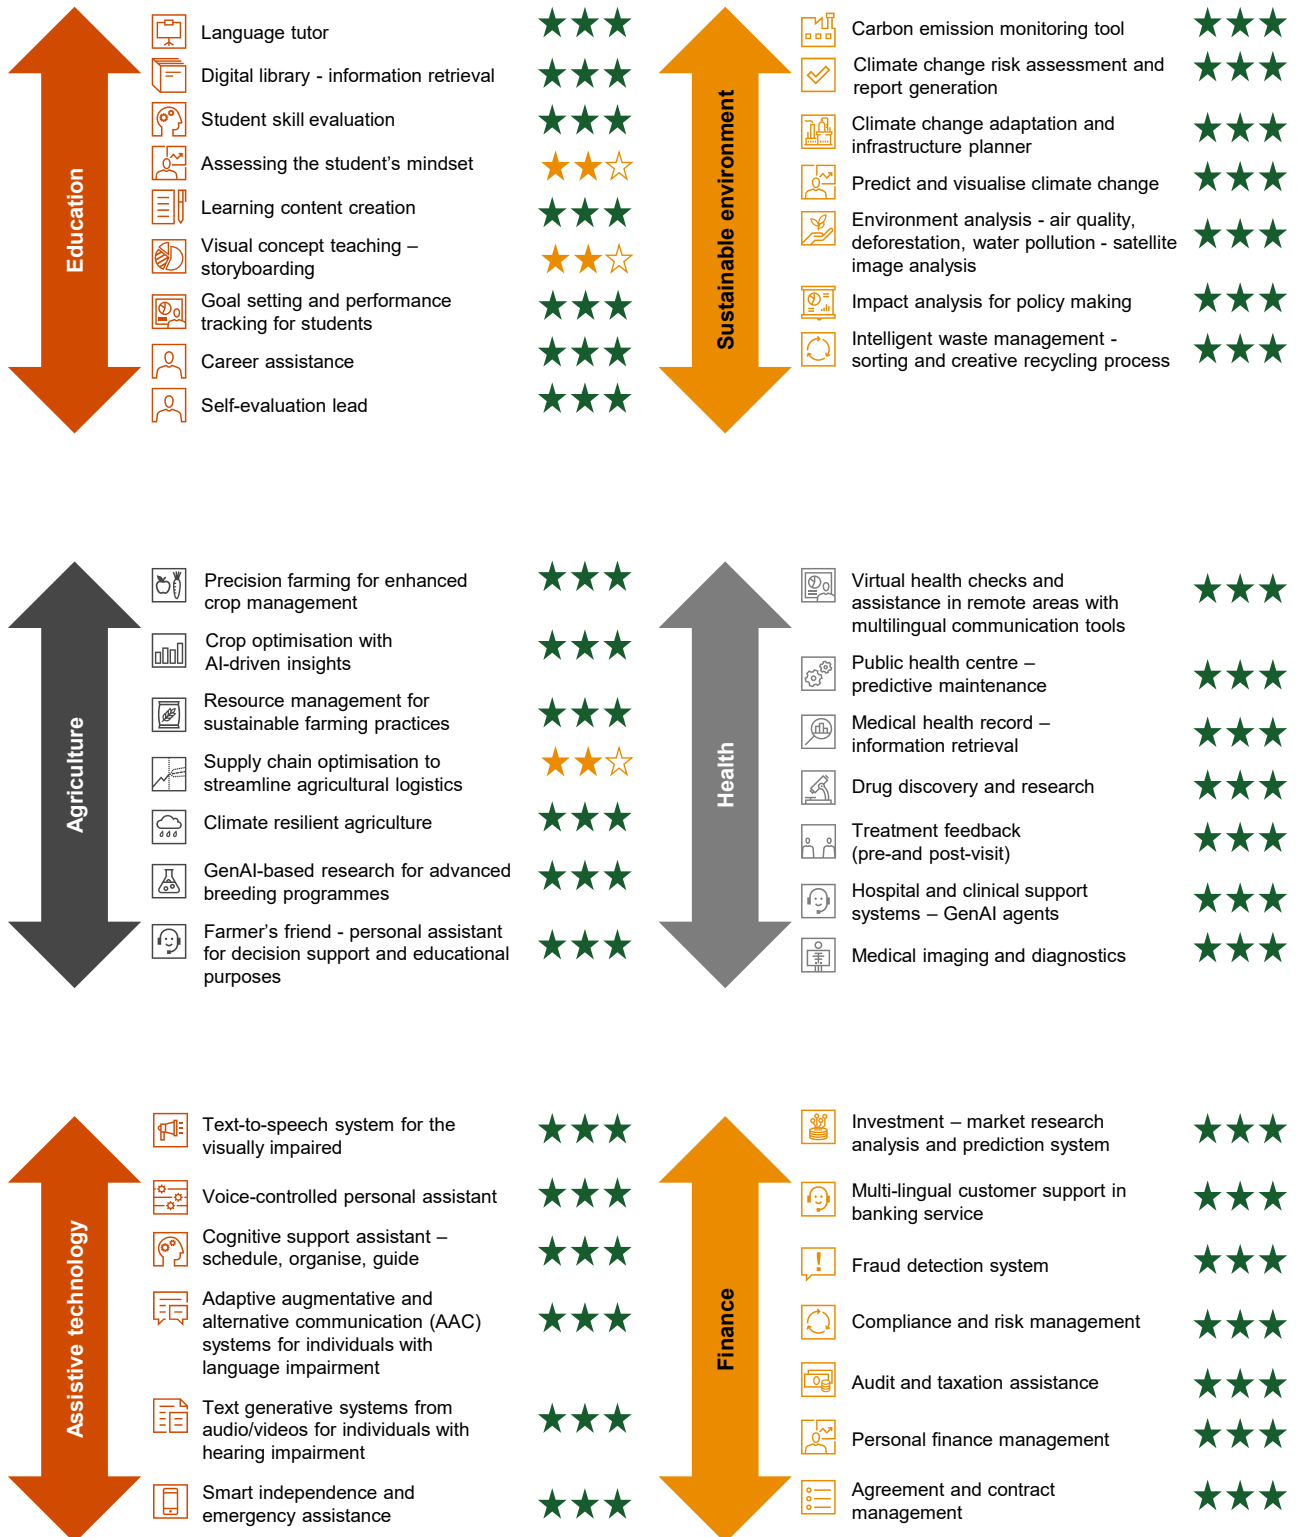
Governments worldwide share certain common goals – safety and well-being of their citizens, economic growth, poverty reduction, environmental protection, responsible living, etc. The United Nation's Sustainable Development Goals (SDGs) serve as a universal framework for governments to measure their progress.¹ By harmonising their policies and actions in line with the SDGs, governments can effectively tackle a broad spectrum of societal challenges, ranging from healthcare and education to gender equality and environment protection. Distinguished scholars, industry leaders and technology innovators are convinced that GenAI can assume a pivotal role in elevating the quality of human life.² GenAI has the potential to positively impact critical challenges facing humanity such as inequality, lack of access to quality education, unemployment, lack of healthcare, rapid environmental degradation and violation of human rights.³ For example, GenAI solutions in the education sector can dramatically improve learning by transcending linguistic barriers and curating the best lectures from around the world on a given topic and by creating visually rich multimedia content to explain difficult concepts. A wide spectrum of sectors will experience significant transformation by adopting GenAI solutions (see Figure 2).

