GUIDANCE ON THE USE AGENTIC AI
IN PUBLIC SECTOR HEALTH

AI Agents Handbook

Richard Boyd - CTO Public Sector Health

Handbook Table of Contents

1

What is an AI Agent?

Agents work alongside your human employees, complementing & empowering them

01

What an Agent is NOT

Dispelling Myths & Tackling Change

When introducing Agents, it's crucial to address common misconceptions:

'AI is going to take my job!'

Agents are designed to enhance human capabilities, not replace them. They excel at handling repetitive, low-value tasks, freeing up staff to focus on more strategic, creative, and patient-centric responsibilities. By automating routine processes, Agents support workforce retention and reduce burnout, leading to a more engaged and fulfilled workforce.

'AI is a threat to my team's morale'

With **53% of government workers expecting rising workloads** and only **74% prioritising relationship building**, there's a gap between demands for personalisation and team capacity.

Agents address this by automating repetitive tasks, freeing staff to focus on meaningful, empathetic work-boosting morale and rebuilding public trust.

Source: State of Service, Salesforce, April 2024.

'AI will make mistakes–it's not ready for the real world!'

Like any new team member, Agents require onboarding and training. However, unlike humans, their training happens in weeks, not months. They continuously learn and improve from the data they process, becoming more accurate and efficient over time. It's about managing expectations and understanding that Agents, while powerful, are constantly evolving and improving.

'Agents are just for call centres–they can't help my department'

While Agents excel in handling patient queries and automating tasks in call centres, their applications extend far beyond. From guiding complex processes, to generating insights from data and automating content creation, Agents can adapt to meet a wide variety of organisational needs including internal efficiencies.

'AI will dehumanise our services–patients want a human touch'

Agents help staff be more human. By removing the burden of repetitive tasks, they free up time to focus on the tasks that require empathy, creativity, and relationship building. This leads to increased staff satisfaction, reduced burnout, and improved retention rates, creating a more engaged workforce that can provide the human touch that patients value.

01

So, what is an Agent then?

Agents are digital workers that can transform how we operate and solve challenges. They're more than just software; they're proactive collaborators capable of learning, adapting, and acting autonomously. Think of them as your newest team member, ready to take on tasks, provide insights, and help you achieve your goals. Crucially, Agents can be designed specifically for public services, ensuring secure data handling and responsible use, operating within established governance frameworks.

Superheroes of the digital world

These intelligent digital assistants are reshaping the public sector, empowering organisations to deliver exceptional patient services and achieve ambitious goals. Let's explore how Agentic technology works and what sets it apart:

Purpose-Built for Action

Agents are designed to go beyond answering questions–they take action to deliver meaningful outcomes. With advanced **Natural Language Processing (NLP)** and secure data access, they can automate workflows, process patient requests, and complete transactions. This makes them more than just digital assistants–they are transformative tools that improve efficiency and drive results in real-world scenarios.

Bridging Data & Systems

Agents thrive on data and act as a bridge between fragmented systems. By integrating seamlessly with existing platforms and legacy infrastructure, they unify disparate data sources to provide a cohesive experience. This ensures minimal disruption during implementation and allows public sector organisations to unlock the full potential of their data assets, while maximising the value of technology investments.

More than a Language Model

Unlike generic **Large Language Models (LLMs)**, Agents combine cutting-edge NLP capabilities with secure data access and the ability to act on insights, delivering tailored services that adapt to patient needs. Continuously learning and evolving, agents are designed to excel in specific tasks and deliver tangible results that go beyond traditional LLM capabilities.

Why This Matters

AI is poised to revolutionise public sector operations and enhance citizen services. However, it's important to remember that not all AI or Agents are created equal. To ensure you're getting the most out of this technology, consider the following:

- **Start with clear objectives** and prioritise data quality for accurate and reliable AI training.
- **Embrace a phased approach,** beginning with low-risk, high-impact use cases to demonstrate value and build momentum.
- **Champion human-AI collaboration,** designing agents to work alongside staff for a more engaged workforce.
- **Ensure transparency and accountability** by establishing clear governance frameworks and ethical guidelines.

