COHESIONX

HARMONY IN EXPERTISE

The State of Generative AI in South Africa (2025–2028)

A Strategic Briefing for Enterprise

The State of Generative AI in South Africa (2025–2028) – A Strategic Briefing for Enterprise Leaders

Key Findings

- **Execution Has Overtaken Experimentation:** The era of generative AI (GenAI) pilots and proofs-of-concept is giving way to broad implementation. South African enterprises are moving beyond isolated experiments into programmatic, scaled deployments of GenAI capabilities. The focus is shifting from “What can the model do?” to “What measurable business outcomes can we reliably achieve at scale under proper controls?”

- **Adoption is High, Strategy is Low:** South Africa exhibits a paradox of enthusiastic bottom-up GenAI adoption within a top-down strategy vacuum. Employee use of GenAI tools is widespread – over half of the surveyed firms report that staff are using GenAI for work – yet **86% of companies lack a formal AI strategy. Only** ~13% have implemented any governance or ethics frameworks. This unsanctioned “shadow AI” usage yields pockets of productivity gains but poses significant risks (e.g., data leakage, compliance violations). It also represents the most significant near-term opportunity once brought under a unified strategy and governance, these grassroots efforts can be scaled for enterprise-wide value.

- **Foundational Capabilities Are the Bottleneck:** The major constraints on GenAI value are not the intelligence of the latest model, but the readiness of the surrounding foundations – namely, **AI-ready data, ModelOps processes, integration architecture, and change management**. A majority of organisations (over 55%) assess their data as **“not AI-ready,”** making poor data quality and access the primary execution bottleneck. Similarly, the lack of robust model lifecycle management and integration controls (including security, APIs, and identity) now limits value delivery more than model performance does.

- **Next-Generation AI: Agentic and Specialised:** The frontier of GenAI is moving toward **“agentic AI”** – autonomous goal-seeking agents that can plan, execute multi-step tasks, call APIs/tools, and act within workflows – and toward **domain-specific and small language models (DSLMs)** tuned for specialised tasks. These approaches promise higher accuracy and efficiency for specific use cases at lower cost-to-serve. Enterprises are beginning to complement general large language models (LLMs) with smaller models fine-tuned on their industry or business data, finding they can achieve equal or better results with a fraction of the compute resources.

- **Local Cloud Infrastructure Lowers Barriers:** Long-standing infrastructure hurdles are rapidly being overcome. Global hyperscalers have invested heavily in South Africa – for example, **Google Cloud opened a Johannesburg region in 2024**, and **Microsoft is expanding its data centres with a R5.4 billion investment** – providing low-latency cloud and AI services onshore. A new AI supercomputing facility (a partnership between

Cassava Technologies and NVIDIA) is also underway to boost local AI processing capacity. With improved access to high-performance computing and data residency compliance, the onus shifts to enterprises to leverage these platforms effectively. The historical excuse of "infrastructure limitations" is disappearing; now, **strategy, skills, and governance** determine success.

- **"Augmentation-First" is Imperative in the SA Context:** Given South Africa's **30%+ unemployment rate**, businesses must frame GenAI as a tool for **human augmentation, not as a means of automation-driven job cuts**. The goal should be to elevate employee productivity and create new value, rather than eliminate roles. This principle isn't just ethical – it's pragmatic. Positioning GenAI as augmentative helps secure employee buy-in and avoids reputational harm. Companies that succeed with GenAI will be those that reskill their workforce and redeploy freed capacity toward innovation and growth, rather than simply reducing headcount.

- **Multilingual AI is a Requirement, Not a Nice-to-Have:** With **11 official languages** and a richly diverse consumer base, South African enterprises must deliver AI solutions that understand and communicate in local languages. An English-only chatbot or model misses the majority of the market and risks perpetuating digital exclusion. AI systems must be culturally aware and linguistically competent. This is also a competitive opportunity: organisations investing in multilingual GenAI (for example, customer service chatbots fluent in isiZulu, Xhosa, Afrikaans, etc.) can tap underserved segments and build trust with a broader customer set. In short, a **truly effective GenAI strategy in South Africa "speaks the languages" of its people**.

- **Customer Experience and Software Engineering Lead: Early Value -** The fastest, safest routes to GenAI ROI lie in **customer service and software development**. Multilingual conversational agents (for customer support and sales), intelligent search and knowledge management tools, and AI coding assistants are showing quick wins. These use cases tend to have clear metrics (e.g. call deflection rates, resolution times, developer productivity). They can piggyback on existing platforms (like integrating GenAI into Office 365 or software IDEs), making them low-friction entry points. In asset-heavy sectors like mining and manufacturing, pairing GenAI with predictive maintenance and IoT sensor data is another high-impact play – reducing downtime and maintenance costs in ways that directly save money and preserve jobs.

- **Governance Must Be "POPIA-First" by Design:** South Africa's privacy law, the Protection of Personal Information Act, provides a ready-made blueprint for AI governance. Its principles of **lawfulness, purpose limitation, data minimisation, transparency, security safeguards, and data subject rights** align directly with controlling GenAI. Yet currently, only ~13% of organisations have any AI governance in place. Instituting a robust **AI usage policy** is the quickest, highest-impact step to rein in shadow AI. Such a policy should prohibit feeding confidential data into public tools, mandate human review of AI outputs, enforce bias and fairness checks, and assign clear accountability for AI use. Leaders who embed **POPIA compliance and ethical guardrails**

into their AI deployments not only avoid legal pitfalls but can also turn compliance into a competitive advantage, assuring customers and partners that their AI is trustworthy and secure.

- **Economics Are Under Scrutiny – ROI Discipline is Rising:** After the initial hype, CFOs and CEOs are now pressing hard for GenAI initiatives to demonstrate tangible value. Many early projects with weak business cases or unclear returns are being **paused, rescoped, or cancelled**. GenAI is no longer a blank-check experimental spend – it competes for budget like any other investment. As a result, organisations are focusing on **cost optimisation and credible ROI measurement**. They are right-sizing models (often using smaller models that are sufficient for the task to save on infrastructure costs), negotiating cloud costs aggressively, and establishing clear KPIs (e.g., "the AI chatbot must deflect 20% of calls, saving X rands"). Those that **cannot demonstrate a meaningful impact on revenue, cost, or customer metrics within a reasonable timeframe will see funding dry up**. The winners in 2025–2026 will be the companies that combine visionary experimentation with rigorous financial discipline, continually iterating until GenAI delivers repeatable business value.

- **Execution Excellence Will Define the Leaders:** As GenAI adoption becomes widespread, the differentiator will not be who has access to the smartest model – it will be **how well organisations execute**. Leadership commitment and cross-functional coordination are paramount. The hallmarks of leading enterprises are emerging: **visible executive sponsorship (often CEO-level), a cross-department AI governance committee, a sanctioned enterprise AI platform or toolset, strong data management practices, a portfolio of 3–5 high-impact use cases with baseline metrics, and an internal talent program to build AI fluency and trust**. In contrast, laggards will be those stuck in endless pilots, lacking governance, struggling with data issues, and facing pushback from employees or customers. By late 2026, it will be evident which companies have treated GenAI as a strategic, managed program versus those that have approached it ad hoc, and the performance gap will widen in areas such as customer satisfaction and cost efficiency.

Strategic Assumptions (2025–2028)

These forward-looking predictions serve as "guardrails" for planning GenAI initiatives and anticipating risks within the South African enterprise context.

- **Agentic AI Hype Meets Reality:** *By late 2026,* approximately **40% of enterprise software products** used in South Africa will incorporate some form of AI **agents** or autonomous assistants to automate at least one task or process step. However, **at least 30–40% of early agentic AI projects** (launched between 2025 and 2027) will stall, be **rescoped, or cancelled** before full deployment. The leading causes will be unclear ROI, insufficient data readiness to support the agent, and integration difficulties with legacy systems or business processes.

Interpretation: Many organisations will enthusiastically pilot AI "agents" that can act on behalf of users (for example, an AI that autonomously resolves IT support tickets or manages aspects of the supply chain). But the complexity of making these agents reliable and safe in real business workflows will be underestimated. A significant share will fail to graduate from proof-of-concept due to high error rates, unexpected costs, or user mistrust. CIOs should temper expectations and implement strong oversight for agentic AI, focusing on narrow, well-governed use cases for agents first.

- **Domain-Specific Models Overtake General LLMs:** *By 2027,* more than **50% of GenAI models used in enterprises will be domain-specific or fine-tuned for specific industries/functions**, up from virtually none in 2023. In other words, specialised AI models (trained on a company's or industry's own data and terminology) will become more common in production than generic large-scale models. The **total cost of ownership (TCO)** for these tailored models will be **30–60% lower** (for a given task) than using a giant one-size-fits-all model. Yet, they will deliver equal or superior accuracy on targeted tasks.

  Interpretation: We will see a shift from reliance on a few big-name LLMs to a proliferation of smaller, purpose-built models – e.g., a legal contract analysis model, a mining operations model, and a clinical health model. These smaller models are cheaper to run and easier to govern. Gartner forecasts that organisations will use **three times smaller**/task-specific models than large general models by 2027. This trend will be enabled by the growing ecosystem of base models and tooling that let companies train or customise AI on their own data. Companies should plan for a hybrid AI architecture where different use cases require different models – the era of a single, monolithic AI brain for everything will fade.

- **Synthetic Data Becomes Standard Practice:** *Through 2026,* the generation of **synthetic data** (artificially created datasets) will go from a niche experiment to a **mainstream tool** in enterprise AI development, particularly in regulated sectors. By 2026, **around 75% of large enterprises will be using GenAI to produce synthetic data for one or more use cases** – a jump from under 5% doing so in 2023. Synthetic data will be employed to augment scarce training data, protect privacy (by using artificial data in place of real sensitive records), and to debias models.

  Interpretation: As organisations face data shortages and privacy constraints (especially under laws like POPIA), they will increasingly "make" data to feed AI models. For example, a bank might generate thousands of synthetic customer profiles to test an AI model without exposing real customer info, or a healthcare provider might create balanced synthetic datasets to reduce bias. This practice will become a standard part of the AI toolkit. Leaders should invest in capabilities to generate and evaluate the quality of synthetic data, and expect new roles or services to emerge around managing synthetic data pipelines.

- **Multimodal AI Becomes the Norm:** *By 2027–2028,* roughly **40% of new GenAI solutions** deployed in enterprises will be **multimodal**, meaning they

can simultaneously handle text, voice, images, video, or other data formats in a unified model. (This is up from essentially 0–1% of AI solutions being multimodal in 2023.) These multimodal systems will power use cases such as rich media customer service (combining voice and text understanding with image recognition), automated video content generation, and complex field maintenance assistants (an AI that can "see" equipment via a camera and discuss it), among others.

Interpretation: The first wave of GenAI was mostly text-based (chatbots, text generators), but the next wave integrates multiple data types. For instance, a single AI assistant might take a spoken question from a technician, analyse an equipment diagram image, and output a text report and an annotated image – all in one workflow. This will significantly expand AI's utility in areas such as technical support, design, and education. Companies should watch for vendors offering multimodal AI capabilities and be ready to pilot them where they can add value (e.g. an insurance appraiser AI that can both converse and examine photos of damage).

- **Value Proof Will Be Mandatory for Funding:** *From 2025 through 2026,* CFO oversight will tighten, resulting in **approximately 25% of planned AI project spending** in enterprises being **delayed, reduced, or reallocated** to 2027 unless those projects can demonstrate clear, measurable progress on key performance indicators (KPIs). In practice, AI programs that implement dual-horizon value tracking – e.g., reporting **short-term "quick win" metrics every quarter, along with projected 12– to 18–month ROI** – will be more likely to retain full funding. Those without evidence of early impact will see budget scrutiny and potentially cuts or pauses.