¹ The SDGs in Action | UNDP

² Dwivedi, Yogesh K., et al. 'So what if ChatGPT wrote it?' Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy." International Journal of Information Management 71 (2023): 102642.

³ ChatGPT in the Public Sector – overhyped or overlooked? | Council of European Union | April 2023

Positive social impact of GenAI



Impact on SDGs

- ★★★★ High impact
- ★★★ Moderate impact

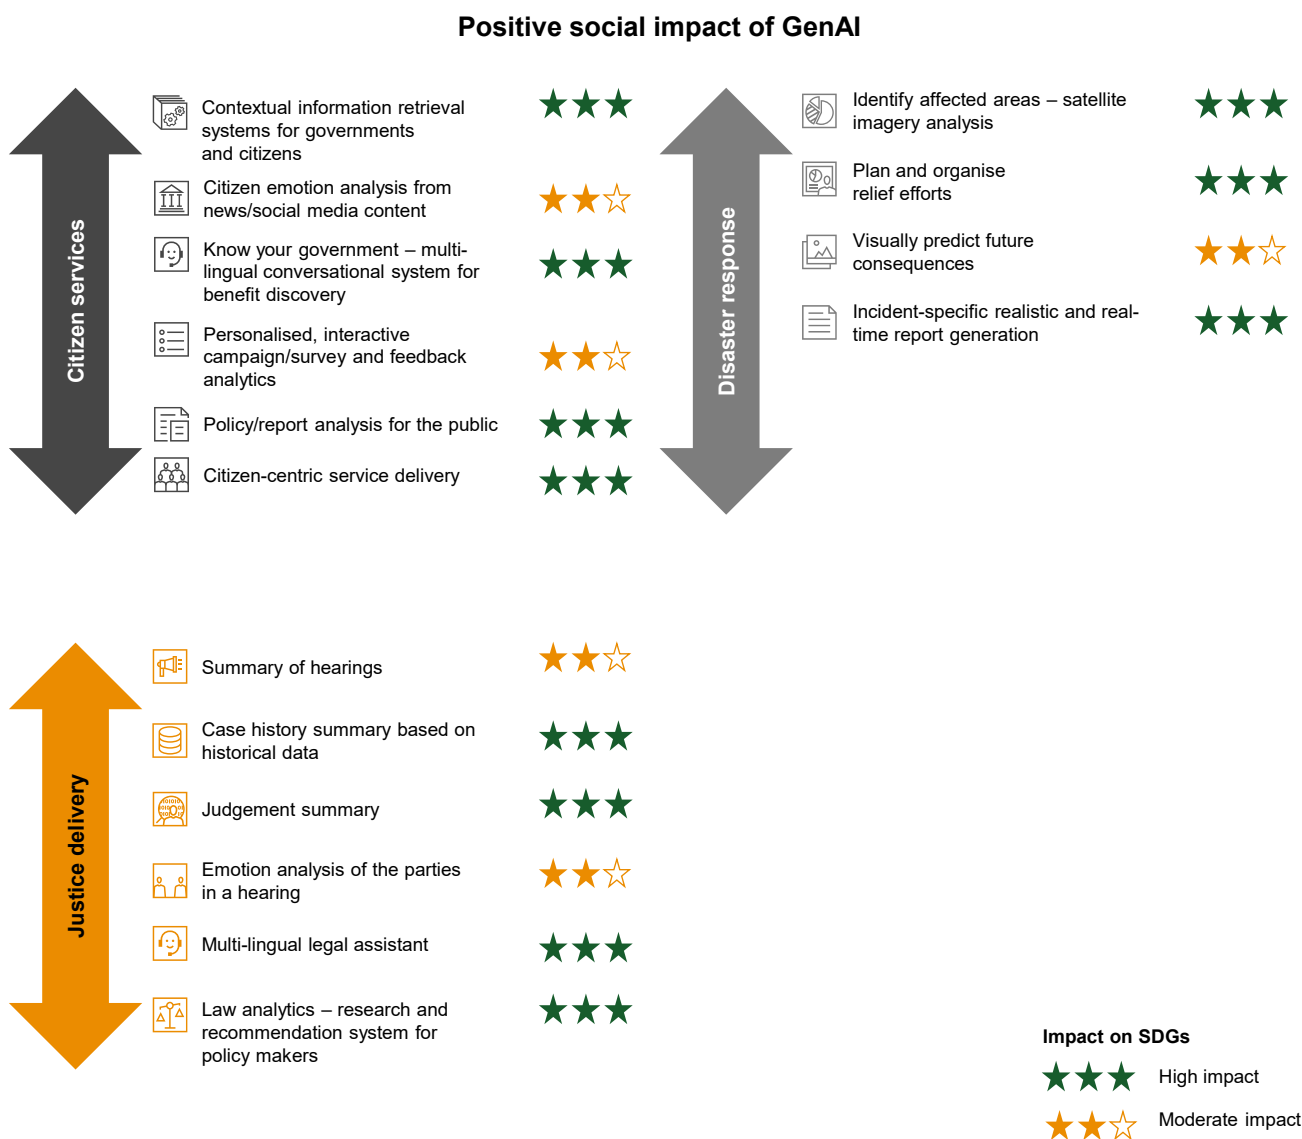


Figure 2: Sectoral use cases of GenAI solutions for next-gen governments

The World Health Organization (WHO) highlights that half of the global population lacks access to essential health services and high medical expenses are pushing many to poverty.⁴ One of the SDGs is facilitating universal access to fair, safe, trustworthy, inclusive and quality healthcare services. For this, governments can use a GenAI-powered healthcare agent, available on a mobile device, to guide citizens in the time of distress. Whenever a citizen needs healthcare-related advice, they could obtain it through the GenAI-powered healthcare app and receive personalised help in a language/dialect they understand, under the supervision of a qualified health care professional.

Integrating GenAI into policies and governance processes would facilitate data-driven decision making for governments, accelerating their progress towards the SDGs and ensuring a sustainable and equitable future for all.

⁴ <https://www.who.int/home/13-12-2017-world-bank-and-who-half-the-world-lacks-access-to-essential-health-services-100-million-still-pushed-into-extreme-poverty-because-of-health-expenses> | WHO | December 2017

Select sectors where GenAI can make an impact in service delivery include:



Healthcare

GenAI can provide personalised health assistance and an interactive health guide to citizens. It can be a useful tool for preventive healthcare. In the case of hospitalised care, it can transform the patient's overall treatment experience and ensure a quick recovery built on transparency and trust.⁵ For example, it can generate discharge summaries with doctor's instructions in the native language of the patient, with annotated medical images and diagnostic tests for better understanding. GenAI systems can be a great partner for running health awareness campaigns containing easy-to-understand educational videos, images and summaries. These systems can even help expedite drug discovery by quickly generating molecular structures and identifying potential treatments.