The best part? Building Agents is easier than you might think. Business stakeholders actively participate in their creation; they take instructions in natural language and don't need even a line of code, reducing reliance on IT and fostering a more patient-centric approach to public health services.

2

Application in Public Sector Health

02

Foundational Activities to Task Agents With

With their unparalleled flexibility and intelligence, Agents transform how the public sector operates. As with any new capability, it's best to start with discrete use cases with proven patterns. Let’s dive into foundational use cases where Agents are already having impact. Consider: “Which activities within my organisation align with these use cases?”

Knowledge Navigator

What It Is

Agents act as intelligent guides, navigating vast knowledge repositories to surface concise summaries from natural language queries. This empowers patients and staff to self-serve, finding accurate information without asking a human.

What It Does

- Simplifies access to information stored across various repositories
- Makes knowledge accessible through natural language to overcome terminology discrepancies.
- Enables users to self-serve rather than manually asking through phone or email channels.
- Securely verify identity to improve accuracy and contextualise responses.

Measures of Success

- Reduce query resolution time
- Increase self-service success rates
- Reduced escalation to human ratio

In-Context Assistant

What It Is

Agents can be seamlessly embedded within existing digital workflows, offering real-time support and insights as users navigate complex processes. They act as 'digital assistants,' ensuring accuracy and efficiency in decision-making.

What It Does

- Provides contextual assistance, such as policy guidance or compliance checks, while users complete forms or processes.
- Acts as an "AI assistant", ensuring accuracy and efficiency in real-time decision-making
- Only uses Knowledge from approved organisational sources.

Measures of Success

- Reduction in errors in workflow processes
- Time saved per task or decision point
- User satisfaction with in-process assistance

Handle Routine Task

What It Is

Agents take over repetitive, time-consuming tasks, such as scheduling appointments, managing emails, or conducting eligibility checks. This frees up human employees to focus on strategic, patient-centric activities.

What It Does

- Streamlines routine workflows, freeing up staff for higher-value activities.
- Handles predictable and low-complexity tasks with high accuracy.
- Works with a high tolerance of ambiguity; through rules decides the best course of action in e.g. schedule conflicts

Measures of Success

- Reduction in average handling time for routine tasks.
- % of tasks successfully completed by the agent.
- Increase in staff productivity, satisfaction & retention.

02

Foundational Activities to Task Agents With (Con't)

Personalised Communicator

What It Is

Agents act as expert communicators, crafting personalised content tailored to individual recipients and their specific circumstances. This ensures patients receive relevant information and guidance in a timely and engaging manner.

What It Does

- Drafts communications, such as an email, using department standards.
- Understands context of the recipient within a process to personalise the message.
- Allows the human in the loop to adjust content to taste.
- Automate generation of content in 1:1 or 1:many messaging.

Measures of Success

- Decrease in case handling time.
- Reduction in errors in workflow processes.

Document Sentinel

What It Is

For inbound documents, an Agent that performs an initial scan of content to check for completeness, like identification information or date validity.

What It Does

- Ingests content sent in emails, documents or PDFs.
- Checks against a set of rules defined by the department on completeness.
- Highlights to the user where information may be incomplete.

Measures of Success

- Decrease in case handling time.
- Decrease in case resolution time.
- Decrease in routine quality checks.

Keep it Discrete

These use cases are often part of a larger, multi-step process. It can be tempting to design a single Agent to handle every step. Best practice is to create Agents with discrete roles, and pick the right Agent for each step of the process.

An Agent is interoperable with other Agents; that means an Agent you create to handle (for example) inbound document checks can 'transfer' the conversation to the in context assistant when necessary, seamlessly to the user.

Assigning discrete use cases to each Agent make them straightforward to iterate as you observe user behaviour.

Better Productivity

62%

of the workday is lost on repetitive, manual tasks*

These are just a few examples of how Agents are transforming government operations. By embracing this technology, organisations can unlock new levels of efficiency, improve patient satisfaction, and empower staff to focus on higher-value work.