  Interpretation: The message from finance chiefs is "prove it or lose it." Unlike the 2023 spending rush fuelled by fear of missing out, the 2025 approach is far more results-driven. A recent global survey indicated that fewer than one-third of companies can directly tie AI initiatives to profit-and-loss improvements today, which is causing scepticism. As a result, AI leaders must institute rigorous benefit tracking. For example, if an AI tool is designed to reduce call centre volume, it needs to demonstrate the reduction and associated cost savings; if an AI system is intended to accelerate R&D cycles, it needs to show how much time is saved. Organisations that build credibility through transparent metrics will continue to secure investment, whereas those that cannot will find AI budgets under the knife of cost optimisation.

- **AI Governance Will Be Non-Negotiable by 2026:** *By 2026,* having a robust AI governance framework will be a **prerequisite for doing business** in many sectors. We anticipate that a large majority of enterprises (certainly all in heavily regulated industries) will have formalised **AI policies, risk assessment procedures, and human-in-the-loop requirements** in place. Regulators and auditors, both local (e.g., POPIA authorities) and international (e.g., clients from jurisdictions with AI regulations), will increasingly demand evidence of effective governance. Specifically, by 2026, any substantial AI deployment will require documented **model fact sheets** (detailing training data, intended use, and limitations), bias mitigation processes, and automated controls such as data loss prevention for AI tools.

Interpretation: The grace period for "move fast and break things" in AI is ending. Much like cybersecurity and data protection have become board-level issues with required controls, AI ethics and governance are heading in the same direction. International developments, such as the EU's AI Act, are raising the bar, and South African companies will need to demonstrate compliance if they want to partner globally. We expect to see roles such as **Head of AI Governance** emerge, along with cross-functional committees that oversee AI initiatives. Companies that lag on governance could face not only regulatory penalties but also exclusion from deals or supply chains due to trust concerns. On the other hand, those who are proactive can turn their strong governance into a selling point ("AI you can trust").

Global and Local Context

In plotting a GenAI strategy, South African CIOs and CTOs must navigate both global trends and local realities. The global context sets the overall technology and market landscape, while the local context shapes how these innovations unfold, considering South Africa's unique opportunities and challenges. Below, we examine each in turn.

The Global Landscape: From Hype to Pragmatism

Position on the Hype Cycle: Globally, generative AI has now passed the "peak of inflated expectations" and is descending into the **"trough of disillusionment"** in Gartner's hype cycle terminology. This is a natural maturation phase. Following the sensational breakthroughs and extensive press coverage in 2023, enterprises in 2024 and 2025 are grappling with the real-world implementation challenges. The exuberance ("GenAI will change everything overnight!") is being tempered by more sober questions around ROI, scalability, and risk. This shift does *not* imply that GenAI is a bust – rather, it signals that businesses are moving beyond fascination with the technology itself and focusing on delivering **tangible value**. Executive sentiment has cooled healthily: the question is no longer *"What magical things can ChatGPT do?"* but **"What specific business problem can this technology solve, profitably and securely?"**

Evidence of this reality check is evident in global surveys. In late 2024, Gartner found that while enterprise spending on GenAI projects was climbing, fewer than **30% of AI leaders' CEOs were satisfied** with the returns on those investments so far. Many CEOs felt outcomes did not yet justify the hype or expenditure. The reasons varied: some organisations chased novelty use cases that didn't move the needle, while others lacked the talent and infrastructure to scale pilots into production. This emerging disappointment at the C-suite level has a direct effect: it often creates a strategic vacuum (as leadership pauses to reconsider AI strategy), which ironically can fuel unsanctioned bottom-up experimentation as impatient employees seek value on their own. Indeed, the global hesitancy to commit fully until ROI is proven is a major contributor to the "shadow AI" phenomenon seen in South Africa – a point we return to in the local context.

Pivot to Foundations: A clear global trend is the pivot from **model-centric innovation to foundational capability-building**. In 2023, the discussion was dominated by which model had the most parameters or the best raw performance on benchmarks. By 2025, savvy organisations will realise that *any* model is only as good as the ecosystem enabling it. Thus, the attention has shifted to what might be called the **"boring plumbing"**: data quality, architecture, integration, and operational discipline. Three foundational enablers stand out:

- **AI Engineering and ModelOps:** Companies are investing in the engineering discipline of AI – establishing standardised processes and tools for developing, validating, deploying, and monitoring models. **ModelOps** (Model Operations) teams or functions are being established to manage the AI lifecycle, much like DevOps manages software development. This includes version control for models, bias and performance testing, monitoring models in production for drift or anomalies, and maintaining audit trails for how models were trained and with what data. Without these practices, even the best model can't be trusted or efficiently scaled. Enterprises now recognise that *operationalising* AI is harder than building a proof-of-concept, and they are budgeting and staffing accordingly.

- **"AI-Ready" Data: Data has emerged as the single most critical enabler—and bottleneck - for GenAI.** A startling proportion of organisations report that their internal data is not prepared for AI use (surveys indicate an average of ~57%). "AI-ready" means data that is clean, well-labelled or contextualised, accessible in the right formats, and compliant with usage policies. GenAI models, especially large ones, are extremely hungry for high-quality data. If your CRM is full of typos, duplications, and gaps, or your knowledge base is disorganised, a GenAI deployment will struggle or produce garbage. Globally, there's a realisation that **a successful GenAI strategy is first and foremost a data strategy**. This has sparked a renewed emphasis on data governance programs, data integration efforts, and innovative solutions, such as data synthesis and augmentation (as mentioned, generating synthetic data to fill gaps or overcome privacy limitations). The bottom line: forward-looking CIOs are now directing significant effort to get their data house in order – treating data as a core asset that can either unlock AI value or completely undermine it.

- **Integration and Architecture:** The world's best AI model won't create value sitting in a vacuum – it needs to be *embedded into business processes and systems*. Thus, integration is crucial. Globally, we are witnessing trends such as the emergence of **AI middleware and API orchestration layers** that facilitate the connection of GenAI services to existing IT landscapes. Security and identity management are part of this. For instance, ensuring that an AI assistant integrated into a workflow only accesses data that a user is authorised to view, or that API calls from an AI agent comply with transaction limits and logging requirements. Enterprises are also adopting more **composable architectures** to handle the rapidly evolving AI tool landscape – they favour systems where AI components (models, vector databases, etc.) can be easily swapped in and out without major overhauls. This helps avoid vendor lock-in and allows for upgrading to new AI technology as it becomes available. Gartner predicts that by 2026, companies using **composable**

architectures will implement new AI features 80% faster than their competitors who are still stuck in monolithic setups. In summary, globally, IT leaders are re-architecting environments to be **"AI-plug-and-play"** – flexible, API-driven, and secure – so that GenAI capabilities can be integrated quickly and safely into enterprise workflows.

Emerging AI Frontiers: Even as companies double down on foundations, the next generation of AI tech is rapidly advancing. Two major innovation vectors are dominating global discussions:

1. **Agentic AI Systems:** These are AI systems endowed with a degree of autonomy – they don't just respond to one-off prompts, they can **take goals and carry out multi-step tasks**. For example, instead of a simple Q&A bot, imagine an AI that you can instruct: "Handle all the routine software deployment steps for this new application," and it plans and executes the sequence (checking code, merging to the repository, running tests, deploying to servers), only alerting a human if something goes wrong. Early examples include AI agents that can browse websites and perform actions (such as making reservations, ordering products, or scheduling meetings) based on a single high-level request. Tech firms have begun releasing frameworks for such agents, and open-source projects are exploring "AutoGPT" and similar concepts. Globally, there's tremendous excitement – and caution – around agentic AI. If successful, it could automate whole swathes of knowledge work. But it also raises new risks (as autonomous actions can amplify errors quickly). Expect agentic capabilities to improve steadily, and more enterprise software vendors to offer "autonomous mode" options in the coming 1–2 years. However, widespread enterprise adoption will depend on proving these agents are reliable and controllable, which ties back to having human oversight and strong governance.

2. **Domain-Specific & Small Models:** As noted in the assumptions, a strong movement is toward **smaller, specialised models** – sometimes called **small language models (SLMs)** or domain-specific models. The initial wave of GenAI was led by giant general models (with hundreds of billions of parameters) trained on broad internet data. Now, we're seeing the rise of models with, say, 1 billion or 10 billion parameters, trained or fine-tuned on highly targeted datasets (such as legal contracts, medical texts, and code repositories). They are often just as effective within their domain and far more efficient to run. Open-source advancements have made it feasible for even individual companies or small communities to develop such models. For example, in 2024, Meta's **Llama 2** open-source model allowed organisations to start with a powerful base model and fine-tune it themselves. By 2025, we will have models like Meta's rumoured Llama 3, Microsoft's "Phi" series, or domain-specific variants like BloombergGPT for finance. Gartner and other analysts forecast that enterprises will increasingly favour these targeted models – by 2027, over half of enterprise AI deployments may use them. The benefits include not only cost savings, but also potentially **better accuracy** on domain-specific tasks and easier compliance (since a model trained on a company's own data can be more transparent and controllable). The global best practice emerging is to maintain a portfolio of models: use broad, general models for tasks requiring extensive knowledge, but use smaller, tailored

models for core business tasks where efficiency, consistency, and domain depth are needed.

Global Market Maturation: The global GenAI market is thus simultaneously expanding and "growing up." Adoption statistics indicate a rapid uptake: by the end of 2024, nearly **80% of companies worldwide will have at least piloted GenAI**, according to McKinsey and others. In many enterprises, GenAI has moved from curiosity to at least some form of usage in departments like customer support, marketing content, or R&D. Yet, on the outcome side, there's a tempered view – only about 15–20% of companies claim significant measurable gains (revenue uptick, cost savings) from GenAI so far. This disparity is what's driving the intense focus on **ROI and productivity metrics** going forward.

Another global facet is **consolidation and competition** in the vendor ecosystem. In 2023, we saw a flood of AI startups offering various point solutions. By 2025, larger players will be incorporating GenAI features into their platforms (e.g., Microsoft's Copilot across Office apps, Google's Duet AI in Google Workspace, Salesforce's Einstein GPT for CRM, etc.). This means that even if a company isn't developing AI itself, it will encounter it through software updates. Every major enterprise software category – including ERP, CRM, HR, IT service management, and analytics – is being embedded with GenAI capabilities. Enterprises globally are grappling with how to **enable or govern these "built-in" AI features**. For example, if Microsoft Word now has an AI that can summarise documents, how do we ensure it doesn't send the document content out to an external cloud without permission? Expect to see more demand for enterprise controls (Microsoft, for instance, offers an option to run Copilot in a privacy-compliant, isolated manner for big customers). Companies may also pursue **multi-cloud and multi-vendor strategies** to avoid overreliance on a single AI provider. For instance, using one cloud for general compute but another specialised AI cloud for training models if it's more cost-effective, or exploring on-premise AI appliances for sensitive workloads. Globally, the ability to mix and match, without being locked in, is viewed as a strategic advantage in this rapidly evolving market.

In summary, the global context is one of GenAI becoming **ubiquitous but also normalised**. It's everywhere (in apps, in employees' hands), but it's expected to prove itself. The emphasis is on integrating AI deeply and responsibly into how work is done, rather than treating it as a magic solution on the side. The next 12–24 months globally will be about separating winners and losers: not in terms of any single AI product, but in terms of which enterprises can *operationalise and govern* AI to actually drive value. These trends and lessons heavily inform what should be done in South Africa, which we turn to next.