Environmental sustainability

GenAI can generate synthetic data sets that mimic real work scenarios, analyse trends and patterns to provide insights for optimal natural resource utilisation and can help create solutions to address environmental challenges. Some key impact areas where GenAI can be used to promote sustainability include:

- analysing and predicting energy consumption, helping to create energy optimisation strategies for organisations and manage distributed energy resources by balancing the flow of electricity in grids, reducing peak demand and integrating renewable energy sources
- analysing and interpreting satellite imagery to detect climate change indicators such as changes in land use patterns and green cover, and detecting and monitoring greenhouse gases
- using 3D models and simulations to facilitate the integration of green technologies, plan sustainable and eco-friendly constructions, and aid virtual testing and optimisation before construction begins.



Individuals with special needs

GenAI can augment accessibility of information and inclusion of specially abled individuals by addressing their unique needs through:

- text-to-speech generation for visually impaired individuals
- transcription of audio or video content for the hearing impaired in their preferred languages.

⁵ Kerasidou, Angeliki. 'Artificial intelligence and the ongoing need for empathy, compassion and trust in healthcare'. Bulletin of the World Health Organization 98.4 (2020): 245.



Education

GenAI can support personalised learning through the generation of customised educational content and adaptive tutoring systems. It can assist in creating interactive sessions using virtual reality (VR) experiments and immersive learning environments (e.g. visually explaining how solar and lunar eclipses occur in VR-assisted interactive sessions), leading to much superior learning outcomes.



Disaster response

GenAI can simulate multiple real-world disaster scenarios with geography-specific context and customised level of complexities to help first responders and planners analyse, prepare, adapt, decide and act on them.⁶ GenAI autonomous agents armed with satellite imagery and geographic information systems (GIS) data can significantly improve a city's response to flood monitoring and management. City administrators can simulate the most likely scenarios, visualise the potential impact in each one and accordingly prepare a tailored action plan.

GenAI technology can help build better disaster management solutions by providing rapid risk assessment capabilities and creating autonomous emergency response agents that can be assigned to a micro-geography or targeted population (such as the elderly).

Synthetic data generation:

GenAI can aid governments in predicting the impacts of disasters like natural calamities and epidemics through synthetic data generation and training of AI models, enhancing disaster prediction and response capabilities. Synthetic data generated heat maps hold the key to predicting weather patterns or a pandemic spread model. This would help governments prepare in advance and plan better to handle disasters and epidemics.

1.3. New growth opportunities with GenAI

Governments can leverage GenAI as a major driver for economic growth and job creation. Goldman Sachs research estimates that GenAI could raise the global GDP by 7%.⁷ GenAI has the potential to contribute USD 15.7 trillion to the global economy by 2030 and to transform sectors such as healthcare, automotive, financial services, transportation, logistics, technology, communications, entertainment, retail, energy and manufacturing.⁸

Historically, technological innovations have always led to the creation of more jobs and have generated a demand for new skills, and GenAI is no different. GenAI is going to be a force in boosting economic productivity and creating new jobs such as prompt engineers, responsible GenAI specialists and GenAI public safety certifiers. Yet, the evolution of GenAI is likely to reshape the workplace by replacing workers involved in repetitive, predictable and structured tasks. Governments would need to invest in capacity building of their workforce to help them 'navigate' the GenAI wave better.

⁶ CrisisAI: A novel hybrid AI system for crisis management | European AI Alliance, European Commission | November 2023

⁷ Generative AI could raise global GDP by 7% | Goldman Sachs | April 2023.

⁸ Sizing the prize. PwC's Global Artificial Intelligence Study: Exploiting the AI revolution | PwC



2. GenAI for next-gen governments

Governments worldwide, using digitalisation, want to make their public service delivery systems efficient, transparent and inclusive. Digital government is the new mantra. The following are the five key milestones in its evolution:

1. digital delivery of individual services
2. scaled up value chain for services
3. automation for improved service delivery
4. proactive service delivery, and
5. hyperpersonalised citizen services.⁹

⁹ Why digital government? The journey to Digital Government 5.0 | PwC

According to the UN's E-Government Development Index 2022, Denmark ranks first amongst 193 countries in delivering online services. It has effectively used information and communications technology (ICT), built sound digital infrastructure, invested in developing human capital and encouraged e-participation.¹⁰ Globally, governments face challenges in digital service delivery due to issues related to digital inclusion, digital literacy, language barriers, data security, data privacy, user trust and poor adoption.

The question that arises is, how can next-gen governments use GenAI for good governance? The barriers to good governance are (i) information asymmetry, (ii) controlled access to information, (iii) difficulty in understanding information (legal or semi-legal style) and (iv) language barriers (information being in another language). GenAI can easily break these barriers. It helps demystify governance by translating complex government documents and rules into a simple, easy-to-understand, conversational dialogue format in a language or dialect citizens are most comfortable with. It thus builds citizen trust and makes them self-reliant. An interesting outcome of GenAI is that it makes people more compliant with regulations as their awareness level increases. GenAI is therefore a powerful tool for ease of compliance. It makes public services user-friendly, promotes inclusion and dislodges rent seekers if there are any. A well-designed GenAI interface can help citizens quickly identify government programmes meant for them, enabling them to apply for benefits with ease. It can take on the role of a dedicated and personalised assistant for every citizen and help them track their requests till the service or benefit is finally delivered.

GenAI can also benefit governments by enabling their managers to take consistent, fast and correct decisions. It can digest voluminous data about laws, rules, exceptions, eligibility and citizen case histories (with permission from the citizen) and precedents. It can thus hyperlocalise and hyperpersonalise services for every individual.




Figure 3: How does GenAI transform governments?

¹⁰ UN E-Government Knowledgebase 2022 | Department of Economic and Social Affairs | United Nations

These developments lead us to a future where government operations are not just efficient and effective but also equitable, inclusive and truly representative of the people they serve.