Always consider the ethical implications of the tasks you plan to assign to Agents. Summarisation of clinical information needs to carefully reviewed and should only be evaluated after lower risk scenarios have been used in production and your organisation is comfortable with the capabilities of Agentic AI. Auditing and monitoring is an essential part of using AI in call circumstances.

*Source: State of Service, Salesforce, April 2024.

02

Foundational Activities to Task Agents With (Con't)

Having explored the potential of Agents, it's time to turn vision into action. Using a structured approach aligned with the **UK Government's AI Opportunities Action Plan**–SCAN, PILOT, SCALE–your department can confidently embrace Agents and deliver measurable impact.

Scan
Identify Applications of Agents

Deep Dive into Applications

- Explore how Agents can enhance your department's work, from internal efficiency gains to improved patient experiences.
- Assess both internal applications (e.g., staff productivity) and external applications (e.g., patient engagement).

Define the Challenge

- Articulate the key problems your department faces, identifying target audiences and desired outcomes.
- Consider specific behaviours to change and define the KPIs that will measure success.

Assess Value, Risk & Complexity

- Use a structured approach to identify high-impact, low-risk opportunities for pilots.
- Focus on areas where AI can deliver immediate efficiency gains or address critical bottlenecks.

Leverage Power Users & DDaT Community

- Parts of your organisation may be more familiar with AI tools, using them in ways you may not be aware of. Leverage this knowledge from early adopters to inform your thinking.

Pilot
Experimenting with Purpose

Don't DIY Your AI

- Avoid the pitfalls of building in isolation or relying solely on tools already in use. Partner with expert organisations that provide low barriers to entry and exit.
- Embrace the one-way door vs. two-way door analogy: Choose pilots that can easily scale if successful but allow you to pivot or stop without significant risk.

Educate & Engage

- Proactively communicate the benefits of Agents to both internal teams and the public, addressing concerns and fostering trust.
- Use pilots as an opportunity to build understanding and excitement across your organisation.

Leverage Experimentation & Innovation Agreements

- The message is clear, do not let lengthy procurements get in the way of innovation, nor should you rely on existing AI tools only.

Treat AI Like a New Hire

- Think of your Agent as a new team member: Don't expect perfection from day one. Both your organisation and the AI will learn and adapt over time.

Scale
Delivering Impact Across Systems

Evaluate Results

- Analyse pilot outcomes against predefined metrics, such as query resolution rates, user satisfaction, and cost efficiencies.
- Use the data to identify areas for improvement and refine your approach.

Iterate & Expand

- Continuously improve the agent based on pilot feedback.
- Decide whether to scale the solution, pivot to a new use case, or retire the project.
- Extend success by applying the agent to new cohorts, use cases, or broader capabilities.

Think Big but Stay Modular

- Develop or procure a scalable AI tech stack capable of supporting multiple agents and millions of patient interactions.
- Ensure solutions are open, interoperable, and reusable, allowing for standardisation across departments.

Centralise Support

- Avoid fragmentation by providing centralised funding, leadership support, and scaling services for successful pilots.

02

Choosing the Right Experiment

When introducing new technology, it's important to select the right use cases for experimentation. The best experiments meet clear, objective criteria, minimising bureaucratic or political influences on decision-making. This focuses outcomes on high-impact opportunities.

In the spirit of innovation, embrace and learn from failures. Experiments that don't deliver predicted results still provide valuable insights, refining strategies for future success.

To maximise impact, prioritise quick, lightweight experiments over resource-intensive initiatives. Targeted use cases set the stage for meaningful progress without overburdening teams or budgets.

Here we explore characteristics for selecting the right targets for experimentation.

Isolated

The use case is standalone, without meaningful dependencies and/or potential impacts to other services.

Simple

Relatively easy to implement, avoiding complex contexts (e.g. many decision points, excessive data sources).

Scalable

Those that are repeatable, enhanceable, or extendable to a broader cohort if successful.

Narrow

Just big enough to serve as a meaningful test, but small enough to be manageable.

Measurable

A service with a clear measurable outcome, and success criteria agreed with stakeholders.

Low Risk

Focus on marginal services or sub-processes to learn lessons outside of a high risk service.