The South African Reality: Opportunities and Challenges

Against the backdrop of global trends, South Africa's GenAI journey has distinctive characteristics. The local landscape is characterised by a mix of enthusiastic adoption, structural gaps, and contextual imperatives that differ from those in other markets. South African enterprises cannot simply copy and paste a Silicon Valley AI strategy; they must tailor their approaches to our domestic environment of skills

distribution, regulatory obligations, and socio-economic priorities. Several key aspects define the South African GenAI context in 2025:

Enthusiastic Adoption Amid a Strategy Void: Perhaps the most striking feature is the **mismatch between high adoption on the ground and a lack of formal strategy at the top.** A major study (the "SA GenAI Roadmap 2025") found that **only 14% of South African organisations have a company-wide AI strategy**, meaning 86% do not. Yet, paradoxically, various surveys report that a majority of mid-to-large firms have employees or teams using GenAI tools in some fashion. For example, a 2025 survey of large enterprises showed that GenAI usage had climbed from 45% of firms in 2024 to **67% in 2025** – making GenAI perhaps the fastest-spreading new technology in recent corporate history.

What are employees doing with it? Common examples include using tools like ChatGPT or Bard to draft emails and documents, generate ideas or first drafts for marketing content, assist with coding or data analysis, and answer domain questions by querying large language models. In many cases, these are individual or departmental initiatives rather than officially sanctioned projects.

This **"shadow AI"** environment – high usage with low oversight – presents a double-edged sword. On one hand, it demonstrates a grassroots recognition of GenAI's productivity benefits. South African professionals are proactively finding ways to work smarter, which is encouraging for a culture of innovation.

On the other hand, it introduces **significant risks**: data privacy breaches (workers might inadvertently feed confidential data into public AI services), regulatory non-compliance (violating POPIA by sharing personal data with an overseas AI, for instance), security concerns (pasting proprietary code or plans into an unknown tool), and inconsistent results (different teams getting different answers with no organizational quality control).

For local CIOs, this pattern is both a **warning and an opportunity**. The warning is clear: **the lack of a strategic framework is unsustainable** – it's exposing firms to legal and operational landmines daily. The opportunity lies in the fact that employees are already invested in the potential of GenAI, providing a great foundation to build upon. The task now is to **channel this enthusiasm into a governed strategy**. Rather than stifle the bottom-up energy (which could harm morale and competitiveness), leaders should formalise it: establish enterprise-approved GenAI platforms that employees can use safely, provide training on how to use AI tools responsibly, and incorporate the best of those grassroots use cases into official workflows. South Africa's biggest immediate GenAI win could come from simply taking what people are already doing informally and scaling it under proper management.

Infrastructure Leap and Data Sovereignty: Historically, many African enterprises lagged in advanced tech adoption partly due to infrastructure constraints – unreliable connectivity, no local cloud data centres (meaning high latency and no guarantee of data residency), etc. This barrier is rapidly shrinking for South Africa. In **early 2024, Google Cloud launched its first African cloud region in Johannesburg**, and Microsoft and Amazon Web Services have each expanded their presence (Microsoft had opened Azure regions in Cape Town and Johannesburg a few years prior and announced a massive new investment in 2025 to increase capacity and build

dedicated AI infrastructure). These moves are **game-changers**. They bring world-class cloud computing and AI model hosting literally onshore, addressing multiple concerns:

- **Latency:** With data centres in Johannesburg/Cape Town, the physical distance between users and servers is small, improving response times for AI services. Real-time AI applications (fraud detection, conversational agents, IoT analytics) benefit enormously from this reduced latency.
- **Data Sovereignty and Compliance:** Many industries, including banking, insurance, healthcare, and government, have requirements or strong preferences to keep certain data within national borders. Before, using advanced AI services often meant that data might travel to Europe or the US. Now, with local regions, companies can ensure sensitive data and AI processing remain in South Africa, aiding POPIA compliance and easing client concerns. For instance, a bank can utilise cloud AI services without customer data leaving the country, provided the configuration is set correctly.
- **Access and Cost:** Local presence of hyperscalers tends to improve bandwidth and reduce some costs (no international bandwidth premiums). It also means local companies can leverage cutting-edge AI offerings (like custom model training services, AI-optimised hardware like GPUs/TPUs) without needing to invest in their own data centres. Additionally, initiatives like the planned **Cassava/NVIDIA AI compute hub** indicate that dedicated high-performance AI infrastructure will be available for hire, which can accelerate complex model development in-country.

All of this means that, as of 2025, **the playing field is levelling** – South African firms have nearly the same technological "canvas" to paint on as their counterparts in more developed markets. The locus of challenges shifts: lack of cloud or compute is no longer a valid excuse for not pursuing GenAI. Instead, companies must focus on **using the infrastructure wisely**. We anticipate a trend of more enterprises migrating AI workloads to these local clouds and possibly exploring hybrid setups (keeping some critical workloads on-premise and others in the cloud). The primary constraints will now be the **soft infrastructure** (skills, policies, integration) rather than the hard infrastructure. As one local tech executive put it, "we've solved the bandwidth and server problem – now it's on us to solve the strategy and talent problem."

Socio-Economic Responsibility – Augmentation Over Automation: South Africa's socio-economic context must shape its AI adoption strategy in a way that might not be as pronounced in other regions. We have one of the highest unemployment rates in the world (hovering around **32%** in recent years). This reality creates a powerful imperative: **business and government leaders alike are explicitly prioritising AI deployments that** augment human workers rather than replace them. There is understandable public sensitivity (and even political concern) that automation could further exacerbate unemployment and inequality if done recklessly.

In practice, this means South African enterprises are being cautious about framing AI as a cost-cutting, labour-reducing tactic. Instead, the emphasis is on how AI can **enhance productivity, create new roles, and free staff from mundane tasks,**

allowing them to focus on more value-added activities. For example, a mining company might use AI to automate equipment monitoring and repetitive safety checks – not to eliminate those jobs, but to redeploy those workers to higher-order safety analysis and maintenance planning, thereby improving overall safety while retraining employees for more skilled roles. Banks here talk about AI giving their service agents better information and tools (so each agent can handle more clients effectively) rather than simply halving the call centre staff.

This approach is more than just altruism; it aligns with a sustainable business strategy. If AI projects were rolled out that led to significant layoffs, companies could face backlash, strikes, or damage to their brand in the South African context. Moreover, neglecting the human element often leads to AI project failures (people refusing to adopt the tools, or reintroducing manual workarounds). By contrast, an augmentation-focused strategy naturally includes strong **change management and upskilling**, which increases the chances of AI success. It's encouraging to see that this narrative is being echoed at the national policy level. The Department of Higher Education's partnership with Microsoft, signed in October 2025, specifically aims to **train students in AI engineering and digital skills** through Technical and Vocational Education and Training (TVET) colleges. Such public-private initiatives are geared toward making sure the workforce is AI-ready and that the benefits of AI (in productivity and new services) don't leave the population behind.

For South African CIOs and CTOs, the takeaway is clear: any GenAI deployment plan should include a **workforce impact assessment and plan**. Identify how roles will change, what tasks will be automated, and how affected employees will be trained for new responsibilities. Set a tone that AI is there to **assist your staff, not replace them**. Some organisations have even created joint IT - HR working groups on "future of work" to oversee this integration of AI and people. Those who successfully navigate this will not only avoid negative repercussions but likely unlock the full innovative potential of combining human judgment with AI efficiency.

The Multilingual Mandate: South Africa's linguistic diversity is a defining aspect of its market. English may be the primary business language, but the majority of the population communicates in a variety of home languages (isiZulu, isiXhosa, Afrikaans, Sepedi, Setswana, and others). A **GenAI strategy that ignores multilingual support is fundamentally incomplete in South Africa.** Suppose an insurance company deploys a customer-facing AI chatbot that only understands and responds in English. In that case, it is effectively shutting out or underserving a significant segment of its customers who would prefer to interact in languages such as isiZulu or Sesotho. The same applies to AI in public services – a single-language AI could exacerbate digital inequities.

From a business perspective, providing AI services in multiple languages can be a source of competitive advantage and market growth. It allows companies to reach new customer segments in their preferred language, improving trust and satisfaction. It can also improve the quality of insights – for instance, analysing customer feedback in all languages, not just English, could reveal needs and trends that a monolingual approach would miss.

There are challenges: training models to be fluent in low-resource languages (where digital text data is scarce) is harder. Off-the-shelf global models might not handle, say, isiZulu well because they were never trained on enough content in that

language. But this is a solvable problem with the right investment. We're already seeing local innovation here. Notably, startups like **Bolthale AI and PRAAT** have emerged, focusing on natural language processing for African languages. Partnerships between universities, government, and industry are also being considered to create open datasets and models for South African languages. Globally, large players are expanding language support (for example, by 2025, Google's and OpenAI's models are expected to have significantly increased their non-English understanding); however, fine-tuning them for local dialects and contexts still requires local effort.

For enterprise tech leaders, the mandate is to **ensure that your vendor or solution selection process takes into account multilingual capabilities.** When evaluating an AI platform, ask: Does it support the languages my customers and employees speak? Can it be adapted to do so? During RFPs (Requests for Proposals), include some tests or criteria related to language. If one vendor's model understands Xhosa queries and another doesn't, that should weigh heavily. Additionally, consider collecting and using your own multilingual data to fine-tune models. For example, a bank could use its trove of call centre transcripts in multiple languages to train a custom language model that's far more accurate for local use than a generic one.

In practical terms, we are likely to see South African firms deploying **multilingual chatbots** on popular channels like WhatsApp (which has extremely high usage locally). In fact, a telling example: **Capitec Bank** leveraged a conversational AI on WhatsApp that handles millions of customer conversations a year, many in local languages. They achieved an 78% customer satisfaction score and doubled the efficiency of their contact centre agents by automating routine interactions across languages. This kind of success underscores that meeting customers where they are – in terms of platform *and* language – is key to GenAI delivering value here.

Sectoral Variation and Local Use Case Focus: The composition of the South African market influences which GenAI use cases are most pressing. Our economy has strong mining, manufacturing, and financial services sectors, a significant retail sector, and a sizable public sector, among others. Each of these has unique pain points that GenAI can address, and local early adopters are gravitating to those opportunities.

Broadly, **customer experience (CX)** is a universal priority, and multilingual GenAI for customer service is experiencing rapid uptake (as mentioned, financial institutions are deploying chatbots on messaging apps, telecoms are using AI agents for support, etc.). **Software development and IT operations** also see wide interest in AI assistants (given the global shortage of senior developers, AI that helps write code or manage IT tickets is valuable everywhere, including here).

Beyond that, in **mining and manufacturing**, there's a trend to combine GenAI with **IoT and predictive analytics** – for example, using GenAI-driven analysis of sensor data to predict equipment failures, or employing computer vision AI to monitor safety conditions in mines. These industries are looking for productivity and safety improvements that directly impact the bottom line (reducing downtime, preventing accidents). A local mining trial is utilising AI systems to monitor vehicle emissions and predict machinery breakdowns before they occur, which helps avoid costly unplanned stoppages and supports environmental compliance.

In **retail**, personalisation and supply chain efficiency are key. South African retailers are exploring AI-generated marketing content (to tailor promotions to different customer segments) and using AI for demand forecasting and logistics optimisation. An interesting stat: the market for AI in retail in SA is projected to grow almost tenfold (from around $31 million in 2023 to $282 million by 2032), reflecting anticipated investment in these areas. Early use cases include route optimisation in delivery networks – crucial in a large country where fuel and transport are major cost drivers. AI route planning tools have demonstrated up to **35% fuel savings and 40% reductions in transport costs**, which is very attractive to retail and distribution companies facing thin margins.