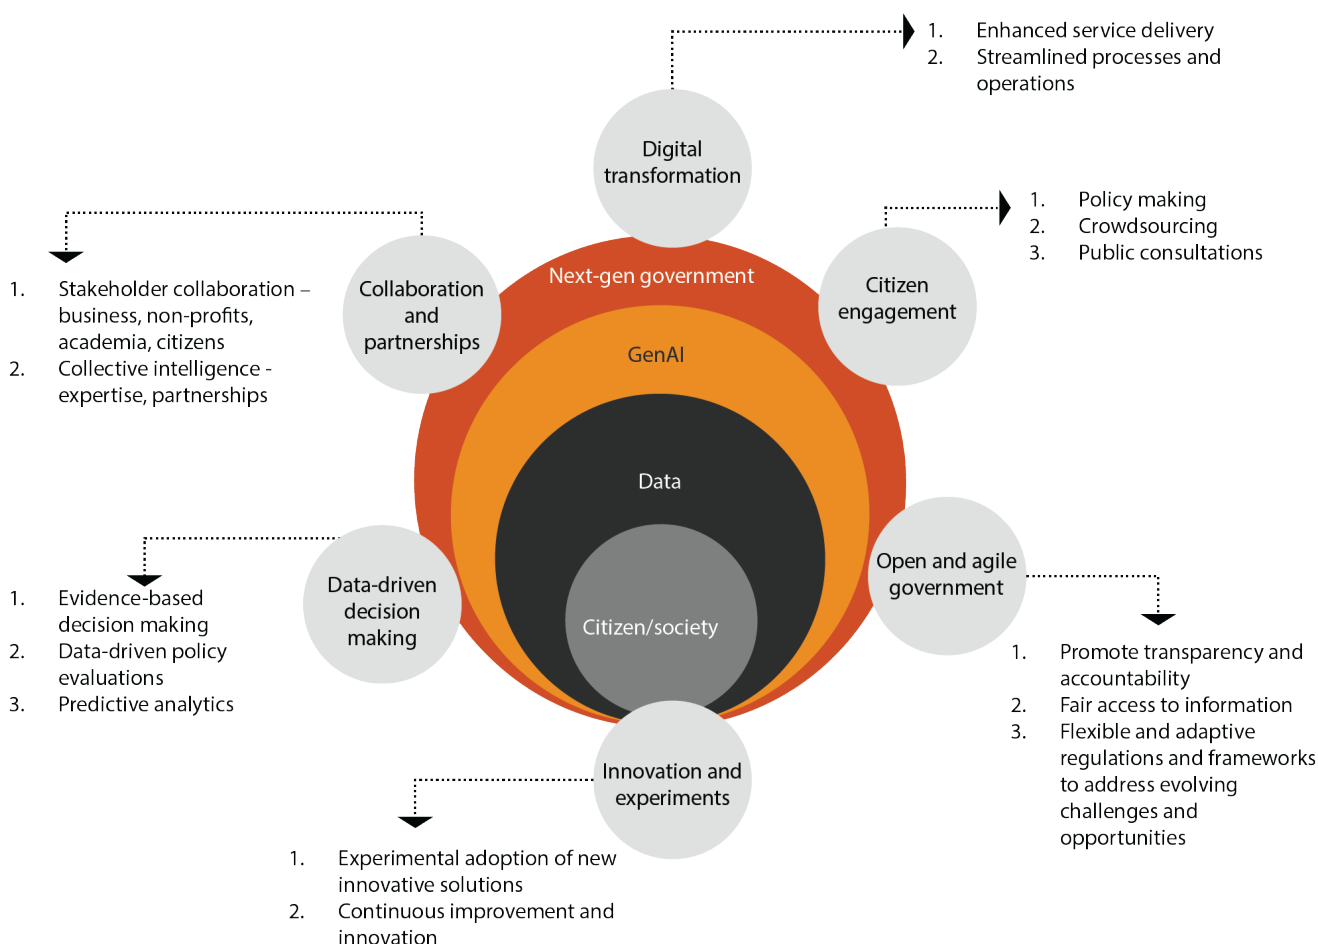


Figure 4: What does a next-gen government facilitate?

A next-gen government is envisioned as a model that keeps citizens at the centre of governance and focuses on inclusion, barrier-free access, hyper-personalisation and hyper-contextualisation of individual needs.

2.1. Decision making

Policymaking: GenAI interfaces trained on historical, economic and social datasets can quickly identify emerging trends and patterns and provide insights for policymaking.¹¹ It can also help in testing alternative policy models. Governments can employ GenAI at different phases of the policymaking lifecycle, namely problem identification, policy formulation, impact assessment and information dissemination.

Citizen grievances and concerns: GenAI tools can provide insights into citizen concerns and popular issues by analysing a large corpus of data from social media platforms. These tools can be used to address citizen grievances by intelligent routing to the right stakeholder. GenAI can provide precise and personalised information to a citizen and thus help in proactive citizen service delivery.¹²

¹¹ Building resilience in government using data science | The Alan Turing Institute | June 2023

¹² Generative AI and public sector | Public Sector Center of Expertise | Microsoft

2.2. Service discovery – citizens and businesses

In some cases, eligible citizens remain deprived of welfare programmes intended for them due to geographical or information disadvantage. Sometimes the complexity of application forms and documentation needed for applying for a government benefit or subsidy is a deterrent for socially disadvantaged groups. The targeted population may remain unaware of a welfare programme due to lack of information access or absence of information in its local language. This usually leads to mushrooming of middlemen/brokers and creates rent-seeking opportunities.

GenAI can eliminate the need for middlemen and provide easy access to relevant information about all government programmes in a language best understood by the beneficiary. For example, an elderly farmer can apply for old age pension with the help of GenAI. GenAI can gather the required information from a beneficiary in a simple conversational mode where answers can be given by the user in voice, text or image (such as ID).

Similarly, while seeking government approvals, licences and permissions, businesses can use GenAI systems which can provide accurate and bespoke information. Sometimes a great deal of entrepreneurial economic value gets locked up due to delays and lack of access to right information. This problem can be addressed by using GenAI-based personalised agents.

2.3. Removing information asymmetry and boosting compliance

Government laws, policies, rules and guidelines are often complex and difficult to understand for most people. This is further complicated by the existence of multiple court rulings and interpretations on each subject matter. Often, the institutional memory of government organisations resides across multiple siloed document repositories. GenAI can easily consume this voluminous data and be trained to deliver the right information about applicable government regulations and programmes for a given situation.

GenAI can be a great assistant to government employees as well and can help them make efficient and better decisions. Helping citizens and businesses navigate the regulatory landscape will promote compliance, reduce litigation and improve governance.

2.4. Personalised citizen service delivery – education and healthcare

Education and healthcare stand to benefit the most if personalised care can be offered to every individual. GenAI allows hyperpersonalisation at scale. Each individual student can be offered customised lessons depending upon their interests, learning strengths, weaknesses, personal history of learning and aptitude using GenAI-based learning systems. The visual educational content can also be customised to suit every student's cultural background, region and preference for pace, language and depth.

Similarly, personalisation can be a gamechanger in healthcare. A recent study published in the Journal of the American Medical Association (JAMA) found that GenAI systems tend to be more empathetic towards patients than doctors while answering their queries.¹³ GenAI can solve challenging problems in the healthcare industry by leveraging LLMs to summarise medical reports, describe diagnostic processes, explain treatment options and build conversations with patients in the language of their choice. GenAI tools can transform medical and conversational data into a form that a clinician can use for further consultation. For example, when a GenAI-powered wearable device identifies an abnormal condition such as elevated heart rate or sudden excessive sweating, it can communicate with the patient to learn how they are feeling and share that conversational data with the healthcare professional for better diagnosis.

¹³ Ayers JW, Poliak A, Dredze M, et al. 'Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum'. JAMA Intern Med. 2023;183(6):589–596. doi:10.1001/jamainternmed.2023.1838

2.5. Justice delivery

The justice delivery system can be one of the leading beneficiaries of GenAI. With courts facing a huge backlog of cases, GenAI can help reduce the judicial burden by a great deal. It can help in expediting the delivery of judgments by analysing legal documents, producing concise and understandable summaries, answering complex legal queries, citing case laws, and assisting in transcribing and auto-correcting court proceedings. It can thus free up a lot of valuable time for judges by creating a customised order template and by pre-filling key details of a case by adopting a judge's writing style. By becoming a personalised assistant to court staff and judges, GenAI can substantially speed up judgments.

While GenAI can deliver high-quality outcomes for justice delivery, they are not reliable for complex judicial tasks like legal research and analysis.¹⁴ Since the outcomes have direct implications on human lives and society, it is essential to establish safety guidelines for GenAI and educate the stakeholders on safe and ethical use of this technology.

2.6. Autonomous agents under human supervision

Governments face challenges in analysing large volumes of data for their needs. GenAI can create scores of autonomous agents trained on domain-specific data (e.g. past data on floods and rain to assist a municipal flood control unit) to work under human supervision. These agents would be equipped with rapid decision-making capabilities and capable of handling multiple scenarios.