Informative

Helps develop understanding of Agentic AI, furthering the organisation's knowledge.

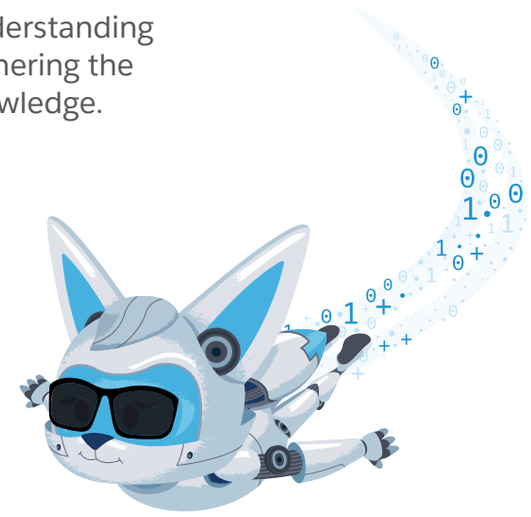

Comfort With Gen AI Use Cases is Net Positive

Comfort With the Following Generative AI Use Cases

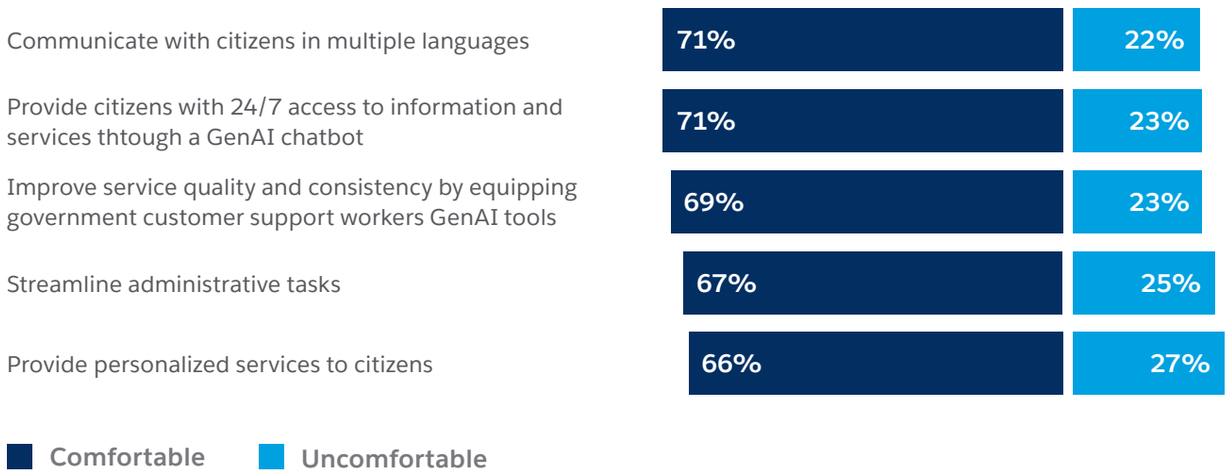

Source: 2024 Global Report: Trust Imperative 4.0, BCG, April 2024.

3

Launching Agents in Your Organisation

03

Building Agents

Recap on what we've covered

In the previous chapters, we explored what Agents are, their transformative potential in the Public Sector Health, and the framework for identifying the best use cases. By now, you understand the foundational principles of Agents and their role in driving efficiency, innovation, and better patient outcomes.

This chapter focuses on how to bring Agents into your organisation. We'll guide you through:

- **Building Agents:** A new approach to development that prioritises collaboration and flexibility.
- **Testing and Iteration:** Using lightweight experiments to refine your Agents to meet your needs.
- **Scaling with Confidence:** Expanding successful pilots into impactful, scalable solutions.

Whether you're building your first Agent or refining an existing one, this chapter equips you with the tools and strategies to onboard effectively, aligning with your organisation's goals.

Delivering Change, the Agentic Way

Traditional software development is often siloed and sequential, involving multiple specialist teams:

- **Business Analysts** conduct user research, creating user stories that define requirements.
- **Developers** build the capability based on those requirements.
- **Testers** then ensure the product is ready for production.