The **public sector** here, while often slower in adopting technology, is also waking up to the potential of GenAI. Government agencies are looking at AI assistants for citizen services (imagine a multilingual virtual agent handling queries about government services 24/7) and AI for internal efficiency (analysing large datasets in healthcare or education). There's also a national interest in how AI can help address social issues, such as education – for instance, AI tutoring systems that can assist students in multiple languages.

Crucially, South African enterprises are recognising that the **highest-impact GenAI use cases will often be those that directly tackle local challenges or leverage local strengths.** This might mean focusing AI on areas that improve financial inclusion (like a bank using AI to help people with limited literacy navigate loan options in their home language), or on cost-saving in essential industries (like energy management AI in our power sector context). There's a pragmatic streak: "What will make a real difference in our operating environment in the next 12–18 months?" rather than just doing AI for AI's sake.

To summarise the local context: South Africa enters the GenAI era with **strong adoption momentum but a need for strategic coherence**. The removal of infrastructure barriers and the urgency of socio-economic considerations create a clear call to action for leaders. Suppose we can marry the global best practices (data readiness, governance) with local priorities (augmentation, multilingual reach, and specific sector focuses). In that case, South African enterprises can leapfrog from sporadic AI experiments to **sustainable, scaled AI-enabled transformation**. The following sections of this report aim to provide exactly that playbook – addressing use case selection, operating models, economics, sector strategies, risk management, and more – to help convert the promise of GenAI into practical results in the South African enterprise over the 2025–2028 horizon.

High-Impact Use Cases for the South African Enterprise

Having set the strategic context, we now turn to the practical question: **Where should South African organisations apply generative AI for the most value?** In the next 12 to 18 months, certain use cases stand out as high-impact, relatively low-risk entry points. These are opportunities that align with both global trends and local needs – delivering tangible improvements in customer experience, operational efficiency, or employee productivity, while leveraging the new infrastructure and mitigating common concerns (like language and data sensitivity).

Below, we outline several key use case domains, each accompanied by real examples or early results in South Africa. For each domain, we describe the applications, the benefits, and local considerations to ensure success.

Transforming Customer Experience & Service

Why this is a priority: In competitive markets such as banking, telecom, and retail, customer experience is a significant differentiator. South African consumers are increasingly digital and expect quick, personalised service. Generative AI – especially conversational AI – offers a way to enhance customer service capacity and quality significantly. Moreover, solving customer queries in all local languages at scale is an area where GenAI can excel, directly supporting financial inclusion and customer satisfaction goals.

Applications: The primary use case is deploying **sophisticated, multilingual chatbots and voice bots** for customer service and sales support. These GenAI-powered agents go far beyond the basic scripted FAQ bots of the past. They can handle complex requests, retrieve information from various backend systems (such as account information and product details), and guide customers through procedures (like applying for a service, troubleshooting a device, or understanding a policy) in a conversational manner. Importantly, they can do so in the customer's preferred language. Advances in large language models mean these bots can understand nuance and generate natural-sounding responses.

For example, a customer might ask, in Sesotho, "I lost my bank card, what do I do?" The AI assistant can respond immediately with the steps to freeze the card and request a new one, phrased in polite, clear Sesotho – or even guide the user through the process interactively. Globally, such GenAI customer service implementations have shown they can **reduce support costs by 30–45%** (by deflecting calls and handling routine inquiries) while **improving response times** dramatically (since the AI is instant and available 24/7).

Local Example – Banking (Capitec): A standout case is **Capitec Bank's WhatsApp-based AI chatbot.** Capitec, one of SA's leading retail banks, integrated a conversational AI to engage with customers on WhatsApp, which is used by the vast majority of mobile phone users here. The results illustrate the impact: this AI system now handles about **15 million customer conversations per year**, with roughly **7 million of those interactions happening through WhatsApp**. Customers can check balances, get information on services, and resolve common issues via a simple message – no need to wait on hold for a call centre. Despite automating these interactions, the experience has been positive: Capitec reports a **78% customer satisfaction (CSAT) score** for the AI-driven service.

Additionally, because the AI handles the easy questions, human agents in the contact centre are freed to focus on complex cases, effectively **doubling the efficiency** (each agent can now manage many more customers with the AI picking up the slack). This hybrid model of AI + human service is scaling customer support without a proportional increase in headcount. The key factors in Capitec's success have been meeting customers on *their preferred channel* (WhatsApp) and supporting *conversational language*, making the service feel natural and accessible.

Considerations for success: When implementing GenAI in customer service, companies should ensure robust **integration** with live agent systems for smooth handoff (if the AI cannot handle an issue or detects an upset customer, it should seamlessly escalate to a human with context). **Training and fine-tuning** the AI on local intents, slang, and product knowledge is crucial – out-of-the-box models will not recognise local banking terms or recent products until they are fed that data. Companies must also put guardrails so the AI doesn't stray from known answers (to avoid hallucinations that could mislead customers). Lastly, **multi-language support** should not be an afterthought – it requires gathering diverse training data and possibly leveraging translation pipelines or multilingual models. But the payoff, as seen with Capitec, is a much broader reach and higher customer engagement.

Driving Operational Efficiency & Automation

Why this is a priority: Across industries, numerous routine, time-consuming processes consume employee time or delay outcomes. Generative AI can automate or accelerate a range of such tasks, from generating reports to assisting in software development. In heavy industries (manufacturing, mining, utilities), even marginal efficiency gains can translate to big cost savings or output increases. In office environments, automating administrative drudgery can free staff for more strategic work. The South African context adds another angle: efficiency improvements can help companies remain competitive and retain jobs by reducing costs without laying off workers, a way to achieve "lean" operations in a humane manner.

Applications: On one end, GenAI excels at **content generation and summarisation**, which applies to tasks such as writing routine reports, drafting standard communications (emails, meeting notes), creating first drafts of policy documents or marketing copy, etc. A lot of back-office work (HR letters, finance quarterly summaries, internal newsletters) can be sped up with AI-generated drafts that humans then review. Organisations are already using tools integrated in Office suites for this – e.g., an AI to auto-generate a project status report from bullet points, or summarise a 10-page memo into a one-page brief for quick reading.

In the realm of **software development and IT operations**, the impact is particularly large. AI **code assistants** (like GitHub Copilot, AWS CodeWhisperer, or others) are being used by developers to auto-complete code, suggest solutions, and even generate entire functions or test cases. These tools have been reported to make developers **50% faster on average tasks and up to 90% faster on repetitive boilerplate coding**. In practice, this could mean that a task that used to take a week might be completed in a few days, or a bug that a developer might spend hours resolving can be resolved in minutes with an AI hint. For IT support, GenAI can draft scripts or queries to quickly diagnose issues or serve as a knowledge assistant to IT staff by pulling answers from documentation.

Moving to heavy industry, **predictive maintenance** is a prime example of GenAI-related applications. By analysing large volumes of sensor data from machinery (such as pressure, temperature, and vibration readings), AI models can predict when equipment is likely to fail or require servicing. Traditional predictive analytics exist, but GenAI can help by interpreting unstructured data (such as maintenance logs or technician notes) alongside sensor data to provide richer insights. Also, GenAI can

provide explanations and recommended actions in natural language ("The vibration pattern on Pump 4 suggests it will fail in ~10 days; schedule maintenance and check part X"). South African **mining companies** are piloting such approaches. In one local mine, AI systems monitor the real-time telemetry of haul trucks and drilling equipment in real-time. **Early results show a significant reduction in unplanned downtime** – meaning fewer sudden breakdowns – which directly translates to maintaining production output and protecting revenue. Additionally, optimising maintenance schedules means safer operations (addressing issues before they become critical failures) and can help preserve jobs by avoiding the downtime that might lead to cost-cutting pressures.

Local Example – Mining: One of South Africa's largest mining firms recently introduced a GenAI-assisted maintenance system in a pilot project. They equipped vehicles and processing equipment with IoT sensors feeding data to an AI platform. The generative AI aspect involves analysing complex patterns and communicating findings, such as correlating an increase in engine vibration with a subtle change in acoustic sensor readings to predict a likely engine belt failure. The AI then generates an alert for maintenance teams, providing a plain-language explanation and even a suggested parts list for the repair. As a result, the mine has been able to **prevent several potential breakdowns**. Each averted breakdown saved many hours of halted operations and millions of rand. Moreover, by avoiding catastrophic equipment failures, they protect their workers from the hazards associated with those events and avoid emergency repair costs. This approach aligns with sustainability goals too – by maintaining equipment proactively, they extend asset life and prevent environmental incidents (like oil leaks or emissions spikes that could occur during a breakdown). The company is now looking to scale this throughout its operations, and other mining houses are taking note.

Key considerations for success in back-office automation include **employee training and acceptance**. Staff need to understand that these AI tools are designed to assist, not evaluate their work. Encourage them to use the AI for first drafts and menial tasks, but always to review the outputs. There should also be clear guidelines on what AI can or cannot be used for (especially regarding sensitive data in documents). In software teams, integrate AI assistants into the development environment and establish best practices (e.g., code review processes remain, regardless of whether the code originated from a human or an AI). Monitoring code quality and security is crucial, since an AI might introduce vulnerabilities if not overseen.

In industrial IoT use cases, ensure that you have **high-quality data** and that domain experts collaborate with data scientists to train models. GenAI predictions in maintenance should be validated with engineers to build trust. It's also important to have a feedback loop: when the AI makes a prediction, track if it was accurate and adjust accordingly (this also helps in proving ROI – keep metrics of downtime reduction, etc.). Additionally, operational AI systems should be implemented with fail-safes; if the AI fails or is uncertain, the default should be not to disrupt operations or to alert a human.

Enhancing Productivity & Employee Enablement

Why this is a priority: Knowledge workers – including analysts, consultants, HR officers, and lawyers – spend a significant portion of their day searching for information, compiling knowledge from various sources, and drafting communications. GenAI can serve as a powerful assistant to amplify individual productivity, helping employees do more in less time and with less frustration. Given South Africa's sometimes constrained skills pipeline (we have fewer experts in certain fields, making each skilled employee's time highly valuable), augmenting each professional's capabilities can have a significant impact. Moreover, deploying AI internally as a tool can also reduce the temptation for employees to go rogue with unsanctioned tools (addressing shadow AI). If you give them a great internal tool, they won't need to use random external ones.

Applications: One major application is **enterprise knowledge management and search.** Think of a GenAI-powered internal assistant that an employee can query in natural language to find information across the company's documents, intranet pages, emails, meeting transcripts, etc. For example, an employee could ask, "Has our company ever done a project with client X in the insurance sector? Summarise the outcome." The AI could search through proposal archives, project reports, and CRM notes to deliver a concise answer, complete with relevant details or documents linked. This goes beyond traditional keyword search by actually understanding context and synthesising answers. It's like having a knowledgeable colleague who's read everything in the company library.

Tools like **VectorMind Atlas** and **Google Workspace Duet** are among the first wave enabling this. They connect to your organisation's SharePoint, email, Teams/Chat logs, etc., and allow queries like "Summarise all the feedback from client Y in the last 3 months" or "Draft a PowerPoint slide deck outline based on our Q2 report". These AI "copilots" can also take on tasks such as generating draft responses to emails or creating initial spreadsheets from raw data. The benefit is breaking down information silos – employees no longer need to know exactly where information is stored or who to ask; the AI can retrieve and consolidate it. This accelerates decision-making and reduces duplicate work.