Intelligent agents can work alongside existing government systems to understand the context and then accordingly trigger the right business processes through a dedicated AI platform. For example, a smart irrigation control system, integrated with GenAI, can be remotely activated by an intelligent agent when certain conditions are met. The irrigation system, guided by GenAI algorithms, will be able to respond and adjust irrigation parameters based on monitored parameters, optimising water usage.

GenAI can help create explainable solutions and highlight the key factors that contributed to decision making. For example, a public service official can use a GenAI application to evaluate the data submitted by citizens and discover services they are eligible for. The model can be traced back to the checkpoints to help explain how the decision was made.



¹⁴ GenAI before the courts: the legal risks in using artificial intelligence | UNSW Sydney | Jan 2024



3. Adopting responsible GenAI

GenAI systems are built on complex mathematical algorithms and trained on large volumes of content extracted from the internet. GenAI models rely on statistical patterns and probabilistic computation to generate responses. They do not have a human-like understanding. This probabilistic method leads to propagation of existing biases in the training content. There are several other issues with GenAI.

GenAI is known to 'hallucinate' and generate responses which are entirely fictitious and at times deceptive. Also, there are concerns around data privacy, consent, copyright and legal ownership, as GenAI systems are trained using data from the internet that may include data about people, organisations and protected intellectual property (IP). In addition, there are concerns around the carbon footprint of GenAI systems as they are built upon LLMs which have high energy requirements.

Thus, while GenAI systems are very useful to society, they also pose challenges and risks.¹⁵

¹⁵ Stahl, Bernd Carsten, et al. 'Exploring ethics and human rights in artificial intelligence—A Delphi study'. Technological Forecasting and Social Change 191 (2023): 122502.

Next-gen governments need to clearly understand these risks and work towards mitigating frameworks with the help of industry and academia. They would need to establish policies and regulations that allow safe usage of GenAI for social good.¹⁶ These policies would need to draw upon responsible AI principles^{17,18} such as transparency, fairness, inclusiveness, security, safety and reliability. GenAI systems would also need to align with the human principles of ethics, equity, privacy, transparency and accountability.¹⁹

3.1. Challenges and risks with GenAI adoption

The characteristics, efficiencies and risks of GenAI systems²⁰ are dependent on the following factors:

- the **data** used for training
- the **model** used (foundational models or fine-tuned models for specific behaviour)
- the **prompts** given (meaningful and contextual)
- and the **user’s** level of awareness (knowledge and training).

To effectively manage and mitigate GenAI risks, the above four factors need to be carefully analysed. Figure 5 describes how these factors influence risks.

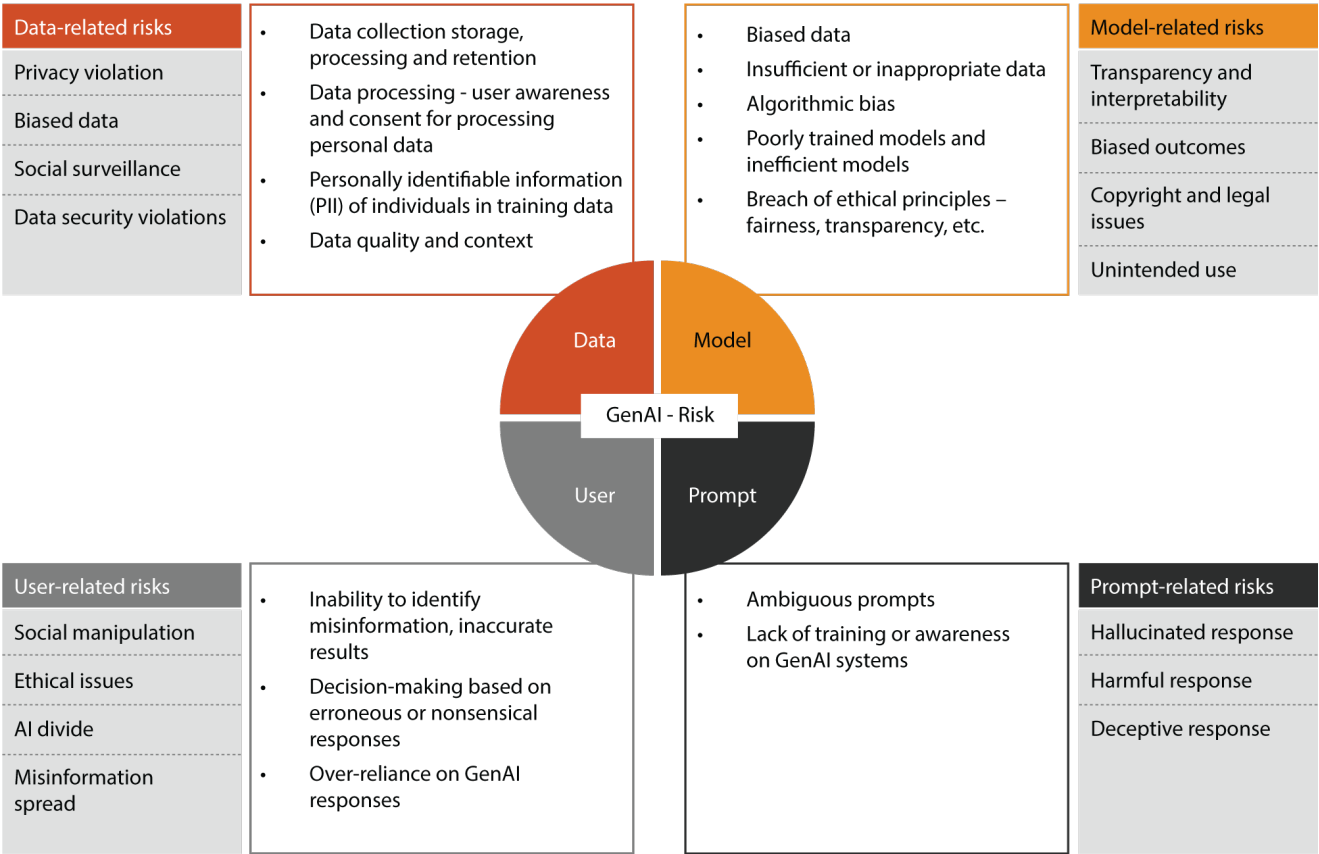


Figure 5: Causes of risks in GenAI systems

¹⁶ Policymakers focus on making generative AI safer for all | PwC | June 2023
¹⁷ Empowering responsible AI practices | Microsoft AI
¹⁸ Responsibility - Our Principles | Google AI
¹⁹ Önder, Murat, Israel Nyaburi Nyadera, and Md Nazmul Islam. The Palgrave Handbook of Comparative Public Administration: Concepts and Cases. Palgrave Macmillan, 2022.
²⁰ Managing the risks of generative AI | PwC | May 2023

Let us consider a GenAI-powered intelligent agent deployed for healthcare service delivery. The GenAI system would have capabilities to predict health issues, offer recommendations and assistance, facilitate doctor consultations, and disseminate information on remedies and healthy lifestyle choices.

Such a GenAI-powered system would need scrutiny from a risk perspective.²¹ Table 1 presents a thorough view of risks associated with such a GenAI-powered intelligent agent.