This structured approach is designed to clearly demarcate responsibilities, create predictability, and allow progress tracking. It also introduces inefficiencies:

- Work is fragmented across teams, creating barriers to communication.
- Quality issues emerge late in the development process when fixes are costly and time-consuming.
- Without a shared understanding of the software product's purpose, clarity is only achieved during final testing stages.

To address this, organisations have embraced the principle of 'shifting left'–identifying and resolving quality issues earlier in the development lifecycle, where fixes are faster and less expensive. Agentic technology fits well with this principle.

03

Building Agents

A New Way: Building with Agents

Agents revolutionise this process by breaking down barriers and enabling collaboration across disciplines. Instead of interpreting user stores and transposing them in to code, Agents are built using natural language instructions. This is an example of how to create a instruction for an Agent:

```
Your job is to provide guidance on quit smoking
programs. You can answer questions about the
program, but you'll need to pass them a link to
the Enrol Online' form if they wish to enrol.
```

This fundamentally changes how development happens:

- **Simplified Articulation:** There's minimal heavy lifting required to transpose requirements into working software.
- **Early Collaboration:** Business stakeholders review and refine natural language instructions as soon as they're written, avoiding misunderstandings only uncovered during formal testing.
- **Clear Scope:** The Agent's responsibilities are clearly defined, and any adjustments are made early without costly rework.

Principles for Building Agents

These principles enable you to build differently with Agents:

Don't Boil the Ocean

Agents are most effective as discrete, composable services. Start small and focus on specific tasks or processes. Domain driven design is a great technique to make services composable.

Build Together

Collaborate early and often. Instead of weeks of back-and-forth in a tracking tool, write and refine instructions together with business stakeholders. Seek out collaboration tools that do not confine you to singular AI models and technologies.

Involve the Business Early

With natural language at the core, "business vs IT" becomes "business with IT," fostering a shared understanding of what the Agent will do.

General AI Metrics

29%

INCREASE IN PRODUCTIVITY DUE TO PREDICTIVE AI

(% time saved, e.g. AI that uses data to anticipate outcomes; examples: Einstein reply recommendations, Einstein lead scoring)

29%

INCREASE IN PRODUCTIVITY DUE TO GENERATIVE AI

(% time saved, e.g. AI that generates content autonomously using large language models)

30%

INCREASE IN ACCURACY OF AI CONTENT GENERATED BY SALESFORCE

(due to grounding in your business data compared to general purpose tools like ChatGPT, % improvement in accuracy)

Source: Salesforce's FY25 Customer Success Metrics - AI Reports

The Result

By adopting this collaborative, iterative approach, you unlock faster development, reduce costs, and create Agents that align with organisational needs. It's not just a different way to build software–it's a better way.

Next, we unpack measuring and iterating Agents.

03

Evaluating Agents

Building an Agent is just the beginning–iteration is where the real impact is made. Testing your Agent before releasing it to users is critical (and there are tools that can help like Salesforce Agent Testing Center), but no simulation fully replicates real-world scenarios. The good news? Agents are highly adaptable and are refined quickly and safely.

Key Tools for Measuring & Iterating Agents

Agents provide auditable logs, just like traditional technology solutions. These logs are invaluable for monitoring and improving performance:

Feedback

- User input, like a 'thumbs up' or 'thumbs down,' reveal what's working and what's not.

Question to Ask: Which journeys are receiving the most negative feedback, and why?

Intent

- Tracks the Agent's understanding of user instructions.

Action Point: Ensure the correct intents are being identified. If no intent is found, refine the Agent's instructions or training.

Build Together

- The discrete use cases the Agent identifies to complete a request (e.g., "Inbound Document Check").

Value: Shows which journeys are popular and highlights areas needing improvement to meet user needs.

Model Token Usage

- Tracks the number of tokens consumed by generative AI during specific journeys.

Optimisation Opportunity: Identify high-consumption areas and refine instructions to reduce usage without sacrificing quality.