A concrete example is using an AI assistant to onboard new employees. Instead of reading dozens of documents, the new joiner can ask the AI questions about company policies or previous work to get up to speed faster. In complex fields like law or engineering, junior staff can query the AI for explanations of terms or historical case references, getting instant guidance that would otherwise require bothering a senior colleague.

One advantage in South Africa is the presence of **local implementation partners and expertise**. A company called **CohesionX** was notably selected in Microsoft's "Global AI Founders" partner program, meaning they are recognised for advanced capability in deploying AI solutions. Firms like these are helping SA businesses integrate AI securely and effectively. This indicates that we have access to the know-how needed to customise and fine-tune these productivity AIs for maximum impact (for example, ensuring they connect to the right knowledge bases and comply with local data policies).

Use Case Examples: Imagine the HR department using an AI assistant to instantly pull up “all policies related to maternity leave and summarise any changes in the last 5 years” when updating employee handbooks. This task could take days of manual digging, but is completed in minutes by AI. Or a consulting firm enabling their analysts to ask, “What did we recommend to the telecom client in 2021 about 5G strategy?” and get an answer with direct excerpts from that old report, saving hours of search. Another scenario: a customer support agent, mid-call, can quietly use an AI tool to find relevant troubleshooting steps or prior similar cases by querying internal documentation with natural language – thereby improving first-call resolution rates, as they can obtain answers on the fly.

Key considerations for success include data accessibility and security. To build an effective internal AI assistant, you need to connect it to potentially sensitive corpuses of data (contracts, internal memos, financial reports). Strong access controls must be in place: the AI should respect user permissions (i.e., not display employee content they shouldn’t see), and the data should not leave your environment in plain form. This typically requires either hosting the model internally or using cloud offerings that guarantee data isolation. IT teams must work closely with solution providers to configure connectors and index data sources properly, while InfoSec teams ensure compliance (for instance, filtering out personal data fields from being indexed if not needed).

User training is also important – employees should learn how to phrase queries effectively to the assistant and, critically, verify important outputs. The AI might occasionally err or hallucinate; a culture of “trust, but verify” should be cultivated. Perhaps even integrate a feedback mechanism: if the AI gives a wrong answer, users can flag it, and that feedback can help improve the system.

Finally, **change management**: introducing these tools may alter how people approach their work (e.g., writing drafts versus editing AI-generated drafts, or searching knowledge via AI versus manually browsing folders). Some may resist or feel threatened. Organisations should clearly position AI as a helper (not a monitor or a replacement) and provide reassurance that judgment and final decisions remain in human hands. Often, demonstrating quick wins (such as how much faster a task can be completed with AI assistance) and highlighting stories of employees who utilise it to excel can drive adoption.

Innovating in Marketing & Sales

Why this is a priority: Marketing and sales functions thrive on personalisation, creativity, and timely engagement – all areas where GenAI can turbocharge results. South African companies, whether B2C or B2B, are seeking to increase revenue and strengthen customer relationships. Generative AI provides tools to produce tailored content at scale and extract insights from customer data, ultimately improving conversion rates. Additionally, with the proliferation of digital channels in South Africa (social media usage is high, and e-commerce is growing), marketing teams must produce more content than ever, and sales teams must manage increasingly data-driven processes – GenAI is a natural ally in these endeavours.

Applications: In **marketing**, generative AI is being used to create content and campaign assets efficiently. This spans writing personalised email copy for thousands of customers, generating social media posts in the brand's tone, crafting product descriptions for e-commerce listings, and even producing visuals or basic video snippets with AI models (for example, using AI image generation to create variant ads or illustrations without a designer's constant involvement). The key advantage is **scale with personalisation** – AI can generate hundreds of versions of a message, each tailored to different customer segments or even individuals (based on their profile or past behaviour), something human marketers cannot do manually.

For instance, an online retailer can utilise GenAI to generate personalised product recommendations and descriptions in marketing emails for each customer, reflecting their browsing history or purchase history. Or a bank's marketing team might use AI to create ad copy variations: one version emphasising rewards to younger customers, another emphasising security to older customers, etc., and test which performs better.

In **sales**, particularly in B2B or high-touch B2C sales (like automotive or real estate), GenAI can assist with drafting proposals, analysing client data to suggest upsell opportunities, and even guiding sales conversations. AI can analyse CRM data to identify which leads are most promising or which customers may be ready for a new product based on subtle signals. It can also help salespeople by generating draft responses to client inquiries or summarising the outcomes of the last call with a client and suggesting next steps.

Another emerging sales use is in creating **dynamic pricing strategies** – AI models that adjust prices or offers in real-time based on supply, demand, and customer willingness-to-pay signals. While dynamic pricing algorithms have existed, GenAI can add a layer by generating strategies or explanations for pricing decisions that human managers can review (for instance, "our data suggests this customer segment is more price-sensitive, we should offer a 10% discount on bundle X to close deals").

Local Example – Retail and Logistics: Though case studies are still emerging in SA, the retail sector shows strong potential. The stats we cited earlier – SA's AI-in-retail market growing at ~27% CAGR – suggest many retailers are investing in AI solutions. A focal area is **supply chain and delivery optimisation,** which directly impacts customer satisfaction (fast, on-time delivery) and cost. AI-driven route optimisation, for example, is being explored by the logistics arms of large retailers and courier services. One local delivery company has integrated an AI scheduling system that plans driver routes for deliveries, taking into account factors such as traffic patterns, parcel urgency, and fuel efficiency. They reported potential savings of up to **35% in fuel costs and around 40% reduction in overall transportation expenses** after testing AI-optimised routes versus manually planned ones. These are significant numbers for any operation with a fleet. Those savings can be reallocated to marketing, passed on as price benefits to customers, or simply used to improve profitability. The same solutions also improved delivery times, which in turn feeds back into better customer satisfaction – a virtuous cycle for the retailer.

On the content side, some South African e-commerce companies are utilising AI to generate product content quickly, thereby expanding their online catalogues without hiring large content teams. Local marketing firms are adopting AI tools to manage

social media for multiple clients, where the AI drafts posts and the human polishes and approves them, enabling the firms to serve more clients effectively.

Key considerations for success in marketing content generation include brand **consistency and quality control**. AI can produce content at scale, but it needs guardrails to ensure it stays on brand (tone, style) and doesn't inadvertently produce something offensive or inaccurate. Companies should invest time in "training" the AI on their brand guidelines and have humans in the loop to review at least in the initial phases. Over time, the AI can be trusted with greater autonomy for low-risk content; however, high-impact campaigns should still involve creative directors.

There's also the risk of **content saturation** – just because AI lets you send 10 times more emails doesn't mean you should. Marketers must use AI to increase relevance, not just volume, or consumers will tune out. So, integrate AI with your analytics: monitor engagement metrics closely to ensure AI-driven content is actually effective.

For sales applications, one key consideration is **Salesforce adoption** – salespeople may be resistant to change. You'd need to introduce AI assistance (like an AI that summarises their calls or suggests follow-ups) in a way that complements their workflow and demonstrates a clear benefit; otherwise, they may ignore it. The AI suggestions should be explainable so that a salesperson trusts them ("AI recommends calling Client X this week because their company just announced an expansion and they might need more of our product").

Data privacy is also a significant concern here: personalised marketing and sales AI rely heavily on personal data. Ensuring compliance with POPIA (e.g., obtaining proper consent for data usage, allowing opt-outs) is critical. However, when done correctly, customers often appreciate the enhanced relevance of their engagements.

In summary, marketing and sales can use GenAI to work smarter – crafting tailored messaging and identifying opportunities from data – thereby driving revenue growth. The human creative spark and strategic oversight remain vital. Still, the heavy lifting of content production and data crunching can increasingly be offloaded to AI, allowing our marketers and sales teams to focus on strategy, relationships, and innovation.

To give a structured view, **Table 1** below maps some of these key use cases against their current maturity in South Africa and the potential impact they can have in the next year, along with enabling factors to consider:

Use Case	Primary Business Function	Current Adoption Level (SA)	Potential Business Impact (Next 12 Months)	Key Enablers & Local Considerations
Multilingual Customer Service Bots	Customer Service, Sales	Early Production in leading firms (common in banks/telcos)	**Transformational:** Can drastically improve service reach and consistency while cutting support costs	Requires investment in models supporting local languages; leverages high WhatsApp penetration in SA;

Use Case	Primary Business Function	Current Adoption Level (SA)	Potential Business Impact (Next 12 Months)	Key Enablers & Local Considerations
				crucial for inclusion and rural customer reach. Must align with POPIA for customer data.
AI Code Generation Assistants	IT / Software Development	Scaling (many dev teams piloting tools)	**High:** Boosts developer productivity, reduces backlog; faster feature delivery and fewer bugs	High developer adoption already; need governance on code quality and IP (ensure AI does not introduce vulnerabilities or license issues). Training devs to effectively use AI is key.
Enterprise Knowledge Search	All functions (Ops, HR, Legal, R&D)	Early Production in some large enterprises	**High:** Breaks down silos, speeds up decision-making; employees spend less time searching, more time doing	Leverages new local cloud infrastructure (for secure data indexing); addresses shadow AI by providing a sanctioned knowledge tool. Depends on the quality of the data and its classification to be effective.
Predictive Maintenance	Operations, Manufacturing, Mining	Piloting & Experimentation	**High:** Proven ROI in asset-heavy industries; prevents downtime, reduces maintenance costs significantly	Requires robust IoT sensor data and engineering buy-in; aligns with sustainability and safety (avoiding environmental incidents and injuries). Start with critical equipment for quick wins.
Personalised Marketing Content	Marketing, Sales, E-commerce	Piloting & Experimentation (some retail/banks testing)	**Medium to High:** Can increase campaign engagement and conversion; might	Automation of content creation at scale requires a strong first-party data

Use Case	Primary Business Function	Current Adoption Level (SA)	Potential Business Impact (Next 12 Months)	Key Enablers & Local Considerations
			boost revenue per customer through better targeting	strategy to personalise effectively; it must be carefully managed for brand voice consistency and POPIA compliance.
Back-Office Document Automation	Finance, HR, Admin	Scaling (common tools in use for drafting)	**Medium:** Delivers incremental productivity gains, error reduction in routine paperwork; frees staff for higher-value tasks	High adoption for tasks like report writing and form generation via MS Office plugins; these are low-risk and easy to implement. Ensure staff are trained to review AI outputs to maintain accuracy and reliability.

Table 1: High-Impact GenAI Use Cases – South African Enterprise (2025)

As shown, several use cases are already moving from pilot to production among forward-thinking South African organisations, with transformative potential especially in customer experience, software productivity, and operations. These can serve as focus areas for enterprises kick-starting their GenAI journey. In the following sections, we will delve deeper into **how to execute** these and other initiatives in a governed, economically sound way – covering governance, risk management, return on investment, and roadmaps for implementation.

The Governance Imperative: Navigating Risk and Regulation

For CIOs, CTOs, and any executive championing AI, it's critical to recognise that **the transformative potential of GenAI is inseparable from the responsibility of managing its risks.** Deploying generative AI in an enterprise without proper governance is akin to letting employees drive a powerful new machine without training, licenses, or safety checks. The result can be anything from data leaks and compliance violations to ethical landmines and public relations disasters. In South Africa, particularly in light of robust privacy laws and socially sensitive contexts, a strong governance framework is not merely a bureaucratic hurdle – it is a **prerequisite for the adoption of sustainable, trusted, and legally compliant AI.** Far from impeding innovation, good governance actually **enables** scale by building the trust and safety needed to proceed confidently.