Table 1: Risk involved in deploying a GenAI-powered intelligent assistant

Use case						
The government has deployed a GenAI-powered intelligent agent to provide timely healthcare information and healthcare-related support to citizens. The GenAI system can predict health issues, provide recommendations and prompt guidance, schedule appointments with healthcare professionals, help plan and allocate resources, and manage social media channels for building awareness.						
Capabilities of the GenAI system						
<ul style="list-style-type: none">Converse with users in their preferred language.Analyse medical reports and images and explain the diagnosed medical condition to the user.Support doctors in decision making by helping gather and process additional information from the patient.Provide medical recommendations, advise on healthy lifestyle choices and increase public health awareness.						
Risk	Source of risk	Scenario	Sample prompt	GenAI system behaviour	Direct effect on	Anticipated adverse impacts
	(What causes the risk?)	(How do the risky scenarios occur?)	(How is it happening?)	(How is the system behaving?)	(Who is affected?)	(What are the impacts?)
Bias/ discrimination	Bias in data	The GenAI system has been trained on data from individuals with lighter skin tones. As a result, it consistently misdiagnoses skin conditions in individuals with darker skin tones.	How can I treat this skin condition? (shares a picture of the affected part)	Produces biased predictions (presence of skin disease)	Citizen	Misleading health predictions and recommendation

²¹ Wach, Krzysztof, et al. 'The dark side of generative artificial intelligence: A critical analysis of controversies and risks of ChatGPT'. Entrepreneurial Business & Economics Review 11.2 (2023).

Risk	Source of risk	Scenario	Sample prompt	GenAI system behaviour	Direct effect on	Anticipated adverse impacts
	(What causes the risk?)	(How do the risky scenarios occur?)	(How is it happening?)	(How is the system behaving?)	(Who is affected?)	(What are the impacts?)
	Bias in model	The GenAI-powered virtual assistant consistently recommends solutions without considering a citizen's socioeconomic background. For example, economically weaker citizens are advised costlier treatments, which prevents them from taking necessary action.	I'm experiencing symptoms such as frequent headaches, nausea and muscle pain in my arms.	Recommendation with socioeconomic disparities	Citizens	Discriminatory healthcare recommendation
	Bias in algorithm	The GenAI system recommends specialists based solely on geographic proximity, which is a disadvantage for patients in rural and distant places with limited access to healthcare.	Suggest a specialist for my health condition.	Algorithmic bias in specialist recommendations	Citizen	Unequal access to healthcare services and resources
Data privacy	Improper data handling or security breaches	The health data collected from citizens during the engagement period might be exposed to third parties, leading to sensitive information leaks and privacy issues.	Can I get a report of my health history in the past two years?	Collecting and storing health data for a longer period without proper consent Breach of personal health records without a proper authentication mechanism	Citizen	Leads to user privacy violations, financial loss, reputational loss and a setback in mental well-being. Reduces trust and confidence in the deployed system.

Risk	Source of risk	Scenario	Sample prompt	GenAI system behaviour	Direct effect on	Anticipated adverse impacts
	(What causes the risk?)	(How do the risky scenarios occur?)	(How is it happening?)	(How is the system behaving?)	(Who is affected?)	(What are the impacts?)
Hallucination	Incorrect response generation	The response generated may be relevant to the symptoms mentioned by the user, but it may not be related to the actual medical condition of the user, leading to unnecessary panic and medical tests.	What could be the reason for my fatigue?	Generation of false symptoms and misleading advice	Citizen	False medical information, harmful advice
Misinformation	Lack of fact-checking in the response generated	The response generated may include an unproven home remedy, leading the citizen to ignore necessary medical treatments.	How can I treat tonsillitis naturally at home?	Spread of unverified information like home remedies	Citizens	Incorrect health decisions, inappropriate actions
Explainability	Lack of transparency in GenAI decisions	Patients receive a diagnosis report without any explanation, leading to doubts and confusion about the system's accuracy and credibility.	What is my diagnosis? Can you provide me with a report on the accuracy of your predictions?	Uncertainty	Citizens, healthcare professionals	Mistrust and inability to understand the relevance of decisions made
Responsibility	Unclear accountability	A citizen followed the recommendation made by a GenAI system and experienced adverse effects. Who is responsible? Is it the patient, the government or the developer?	How can I treat bronchitis at home without antibiotics?	No clarity on who is accountable	Citizen, doctors, product owner	Conflicting roles in decision making

Risk	Source of risk	Scenario	Sample prompt	GenAI system behaviour	Direct effect on	Anticipated adverse impacts
	(What causes the risk?)	(How do the risky scenarios occur?)	(How is it happening?)	(How is the system behaving?)	(Who is affected?)	(What are the impacts?)
Social manipulation	Malicious use of AI-generated content	The system might be influenced to disseminate false information about a disease outbreak, causing panic and ineffective response measures.	Tell me about the spread of dengue in my locality?	Spread of false information, leading to public risks in health and governance.	Citizen, government	Harm to society and the community due to spread of false information and risky action insights during an outbreak
Digital divide	Unequal access to technology	Citizens without smartphone and internet access will miss out on health consultations and clinical support systems that are specially built for efficient service delivery	Can you book me an appointment at the nearest healthcare centre that offers treatment for the medical diagnosis you predicted?	Widening the gaps in access to health services	Citizen, government	Exacerbated healthcare disparities
IP rights	Ownership of GenAI-generated content	The GenAI system is trained on data from various sources. Disputes or confusion arises when the government claims ownership of the GenAI's medical recommendations.	What are the best treatment options?	Ownership disputes over GenAI response	Citizen, doctors, government, product owners	Ambiguity in content ownership

The above table demonstrates that the risk categories are not mutually exclusive. In the biased data example, the failure could be attributed to lack of representative data for training, or it could also be viewed as a 'model bias', where the model is unable to capture optimal features for accurate prediction. Thus, the data, model and prompt are interconnected aspects of the GenAI system. Risks associated with each tend to influence others and understanding these interrelationships is important for developing a robust risk assessment and mitigation framework for GenAI.

3.2. Risk assessment and mitigation framework for GenAI

Next-gen governments would need to adopt responsible GenAI principles and use GenAI efficiently for public service in a safe and responsible manner. Before deploying GenAI systems in public service, governments would need to assess how the system has been developed (including assessing its training data) and how it is going to be used in the decision process. Consider a GenAI-powered digital platform that processes a loan waiver request from a farmer. The procedure would involve multiple levels of screening using GenAI tools. Before deploying such a solution, it is essential for governments to enable a fair, accountable and transparent process. It is also important to educate frontline government officials about the potential blind spots, pitfalls and risks of using the GenAI system. The frontline managers would need to be empowered to overrule GenAI system advice or decisions where they go against the objectives of deployment.

To ensure a thorough assessment of risks associated with the GenAI system, governments may use a three-stage risk assessment and mitigation framework (see Figure 6). To ensure safe and reliable adoption of GenAI systems, governments would need to develop policies and frameworks that identify and manage risks, assess their impact and employ risk mitigation strategies.