Service Metrics

30%

INCREASE IN AGENT & SUPERVISOR PRODUCTIVITY DUE TO INVESTMENT IN AI
(% increase in work accomplished by agents & supervisors in a given week)

30%

CURRENT % OF CASES OR CUSTOMER ISSUES SOLVED BY USING AI
(% of cases)

Source: Salesforce's FY25 Customer Success Metrics - AI Reports

Measuring Success Against KPIs

Regularly monitor performance against the KPIs established during the planning phase. Combined with audit logs, these metrics offer a comprehensive view of how well the Agent is meeting its goals. Examples include:

- Query resolution rates
- User satisfaction scores
- Cost efficiency improvements

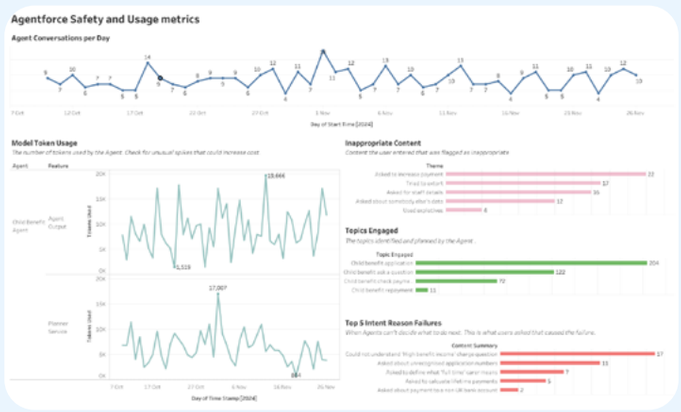

03

Scaling Agents

Change Management for Agents

Agents aren't bound by traditional IT change management processes. They are a decoupled layer across the top of existing data and systems, rather than being tightly integrated. As they're composable and self-contained, improvements are made quickly without exhaustive regression testing.

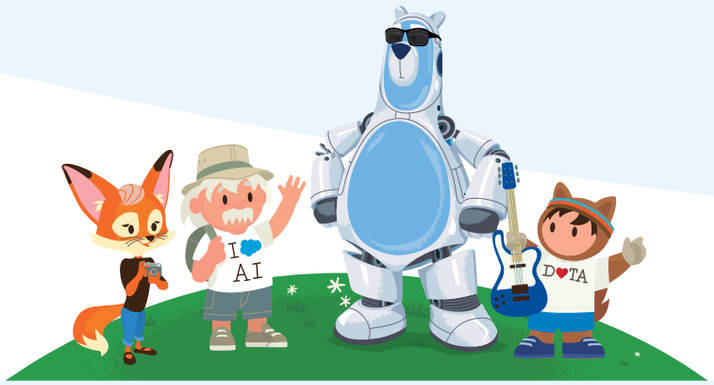

Only 15% of GenAI Experiments Scale Up

Source: Forrester's Priorities Survey, 2024

The Power of Iteration

A successful Agent isn't static–it evolves from user feedback, real-world data, and changing needs. Regular iteration means your agent continuously improves, adapting to user needs.

With the speed and simplicity of Agents, iteration becomes an opportunity to refine, innovate, and achieve better results every day.

How to Adapt Your Processes:

- **Simplify change management and release cycles;** educate teams on how Agents are loosely coupled, de-risking traditional IT change.
- **Act swiftly on real-world feedback to optimise Agent performance;** this is sometimes deploying a simple rewording of an instruction.
- **Ensure regular iterations to maximise effectiveness against KPIs;** make incremental change regularly to identify levers for performance.

What 'Good' looks like

Within this chapter, you've seen why building Agents requires a different approach to unlock benefits. As you build in your organisation, test for:

- Composable, discrete services over monolithic processes.
- No barriers between business analysts and build resources.
- 'Day 1' review of instructions with the business.
- Observability by default; if you can't measure it, don't build it.
- Iterations within minutes and hours, not days and weeks.

In the next Chapter, we'll hone in on how to understand, assess and measure the value, risk and complexity of AI use cases within the wider context of your organisation. Let's get started.