Let's break down the key elements of AI governance in the South African enterprise context:

POPIA as the North Star for AI Governance

South Africa's **Protection of Personal Information Act (POPIA)** should be viewed as the guiding framework for any AI governance program that touches personal data. GenAI systems often ingest and produce data that can be sensitive (e.g., customer interactions, HR memos), so compliance with POPIA isn't optional – it's mandatory and fundamental.

POPIA establishes several core principles for handling personal information. These map closely to the controls one needs on GenAI:

- **Lawfulness, Fairness, and Transparency:** POPIA requires that personal data be processed lawfully and fairly, with the knowledge and consent of the individual where required. Applied to AI, this means you should have a legitimate basis (e.g., consent or contractual necessity) for using personal data in training or running an AI model. You also need to be transparent: if you have a customer-facing chatbot, for instance, it's advisable to disclose to users that they are interacting with an AI and how their data is being used. Fairness further implies that the AI's decisions or outputs shouldn't unjustly discriminate among people (more on bias below).

- **Purpose Specification & Data Minimisation:** Personal data should be collected for a specific purpose and not used beyond that scope; only data strictly necessary for that purpose should be processed. For GenAI, this is a call to avoid the all-too-easy temptation to vacuum up every scrap of data "because it might be useful for the AI." For example, if you are training a customer service AI to answer banking questions, you likely do not need to feed it detailed customer financial histories or national ID numbers. Limit data to what the AI needs to function. And clearly define the purpose – e.g. "This model will be used to assist call centre queries, and not for marketing or credit scoring." Document that purpose and ensure the AI isn't later repurposed without re-evaluating compliance.

- **Security Safeguards:** POPIA mandates that responsible parties secure personal data against loss, unauthorised access, and disclosure. GenAI presents new challenges: an employee could inadvertently input confidential data into an AI service (especially a third-party one), and that data could be stored or even used to train the provider's model (cases of this have occurred globally, such as when employees pasted source code into public AI). To safeguard, organisations need technical and organisational measures. Technical measures might include using **self-hosted or private cloud AI environments** where data doesn't go to external servers without control, implementing **data loss prevention (DLP)** systems that detect and block sensitive information being sent to unapproved AI tools, and encrypting data in transit and at rest for AI systems. Organisational measures include clear policies (see below) and training so that staff know, for instance, not to paste

private personal data into ChatGPT (unless it's an approved enterprise version with proper agreements).

- **Rights of Data Subjects:** POPIA grants individuals the right to access their data, correct it, and object to its processing in certain cases. With AI, one applicable right is the implied right **not to be subject to a purely automated decision that significantly affects them, without human intervention** (this is analogous to provisions in laws like the EU's GDPR). In practical terms, if an AI model were making, say, hiring recommendations or lending decisions, POPIA would effectively require that individuals can request a human review of those decisions. Even outside of legally mandated areas, it's a good governance practice to keep a "human in the loop" for consequential decisions. Additionally, organisations should be ready to **explain AI-driven decisions** if asked. This ties into emerging global norms around AI transparency – even if our law doesn't explicitly say "you must explain your AI's reasoning," being prepared to do so will cover POPIA's fairness and data subject engagement principles.

- **Accountability:** POPIA places the duty of compliance on the organisation (the responsible party). Translated to AI, this means that if an employee uses an AI tool in a way that breaches privacy – for example, uploads a client list to an AI service that then gets compromised – the company is accountable for that breach, not the individual alone. This underscores the importance of executives proactively establishing guardrails for the internal use of AI. One case to note: if employees use AI improperly (like shadow AI tools) and a data breach occurs, **the employer can be held liable** under POPIA. This is already motivating local companies to establish internal AI usage policies and monitoring.

In essence, **POPIA compliance should be built into the design of AI solutions ("privacy by design")**. When developing or deploying GenAI, always consider: are we honouring consent? Are we minimising data? Do we have security? Can we accommodate people's rights?

Key Risks of Ungoverned GenAI

Without a formal governance structure, enterprises face three primary categories of risk:

1. **Data Leakage and Privacy Breaches:** This is the most immediate and severe threat in the GenAI arena, particularly in the context of shadow AI use. A vivid example is an employee taking a chunk of confidential data – such as a customer's account details or an internal strategy document – and posting it into a publicly available AI chatbot to get help analysing or writing something. The moment that happens, control over that data is lost. It may reside on external servers outside our jurisdiction, possibly be seen by AI trainers, or even get inadvertently regurgitated to another user if the AI's memory is not well compartmentalised (there have been reported incidents globally of AI tools revealing snippets of other users' inputs). Under POPIA, such an incident could be a reportable data breach with legal ramifications. Even using

ostensibly safe tools can be risky if not configured correctly – for example, using a cloud-based AI service without disabling the default setting that saves inputs for model training. The risk isn't just external tools; even an internal AI could leak data if it doesn't have role-based access controls (imagine an HR AI that accidentally allows any employee to query salary info because it was trained on the HR database but not permissioned in responses). The bottom line is that governance must enforce **strict rules about what data can be fed into AI and how that data is handled.**

2. **Algorithmic Bias and Discrimination:** Generative AI models learn patterns from data. If those patterns reflect historical biases, the AI can perpetuate or even amplify them. In a diverse society like South Africa with deep inequalities, this is a critical concern. For instance, if an AI is used in screening CVs for recruitment and is trained on past hiring data or general internet text, it might pick up biases (perhaps unjustly favouring certain genders, universities, or regions). Without checks, this could lead to discriminatory outcomes, which are not only ethically wrong but could also violate equality laws. Similarly, a customer-facing AI might not respond equally well to queries in different languages or might have learned subtle, biased associations (such as assuming certain surnames or areas correlate with specific behaviours). This risk requires active measures, including testing AI outputs for bias, ensuring that training data is balanced and representative, and applying algorithmic fairness techniques to mitigate bias. It's a governance responsibility to set up those measures. Failing to do so risks **reputational damage** ("Company's AI is biased against X group" headlines) and **legal challenges** (e.g., if biased decisions result in unfair treatment, there could be grounds for complaints or lawsuits under discrimination laws).

3. **Hallucinations and Misinformation:** GenAI models are notorious for sometimes producing outputs that are factually incorrect or entirely fabricated – what we call "hallucinations." They do so in a very confident, plausible manner, which can be dangerous if users take the output at face value. In a business context, this could lead to serious errors. Imagine an AI assistant that drafts an email to a client and cites incorrect figures or non-existent policy details. If unchecked, an employee might send that out, causing confusion or contractual issues. Or an AI system generates an analytical report for management with some made-up data points that slip through review and lead to a poor decision. Another dimension is external: AI makes it easier for malicious actors to create disinformation (e.g., deepfake content, fake news articles). While that's more an external threat, companies need to be aware – e.g., a deepfake of a CEO could circulate, affecting stock or stakeholder trust, and the company should have a response plan in place. Internally, however, the focus is on **verifying the accuracy of AI outputs**. Good governance means setting a policy that "AI-generated content must be reviewed by a human before use in any official or customer-facing capacity" (especially early on, until trust is established). It also means training users not to trust the AI blindly – encourage a healthy scepticism and double-checking of critical facts. Over-reliance without verification can erode trust with customers and lead to mistakes that must be corrected at a cost.

Beyond these top three, there are other risks: intellectual property leakage (if AI is used to generate code or content, who owns it? does it inadvertently include licensed material from training data?), regulatory non-compliance (e.g., upcoming AI-specific regulations may impose requirements and penalties), and model security (the risk of someone attacking your AI model – say, extracting its training data via clever queries, or feeding it malicious inputs to manipulate its output in harmful ways).

However, most of these fall under the umbrella of governance processes and controls. The message is that **ungoverned AI = unmanaged risk.** It's far better to proactively implement guardrails than to react after a scandal or breach has occurred.

The Necessity of a Formal AI Usage Policy

One of the most immediate and practical tools in AI governance is instituting a **clear internal AI usage policy**. This is essentially an organisation-wide directive that spells out how employees may or may not use AI tools in their work, and under what conditions. We can think of it as analogous to the early days of "internet usage policies" or "BYOD policies" – new tech requires new rules of the road.

Why a formal policy? Because it addresses the shadow AI problem head-on. Employees often turn to convenient AI tools to get their job done, but many may not realise the risks (as we outlined). A policy both *educates* and *regulates*.

A robust AI usage policy in a South African enterprise should include at a minimum:

- **Prohibition on Uploading Sensitive Data to Unapproved Tools:** It should be explicit that no confidential, personal, or otherwise sensitive information is to be entered into any AI service that hasn't been vetted and authorised by IT and compliance. For example: "Do not input customer personal information, non-public financial data, source code, strategic documents, or any information classified as internal/confidential into public AI platforms (e.g., the free versions of ChatGPT, Google Bard, etc.) or any third-party AI service that has not been approved for use by the company." This clause, if followed, would mitigate a significant portion of the data leakage risk.

- **Approved AI Tools and Vetting Process:** The policy should list any officially approved AI tools (e.g., the company may provide an internal AI assistant or have licensed a secure enterprise version of an AI platform) and state that all other tools are off-limits until cleared. It should also outline how a new tool can be requested and evaluated. By formalising this, you simultaneously discourage shadow use and provide employees with a path to acquire new tools if they demonstrate value and can be secured.

- **Mandate for Human Oversight:** As noted earlier, require that AI outputs used for any decision-making or external communication be reviewed by an appropriate human. This could be phrased as: "AI-generated content must be treated as a draft or suggestion. Employees are responsible for reviewing and validating such content before acting on it or sharing it externally." This ensures that accountability remains with humans, and the AI serves as a

helper, not an autonomous actor (except perhaps in narrowly defined cases where automation has been thoroughly tested and approved).

- **Ethical Use Guidelines:** The policy should embed ethical principles. For instance, instruct that AI should not be used in ways that would violate human rights, enable illegal discrimination, or generate inappropriate content. Suppose employees are using generative AI to create presentations or prototypes. In that case, they should avoid using it to generate harassing, explicit, or culturally insensitive content – common sense, but better spelled out to avoid "the AI made a joke that we didn't realise was offensive" type incidents. Additionally, guidelines are in place to prevent misleading people (e.g., not using deepfake technology for any company purpose without transparent labelling).
- **Accountability and Roles:** Clearly state who in the organisation is responsible for AI oversight – e.g., "The Risk/Compliance Officer (or an AI Governance Committee) will maintain this policy, monitor compliance, and address any incidents." And that violation of the policy may result in disciplinary action (to give it teeth). This signals that leadership takes it seriously.
- **Training and Awareness:** Ideally, the policy is accompanied by training sessions or materials. Employees should sign off that they've read it, and maybe go through a short course explaining do's and don'ts with realistic examples: "Here's a scenario, what should Alice do vs not do with AI?"

Implementing such a policy is arguably the **fastest and highest-impact step** to rein in shadow AI risk. It doesn't require new technology – it requires clarity and communication. Many South African firms are currently rolling these out (some banks and insurers have already done so, given their sensitivity to data security). This doesn't solve everything – policies need enforcement and eventually technical controls to back them up (like DLP systems to catch violations). But it's a foundational move.

Turning Compliance into Competitive Advantage

It's worth highlighting that **good governance is not just about avoiding negatives; it can create a positive differentiation for companies**. In conversations with business leaders, especially those dealing internationally, it's clear that **trustworthy AI operation is becoming a selling point**. For example, European clients or partners (facing their own strict AI regulations soon) might demand that a South African supplier demonstrate robust AI governance before doing business. Suppose your company can demonstrate its POPIA compliance and alignment with global best practices (such as documenting algorithms and conducting bias testing). In that case, it may be preferred over a competitor with a more casual approach.