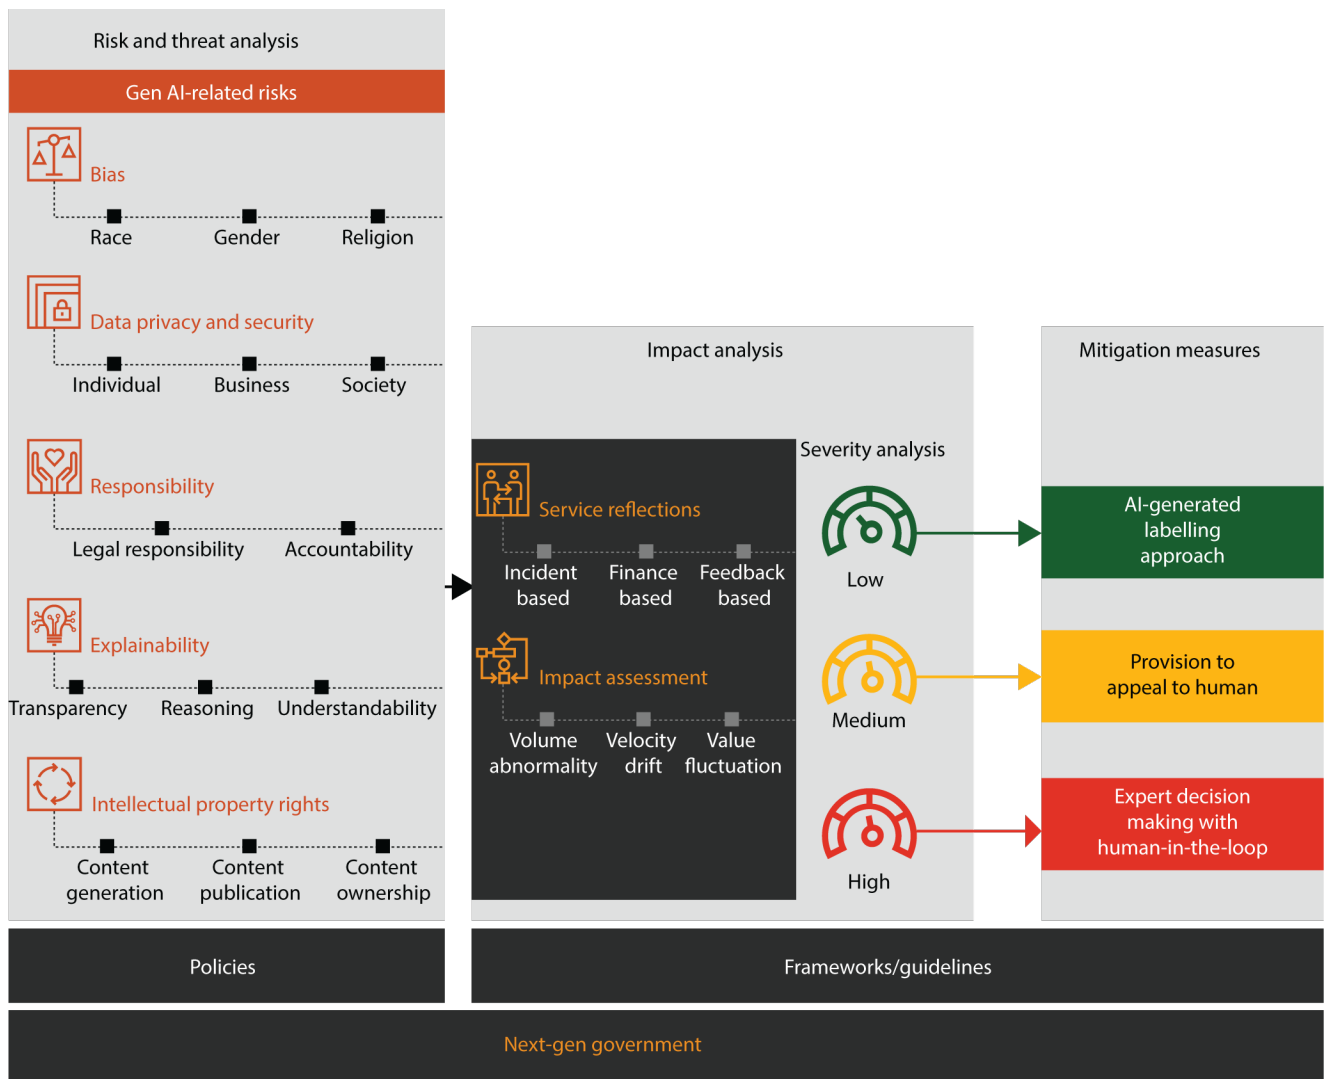


Figure 6: How can next-gen governments analyse and mitigate risks in GenAI systems?

Risk assessment

Governments have to ensure responsible and safe use of GenAI systems in the public domain. They would also need to perform risk assessments of GenAI systems, and examine challenges such as biased outcomes, hallucination, privacy, cybercrimes, accountability, unintended use of GenAI system and intellectual property rights (IPR) violations. The risk assessment will help in categorising the GenAI systems as high, medium or low risk depending on their purpose and the potential of harm to an individual or society. This assessment will aid in ensuring that the GenAI system works as intended. It would also help in identifying potential hacking vulnerabilities and understanding the impact and severity of any GenAI system if it turns rogue.

Impact assessment

GenAI systems deployed for public use would need to be evaluated on three dimensions of impact, in case they are hacked or turn rogue:

1. Value	2. Volume	3. Velocity
How much is the financial impact? How much does it affect the user or society?	How many individuals are affected by the inaccurate responses of a GenAI system?	How quickly can a response impact people's lives?

GenAI models which rank high on any of the above three dimensions would need to be put through checks before they are approved for public use. They would also need more frequent evaluation for safety. Any model which checks two or more of the above dimensions would need to be scrutinised minutely on all safety parameters, including an audit of its training data.

Severity analysis

Governments would need to work with the industry and academia to set up a mechanism to classify public-use GenAI systems as high, medium, or low-risk systems and accordingly create safety guardrails for their usage. European Union AI Act²² provides a framework of risk categorisation and prohibits the use of high-risk AI systems which may violate the human rights or cause harm to health and safety of humans. The medium-risk systems must pass through a conformity assessment to ensure they do not pose risk to rights of the people. The low-risk systems can be deployed with little restriction as they pose virtually no risk to human life.

Risk mitigation and management

Some of the strategies followed to manage the risks of AI systems can be applied to manage risks in GenAI systems as well:

1. Mandatory human oversight	2. Appeal provisions	3. Labelling and disclosure
For high-risk GenAI systems, governments can mandate human subject matter specialist oversight to validate system decisions.	For medium-risk GenAI systems, governments can establish an appellate authority to review the automated decisions of GenAI systems. It may mandate a short timeline for disposing complaints against GenAI systems.	Governments can mandate that every decision, text, audio, video or image generated with the help of GenAI be clearly and accurately labelled.

²² EU Artificial Intelligence Act | European Parliament | June 2023

Table 2: An illustration of various risk levels, potential impacts and recommended mitigation strategies for GenAI systems in the public domain

Risk level	Risk scenario	Impact and severity assessment	Risk mitigation strategy
High-risk GenAI systems (Outcome is likely to directly impact human life and human rights.)	The GenAI system is used in the healthcare industry.	It could provide incorrect diagnosis or medical recommendations, leading to delayed treatment or neglect of a critical health condition.	Human oversight
	The GenAI system is used in a judicial system for faster and informed judgements.	It could deny justice or fail to provide safeguards as per the law.	
	The GenAI system is used for automated selection of job applicants.	It could deny a job to a well-deserving candidate due to biases like gender bias, social stereotypes, age, socioeconomic status, and algorithmic or systematic bias.	
Medium-risk GenAI systems (Outcome will directly or indirectly impact the well-being of humans, e.g. by affecting their socioeconomic status.)	The GenAI system is used in content creation, say a video for building awareness of a newly announced government programmes.	It could violate the IPR of an artist or a writer.	Appeal provisions and timely decision making
	The GenAI system is used for an aspirant's skill assessment and grading.	It could provide a toxic or offensive response, damaging the aspirant's self-confidence.	
	The GenAI system is used to analyse citizen feedback on a newly launched government initiative.	It could provide deceptive or misleading insights, resulting in the continuation of a failed programme.	
Low-risk GenAI Systems (Outcomes have limited potential to cause unintended harm/impact.)	GenAI system is used as a grievance redressal and information agent	It could give wrong or misguiding answers causing avoidable difficulty to citizens.	Transparent and clear labelling of GenAI-created content
	The GenAI system is used to create social media content for promoting or creating awareness about a government initiative.	It could hallucinate and end up spreading misinformation.	



4. Next-gen governance through GenAI for India


India is recognised as one of the global leaders of digital governance. With over 1.42 billion people having digital identities and 900 million people connected to the internet, India is a natural digital leader.²³ The Government of India (GoI), through its **Digital India programme**,²⁴ has established the **National Programme on AI** to promote the adoption of responsible AI in society. The Ministry of Electronics and Information Technology's (MeitY's) **Bhashini** initiative²⁵ focuses on building Indian large language model (LLMs) suited to the vernacular needs of the country. As a founding member of the **Global Partnership on Artificial Intelligence (GPAI)** that unites governments, industry, civil society, academia and startups, India along with 15 member countries has released a policy brief on '**Generative AI, Jobs and Policy Response**' for using GenAI for social good.²⁶

²³ India to overtake China as world's most populous country in April 2023, United Nations projects | United Nations | April 2023

²⁴ Achievements made under Digital India Programme | PIB 2022 | Ministry of Electronics & IT

²⁵ National Language Translation Mission | MeitY

²⁶ Adamoli, Matteo, et al. "Policy Brief: Generative AI, Jobs, and Policy Response." (2023).



These initiatives lay a solid foundation for governments in India to leverage the power of GenAI to move to next generation of governance by delivering hyperpersonalised and proactive services to citizens. However, governments would need to add a few more building blocks to harness the full potential of this technology for governance, economic growth, inclusion and sustainable environmental practices:

- promoting research and innovation in GenAI
- providing economical compute infrastructure for startups and researchers
- building data infrastructure for developing foundational GenAI models
- skilling the youth in GenAI, compute and data technologies
- setting up strong guardrails and safety mechanisms for the use of GenAI across governments
- providing hyperpersonalised services in healthcare and education
- facilitating sustainable economic growth in agriculture
- promoting industrial use of GenAI for improving productivity and sustainability.

Governments (both at the state and central level) would need to focus on the following aspects:

Building internal capacity

Next-generation governments would need to focus on developing indigenous sector-specific GenAI solutions to serve the industry and citizens. Considering the rapidly evolving technology landscape, GenAI governments would need to build internal capacities to keep pace with these developments and the challenges they pose.

Building data and compute infrastructure

Governments would need to open up the vast amount of data they hold to researchers, startups and other stakeholders. Naturally, such access must comply with privacy and data protection requirements and has to be made equitable. These datasets would be critical for building new public services in the fields of agriculture, health and education. Data infrastructure such as the National Agri Stack can be leveraged to provide farmers with easier access to credit, timely farm inputs, and pest and disease control advice. Data from initiatives such as the Small Farmers' Agri-Business Consortium (SFAC) and Open Network for Digital Commerce (ONDC) can directly connect farmers to consumers, enabling better value realisation for farmers.

Governments would also need to facilitate access to and affordability of GenAI-compatible compute infrastructure (primarily GPU compute). A decentralised approach to such compute infrastructure would be needed to catalyse the growth of the GenAI ecosystem beyond tier-1 cities.

Developing strategy, policy and safety frameworks

To be able to harness the full potential of GenAI for public services, governments would need to develop a strategy for AI and GenAI. They would need to work on policies and create safety frameworks to promote responsible use of GenAI in the public domain. Further, as GenAI evolves, the policies and safety frameworks would need to be updated at regular intervals to stay relevant.

These steps would enable India to realise the vision of **Viksit Bharat@2047**.



5. Conclusion

GenAI can be a transformational technology for governments for graduating to the next level of governance, where citizens and businesses get hyperpersonalised and customised services in a proactive manner through an interface and language of their choice. GenAI can unlock new growth for entrepreneurs and businesses by assisting them in understanding complex regulations and compliances in easy and simple language. However, to be able to realise the full potential of GenAI, governments would need to invest in capacity building, skilling, and building common data and compute infrastructure.

Before initiating large-scale implementation of GenAI, it is crucial that governments clearly understand the risks and challenges of this technology. They would also need to develop a risk assessment framework, safety guardrails and a responsible GenAI policy.



Contact us

Mr. Santosh Misra

Partner, PwC India
santosh.misra@pwc.com

Contributors

Mr. Santosh Misra

Partner, PwC India
santosh.misra@pwc.com

Mr. Sri Krishnakanth Prabhala

Director, PwC India
sri.krishnakanth.prabhala@pwc.com

Dr. Geetha Raju

Senior Associate, PwC India
geetha.raju@pwc.com

Editorial support

Vishnupriya Sengupta

Ruchika Uniyal

Dion D'Souza

Rubina Malhotra

Design support

G Gnanaraj

Shipra Gupta



About PwC

At PwC, our purpose is to build trust in society and solve important problems. We're a network of firms in 151 countries with over 360,000 people who are committed to delivering quality in assurance, advisory and tax services. Find out more and tell us what matters to you by visiting us at www.pwc.com.

PwC refers to the PwC network and/or one or more of its member firms, each of which is a separate legal entity. Please see www.pwc.com/structure for further details.

© 2024 PwC. All rights reserved.

Data Classification: DC0 (Public)

In this document, PwC refers to PricewaterhouseCoopers Private Limited (a limited liability company in India having Corporate Identity Number or CIN : U74140WB1983PTC036093), which is a member firm of PricewaterhouseCoopers International Limited (PwCIL), each member firm of which is a separate legal entity.

This document does not constitute professional advice. The information in this document has been obtained or derived from sources believed by PricewaterhouseCoopers Private Limited (PwCPL) to be reliable but PwCPL does not represent that this information is accurate or complete. Any opinions or estimates contained in this document represent the judgment of PwCPL at this time and are subject to change without notice. Readers of this publication are advised to seek their own professional advice before taking any course of action or decision, for which they are entirely responsible, based on the contents of this publication. PwCPL neither accepts or assumes any responsibility or liability to any reader of this publication in respect of the information contained within it or for any decisions readers may take or decide not to or fail to take.

© 2024 PricewaterhouseCoopers Private Limited. All rights reserved.

GG-January 2024-M&C 34854