We see a parallel in cybersecurity: companies with strong security postures win contracts because customers feel safer sharing data or integrating systems. Similarly, a "POPIA-compliant by design" AI status can be marketed. It signals to

customers that you respect their data and rights (building trust), and to regulators that you're low risk.

Also consider talent and public perception – employees want to work at companies that aren't going to land them in ethical quandaries, and the public is increasingly wary of AI missteps. Proactive governance thus also enhances internal culture (employees are aware of the boundaries and are less likely to be concerned about potential issues arising from unclear rules) and brand reputation.

As global frameworks evolve (e.g., the EU AI Act likely coming into effect around 2024/25), being ahead of the curve locally means you'll already comply with many of those standards, making it easier to trade or partner overseas. The EU AI Act will require risk assessments, transparency, and the possible registration of certain AI systems, among other measures. If a South African firm masters those disciplines early via POPIA alignment, it stands ready for global markets that require evidence of "trustworthy AI".

In summary, **robust governance and compliance can be transformed from a domestic obligation into a global business enabler**. By building AI solutions and processes with fairness, privacy, and accountability in mind from the start, South African enterprises not only avoid pitfalls at home but also position themselves as forward-thinking, trustworthy players internationally. That is a strategic edge in an environment where AI-related scandals or breaches could significantly set others back.

Having established the why and what of governance, the next step for enterprise leaders is execution: creating the governance structures (like an AI steering committee, appointing responsible officers), rolling out policies as discussed, and investing in the necessary tools (monitoring systems, bias audit tools, documentation workflows) to operationalise these principles. We will revisit some of these points in the **Risk Management** and **Metrics** sections later. In the **Roadmap,** we will suggest concrete steps for the first 12 months to establish governance.

However, one thing should be clear: **for South African enterprises, AI governance is not a "later" issue – it is an immediate priority that must run in parallel with any AI adoption.** Those who embrace a "governance-first" mindset will unlock GenAI's value far more safely and sustainably than those who treat governance as an afterthought.

Strategic Recommendations for the Next 12 Months

The current generative AI landscape in South Africa – characterised by high interest, rapid grassroots adoption, but also rising risks and scepticism – presents a **clear call to action** for enterprise technology leaders. The window of 2025–2026 is critical: companies that move proactively now can set themselves up as leaders, while those that delay may fall irreparably behind or stumble into pitfalls.

To translate the insights and trends we've discussed into concrete action, we outline five strategic recommendations. These can serve as a focused roadmap for CIOs, CTOs, and their executive peers to execute over the next 12 months (and beyond),

ensuring GenAI is approached in a value-driven, secure, and aligned manner with South African market realities.

1. Formalise an Enterprise AI Strategy to Eliminate "Shadow AI"
Urgency: High (immediate quarter)
What: The most urgent priority is to address the strategic vacuum that has allowed uncontrolled "shadow AI" use to proliferate. This means establishing a **centralised AI strategy** endorsed by top leadership and setting up a cross-functional governance body to enforce it. This is about moving from fragmented, individual-driven experimentation to a cohesive program guided by business objectives and risk awareness.

How: Start by assembling an **AI Steering Committee** or Task Force that includes key stakeholders: IT and data leaders, legal/compliance officers, risk management, HR (for training and change management), representatives from major business units, and, importantly, an executive sponsor (CXO level) to give it authority. This body's first task is to articulate an **AI vision and policy framework**. For example: "Our enterprise will leverage GenAI to enhance customer experience and operational efficiency, within the bounds of strict data privacy and ethical standards, aiming for X% ROI in 2 years."

Then, conduct a rapid audit of current AI usage across the company (including shadow usage). Identify quick wins and high-risk practices. Use this to develop an **official AI use case portfolio** – e.g., decide on 3-5 high-impact use cases to focus on (likely those identified earlier like customer support chatbot, internal knowledge assistant, etc.), and designate which tools/platforms will be used for each (preferably standardized on a small set of approved platforms instead of dozens of disparate ones).

Simultaneously, **publish the AI usage policy** we detailed in the governance section (covering what's allowed, what's not, and who must approve new AI initiatives). Communicate it clearly across the organisation via emails, town halls, intranet FAQs, and other channels. People need to understand why certain behaviours (such as using unauthorised AI apps) are now prohibited and what alternatives or support they have (like an approved internal AI or a process for requesting AI solutions).

Finally, ensure the steering committee establishes a basic **AI project governance process**: any new GenAI project should come through this body for initial review (to assess feasibility, data needs, alignment with strategy, and risk mitigation), and then proceed with that oversight. This might sound daunting, but it can be straightforward – perhaps just a one-page submission and a committee green light. The point is to funnel AI efforts through a controlled gate rather than having them scattered.

Outcome: By doing this, the organisation shifts from chaotic AI adoption to coordinated implementation. Employees and managers will know: GenAI is something we're embracing, but in a managed way. This significantly reduces the chance of nasty surprises (like discovering a department exposed data via an AI tool), and it focuses resources on the most promising initiatives instead of random experiments. Essentially, you're turning GenAI from a rogue phenomenon into a formal program akin to any other strategic initiative.

2. Prioritise the "AI-Ready" Data Foundation
Urgency: High (parallel with strategy formation)

What: Generative AI cannot deliver scalable value on a foundation of poor-quality or inaccessible data. We've highlighted that ~57% of organisations consider their data not AI-ready – this is a bottleneck we must remove. Thus, an immediate strategic thrust should be a **data readiness initiative**: invest in data quality, integration, and governance improvements that specifically enable the use of AI.

How: Kick off a targeted **Data Audit** focusing on the data assets required for the top-priority AI use cases identified. For instance, if one use case is a customer service chatbot, examine the customer support knowledge base, past chat/email logs, FAQ documents, etc. Are they up to date? Are they stored in a way that an AI can retrieve info easily (perhaps they need to be consolidated into a single repository)? If another use case is predictive maintenance, audit the sensor data – is it clean, labelled, complete, streaming reliably?

From these audits, develop a short-term **Data Improvement Plan** that may involve cleaning and tagging unstructured text (such as adding metadata to documents), merging duplicate records, filling gaps (possibly by augmenting with external data or synthetic data where permitted), and setting up pipelines to gather and update relevant data continuously. Consider employing data cleaning tools or, alternatively, a task force to manually correct key datasets if needed (a time-consuming but sometimes necessary approach for critical data).

Simultaneously, **strengthen data governance structures**: ensure data ownership is assigned (each major dataset should have a business owner and a data steward on the technical side), and establish/update policies for data access specifically for AI (for example, create a process to provision a secure data environment for AI model training where needed, rather than ad-hoc copying of data).

On the technology side, implementing or optimising a **data platform** could be vital – e.g., ensuring you have a data lake or warehouse where the various departmental data needed for AI can be aggregated under proper access controls. Additionally, consider deploying specific AI-related data tools, such as **vector databases** (useful for semantic search, which GenAI utilises) and **feature stores** (if working on predictive models), as relevant to your use cases.

The key is to sequence this work with your AI projects: you might choose one flagship project and prioritise the data work for that, thereby delivering immediate value and creating a template for others.

Outcome: By the end of 12 months, aim to have the core data for at least one or two strategic AI applications in a state that project teams confidently say, "We have what we need, and it's good quality." Additionally, the organisation should see improved data practices – fewer silos (since you probably had to negotiate some data sharing for AI), more awareness of data as a valuable asset, and hopefully a reduction in time spent by staff just trying to find or clean data for analysis (which is often a hidden cost). This sets the stage for all future AI efforts – it's a one-time heavy lift that will pay dividends repeatedly. Moreover, it likely benefits other analytics and reporting needs too, not just AI.

3. Champion an "Augmentation-First" Workforce Strategy
Urgency: High (start immediately, cultural shift over 12+ months)
What: In the South African context, it's absolutely vital to shape the narrative and practice around GenAI as augmenting human workers, not replacing them. A fearful

or resistant workforce can sabotage AI adoption, whereas a supported and AI-educated workforce can supercharge it. Therefore, tech leaders should partner with HR and business unit heads to drive a comprehensive **AI literacy and upskilling program** for employees at all levels.

How: Begin with communication from the top: messages from the CEO or relevant executives that **"AI is here to help you do your job better, not take it away."** Link it to the company's values – e.g., innovation, growth, and also to national interests (e.g., "as a proudly South African company, we are using AI to empower our people and better serve our community"). Actions must follow words: announce training opportunities and encourage staff to participate.

Implement a tiered **AI Education Program** for general staff, focusing on basic AI literacy (what is GenAI, what it can and can't do, and examples of how it might assist in daily tasks). For power users or specific roles, more advanced workshops are available (e.g., how to prompt AI tools to achieve good results effectively – "prompt engineering" as a new skill; or for data analysts, training on leveraging new AI-driven analytics tools). This could involve e-learning modules, live demos, inviting external experts for seminars, or even sending people to industry courses. There are many free or affordable resources now which are helpful.

Crucially, tie this to **employee involvement in AI projects**: when rolling out that new internal chatbot or code assistant, involve end-users early. Run pilot programs where a small group of employees test the tool in their routine and give feedback. This accomplishes two things: it improves the tool through practical feedback and creates internal champions who will later help their peers adopt it. Perhaps set up an internal community of practice – a Teams channel or Yammer group where employees share tips on using AI in their jobs, ask questions, etc.

Another aspect: work with HR to identify if any roles may undergo significant changes due to AI and plan accordingly. For example, if AI automates certain level-1 support queries, then design a plan to **reskill those support staff** for higher-level customer engagement tasks (using the time AI saved). Or if AI will produce first drafts of marketing copy, train your copywriters to be great editors and AI supervisors rather than worrying about losing writers. Align performance metrics to encourage using AI (e.g., include an objective like "Identify at least one way to use AI to improve your work process" in performance plans, to signal it's encouraged to experiment).

In essence, create a culture where using AI tools is viewed as a skill and a positive asset, much like proficiency in Excel or effective communication. Celebrate successes: if someone used AI to significantly improve a metric (say, a salesperson closed deals faster by using an AI to draft proposals quickly), share that story company-wide.

Outcome: Over 12 months, you should see a shift in how employees talk about AI – from anxiety or hype to practical enthusiasm. People will start coming up with their own ideas of tasks AI could help with. The adoption of AI tools that the organisation officially provides will be higher because employees understand them and feel confident (a lack of training is a common cause of tech project failure). Moreover, this strategy aligns the company with the socio-economic need to upskill workers: you're not just expecting employees to "figure AI out" or worse, quietly planning to

cut jobs – you're actively investing in them to grow alongside AI. This yields more innovation from the ground up, as employees who know their domain best will start to contribute AI improvement suggestions. It also positions the company as a progressive employer (helpful for attracting talent who want to work with cutting-edge tech in a supportive environment).

4. Implement a Robust AI Usage Policy Grounded in POPIA
Urgency: Immediate (policy should be drafted and rolled out within 1-2 months)
What: As elaborated in the governance section, a formal internal AI policy is a must-do to curb the biggest risks quickly. Craft this policy with a specific nod to South Africa's POPI Act and the enterprise's own risk appetite.

How: Form a small team (IT, legal, risk) to draft the policy content. There are emerging templates available that you can reference, but ensure they are customised to your environment. Key components we have already listed include: no confidential data in public AI, a list of approved tools, a process for new tool approval, a requirement for human review of AI outputs, ethics guidelines, etc. Obtain executive sign-off and ideally a statement of support, so it's not just an IT memo but a company-wide directive.

Announce the policy through multiple channels. Consider incorporating it into your code of conduct or IT usage agreement that employees must sign. Host a Q&A session to let employees clarify doubts – some might ask "Does this mean I can't use GitHub Copilot at all?" or "Are we banning ChatGPT entirely?" and you need clear answers. Often, it might be "you can use X tool but only on non-sensitive data or tasks, and here's the process if you're unsure."

Put monitoring in place to enforce it: e.g., network monitoring might flag the use of certain AI web services, so you know if the policy is being breached (and then address the issue with that team, not necessarily punitively, the first time, but as an educational measure). If you have DLP (Data Loss Prevention) software, update its rules to detect patterns such as someone attempting to paste a customer ID or large text chunk into a web form (which could indicate feeding it to an AI web app) and block it, notifying the user of the policy. This technical guard can significantly uphold the policy.

Also, include in the policy that any AI procurement (i.e., buying an AI solution or engaging an AI vendor) must undergo due diligence for POPIA and security. That will protect you from accidentally signing up to an AI SaaS that might take your data elsewhere without you realising – now you'll check things like where their servers are, do they train on your data, do they sub-share data, etc.

Outcome: Within a few short weeks of the policy, you should see a behavioural change: employees hesitating and thinking twice before using external AI tools for work matters. Ideally, they redirect those needs to approved methods (or escalate requests for IT to provide solutions). This drastically reduces the immediate risk of data leaks and keeps you out of newspaper headlines. The organisation will be able to confidently state (to its board, auditors, and customers) that "we have a policy in place to manage AI usage, aligned with privacy laws and best practices." That will soon become a baseline expectation.

You might also uncover unknown usage that's valuable – e.g., discovering through monitoring that a particular department was using a tool because it solves a specific

problem; you can then examine that problem and perhaps officially adopt a safer solution for it. In that way, the policy not only restricts but also channels innovation through the right pipeline.

5. Launch Strategic Projects with a Local Focus
Urgency: Within 1-6 months (and continue scaling through 12 months)
What: While the above steps handle governance and foundations, you also need to demonstrate quick wins and tangible value to maintain momentum and executive support. Therefore, identify **2-3 high-impact pilot projects** that address specific needs of your enterprise and the South African market context. These pilots should be scoped to deliver proof of value within a year or less, be realistically achievable with your resources, and align with the themes we've discussed (customer experience boost, cost efficiency, etc.).

How: Based on earlier analysis, prime candidates for many organisations include:

- A **multilingual customer service chatbot/voice assistant**: Especially for any B2C business in SA, deploying a chatbot that converses in English + at least a couple of local languages can greatly improve customer reach and satisfaction. Choose a channel that's popular (e.g., WhatsApp, as we saw with Capitec). The pilot could start by answering a specific set of FAQs in English, Zulu, and Sotho, gauge usage and feedback, and then expand. Success metrics include a reduction in call centre volume, CSAT improvement, and the containment rate of the bot, among others. The local focus here is on the multilingual aspect and potentially the inclusion angle (such as addressing issues like reaching rural customers via low-bandwidth channels).

- **AI-powered predictive maintenance** in a core industrial area: If you're in manufacturing, mining, energy, or transport – identify one critical asset or process where downtime is costly. Implement a GenAI/AI model to predict failures or optimise maintenance schedules for that asset. For example, a mining company might deploy an AI on a specific type of truck or processing mill. Success metrics: reduction in unplanned downtime, maintenance cost savings, and safety incidents avoided. Local focus: aside from cost, it also directly protects jobs (less downtime -> stable production -> stable workforce need) and could tie into sustainability (a well-maintained machine is often greener and less accident-prone).

- **Enterprise knowledge management tool (GenAI-based)**: Internally, deploy an AI assistant that employees can ask company-related questions. Consider starting within a department like IT or HR for internal queries (e.g., "How do I file an expense report?" AI answers can be found in policy documents). Or for a bank, perhaps an internal AI for relationship managers to quickly query product information and customer data. Success metrics include a reduction in time to find information, the number of queries resolved by AI, and employee feedback on the usefulness of the system. Local adaptation: ensure it handles content in multiple languages if your internal knowledge base has that, and that it accounts for local regulations when advising (like it might include compliance reminders if someone asks about a process that is governed by law).

The important thing is these pilots are **scoped – not moonshots but meaningful improvements** – and you plan for measuring outcomes. Baseline the KPIs now (e.g., current call volumes, downtime rates, employee survey on time spent searching for info, etc.) so you can concretely show improvement.

Execute these pilots in an agile manner: use a small team, conduct rapid iterations, and involve end-users early (as recommended). Probably target a 3-to 4-month development and testing period, followed by a soft launch and then expansion.

Also, incorporate your governance from day one in pilots: conduct a mini risk assessment for each, ensure the data is cleaned and secure (pilots are a good test of your data initiative too), and monitor for issues such as bias or errors during the run.

Outcome: Within 12 months, you should have at least one, ideally a few, functioning GenAI-powered solutions in production delivering real benefits. For example, by Q4 you might report: “Our AI chatbot now handles 25% of customer inquiries and maintains a satisfaction score of 80%+, freeing our call agents to focus on complex cases,” or “Our predictive maintenance AI cut machine downtime by 30%, saving R2 million this year,” or “Our internal AI assistant is used by 500 employees who report saving on average 1 hour per week of searching time.” These concrete wins not only justify the investment but also build confidence and appetite for further scaling.

Crucially, each successful pilot becomes a showcase: you can take the lessons (both technical and change management) and apply them to other use cases. They also demonstrate to any remaining sceptics in the company that AI is not just hype – under governance, it works for us and aligns with our business values (like better customer service, operational excellence, etc.). That goes a long way in sustaining momentum into the next 2-3 years of broader AI transformation.

By following these five strategic steps, a South African enterprise can rapidly transition from a reactive, scattered approach to GenAI into a proactive, strategic posture. In summary: establish leadership and guidelines (Strategy & Policy), reinforce the supporting structure (Data & Skills), and implement focused initiatives (Pilots) that prove value.

This combination tackles the challenge holistically – technology, people, process, and governance. It sets a foundation for scaling AI in the following years towards even more ambitious goals (such as enterprise-wide AI integration and advanced analytics), while minimising the pitfalls that often plague early adopters.

In the coming sections, we will delve further into specific areas like measuring ROI (to accompany these pilots), longer-term roadmaps beyond 12 months (once these foundations are laid), sector-specific nuances (how these recommendations might be tailored for finance vs retail vs public sector), and change management details to ensure these strategic moves stick.

However, as a one-year game plan, the recommendations above provide a pragmatic pathway to swiftly and safely adopt GenAI, transforming it from a buzzword into a source of competitive advantage and innovation for the enterprise.

References

1. World Wide Worx, Dell & Intel – **"South African Generative AI Roadmap 2025"** (Study released July 2025). *Summary in IOL Business Report, 17 July 2025, noting 67% enterprise GenAI adoption in 2025, but only 14% have formal strategies and 13% have governance frameworks.*
2. N. Nkwashu, ITWeb – **"CIO Survey reveals shift in GenAI adoption"** (Brainstorm CIO Survey 2024 results, published 18 Oct 2024). *Findings: GenAI adoption among SA CIOs doubled from 16% (2023) to 38% (2024); the top use cases include cybersecurity (50%), productivity (43%), and knowledge management (38%). Emphasises the importance of strategic vision and skills.*
3. Gartner (Press Release) – **"30% of Generative AI projects will be abandoned after proof-of-concept by the end of 2025"** (Sydney Data & Analytics Summit, 29 July 2024). *Statistic: Escalating costs and unclear ROI will cause at least 30% of GenAI pilots to be halted by 2025. A survey of 822 business leaders showed only ~16% saw revenue increases and 22.6% productivity gains from GenAI so far.*
4. Forrester (2026 Predictions) – **"AI spending slows as tech leaders face scrutiny in 2026"** (CFOtech Asia, 9 Oct 2025). *Insight: Fewer than one-third of businesses can link AI initiatives to P&L impacts, prompting CFOs to delay 25% of AI spend to 2027. Emphasis on governance: 60% of Fortune 100 to appoint Heads of AI Governance by 2026; increased focus on function over flair in AI projects.*
5. U. Sachdev, World Economic Forum – **"Enterprise AI is at a tipping point – here's what comes next"** (31 Jul 2025). *Key trends: shift from consumer AI hype to enterprise pragmatism; move from generic to agentic systems and domain-specific models; 62% of leaders cite data integration as top AI obstacle (Deloitte); composability seen as strategic (by 2026, firms with composable tech outpace others by 80% in feature deployment).*
6. McKinsey & Co. – **"The economic potential of generative AI: The next productivity frontier"** (MGI Report, 14 Jun 2023). *Notable figures: GenAI could add $2.6–$ 4.4 trillion annually across various industries. Banking potential $200–340B per year (with ~4.2x ROI in AI leaders); Retail $400–660B. Up to 60–70% of worker activities could be automated by tech (vs 50% estimated previously). Highlights the importance of complementary workforce transformation.*
7. Credence Research – **"South Africa AI in Retail Market – Growth, 2024–2032"** (2023). *Statistic: SA's AI-in-retail market size $31.42 in 2023, projected to be $281.91 by 2032 (27.5% CAGR). Drivers include personalised shopping experiences, supply chain optimisation, and AI chatbots for service. Indicates that retailers are accelerating AI adoption for increased efficiency and enhanced customer insights.*
8. Admire Moyo, ITWeb – **"Standard Bank to make AI training compulsory for staff"** (Interview with CIO, 2 Sept 2025). *Details: Standard Bank is rapidly deploying AI enterprise-wide to boost CX and productivity. All 50k+ employees required to undergo AI ethics and risk training ("dos and don'ts of AI"). Bank established internal "religious risk processes" for AI use (ensuring data doesn't leak, appropriate model use). Emphasises the importance of strong governance and workforce upskilling as part of the AI rollout.*

9. Gartner (Press Release) – **“By 2027, more than 50% of GenAI models used by enterprises will be domain-specific”** (Forecast, April 2025). *Predicts shift from general LLMs to specialised models (up from ~1% in 2023). Cites benefits: higher accuracy, 30–60% lower TCO for targeted tasks, easier compliance in regulated sectors. Also notes* the mainstreaming of *synthetic data – by 2026, 75% of enterprises will utilise GenAI-generated synthetic data for prototyping and privacy purposes.*
10. Gartner (Press Release) – **“40% of GenAI solutions will be multimodal by 2027”** (IT Symposium, 9 Sept 2024). *Forecast: Multimodal AI (text, image, audio combination) rising from 1% to 40% of solutions, enabling richer human-AI interaction and new applications in design, education, and support. Highlights open-source LLMs and autonomous agents as high-impact trends on a 5–10 year horizon, aligning with the hype cycle (GenAI entering the Trough of Disillusionment in 2024 but poised for real value as hype levels out).*
11. Arthur Goldstuck (World Wide Worx) – **“SA organisations embrace GenAI”** (Market research commentary, 25 Apr 2024). *Findings from GenAI Roadmap 2024: 90% of large enterprises in SA are using or plan to use planning GenAI (45% actively using, 45% e evaluating, and have only 10% no plans). 95%+ see positive potential in productivity and competitiveness. However, emphasises the need for a strategic approach, security, and cultural readiness for sustainable deployment. Indicates strong enthusiasm but also underscores the foundations needed.*
12. Gartner – **“Gartner Data & Analytics Principles / POPIA Compliance”** (various Gartner ITL reports, 2024–25). *General guidance aligning AI governance with privacy: recommends mapping AI data flows to privacy principles, implementing bias audits, documentation (“model cards”), and human-in-the-loop for high-stakes AI decisions. Notes that enterprises demonstrating strong AI governance will gain trust of customers and regulators, effectively turning compliance into a market differentiator.*

www.cohesionx.co.za