4

Value, Risk, and Complexity Factors

04

Value Types and Levers

Classify & Quantify Value

Every organisation has a unique value framework shaped by its Mission, Goals, Priorities, and the Risks it manages. Despite this uniqueness, the way value is measured often follows consistent principles, categorised into Quantitative and Qualitative dimensions:

Qualitative Value

This refers to measurable, monetary impacts on the organisation.
It can be further divided into:

- **Cashable Benefits:** Direct financial savings, such as cost avoidance (e.g., reduced legacy system expenses) or headcount reductions..
- **Non-Cashable Benefits:** Indirect gains that improve efficiency, enabling the organisation to redeploy resources. Examples include:
  - **Efficiency Gains:** Improved employee productivity through automation.
  - **Humanistic & Societal Benefits:** Expanding inclusivity and accessibility through innovations like AI-driven translation services or tailored patient interactions, ensuring equitable service delivery for all demographics

Quantitative Value

This focuses on less tangible outcomes that drive mission delivery but may not immediately reflect financial ROI.

- **Patient Satisfaction:** Enhanced user experiences through faster, more personalised interactions.
- **Improved Trust:** Transparent and reliable service delivery builds confidence in public services.
- **Humanistic and Societal Benefits:**
  - Example: Agents can guide patients, book appointments, and offer 24/7 support. This includes providing comprehensive information about available services, healthcare locations, and appointment options, ensuring patients can easily navigate the system and access the care they need.

As always a strong connection between Agents and Business Intelligence is paramount.

Value Planning Parameters

When evaluating Agentic opportunities, it's essential to use a comparative framework. Key parameters include:

Investment Amount

The total funding required to achieve the target outcome, i.e., the cost of the project.

Return on Investment (ROI)

The financial or operational benefit produced by the investment, often measured over a 5-year period.

Time-to-Value (TTV)

How quickly the organisation realises value after implementation. This includes the **Break-Even Point,** where initial costs are recouped, and the ROI turns positive.

Risk

Encompasses **complexity** and **dependencies**, as well as potential risks that could drive up costs, delay delivery, or impact outcomes negatively (e.g., **externalities** like unforeseen policy changes).

04

Three Major Value Methodologies

Prioritise Use Cases & Building an Investment Backlog

Now that we have aligned on how to classify and measure value, how do we prioritise use cases for Agents in a singular value framework? Consider using one of these frameworks to assess value.

Outcome-Oriented

What It Is

Not all organisations are driven by cashable value. Especially in public sector, certain mission objectives or intangible values take priority. An outcome-oriented prioritisation defines one or several goals to optimise with Agents. These could be a combination of the outcomes mentioned earlier:

- Reduce, Eliminate Backlogs
- Scale Resources, Maximise Impact
- Increase Pace of Delivery

Or they could be organisation- specific, aligned to strategic goals or targeting critical pain points.

When to Use It

Apply to address critical challenges or objectives, particularly when they cannot be directly quantified in monetary terms.

Net Present Value

What It Is

The typical approach for “apples- to-apples” comparisons of investment alternatives is with Net Present Value. Decision-makers assess the options before them in exclusively monetary terms, including time considerations, with the Discount Rate serving as a proxy for organisational urgency.

Risk too can be quantified and incorporated into ROI figures as Expected Value (i.e. likelihood x projected loss).

When to Use It

The go-to framework, used to evaluate investment options with limited consideration for skews in any one parameter.

Capex-Constrained

What It Is

Public sector organisations are known to operate under tight budgets– especially now. Meaningful opex savings that might be achieved with enhancements to process or technology are often out of reach, with Capex funds being unavailable.

We can instead rank-order projects based on a composite of required investment, ROI, and TTV. Find an affordable first project with a high relative Opex ROI, and quick TTV. Resulting savings can fund a slightly larger second project, and so on.

When to Use It

Apply this approach when investment funds are severely limited and a reduction in operating costs is critical to the organisation’s going concern.

04

Build Your Own Value Framework and Backlog

A Hybrid, Dynamic Value Model to Align AI Maturity & Value

Identifying the right use cases for Agents is only the beginning. To create meaningful impact, organisations must adopt a dynamic and flexible value framework that evolves alongside their AI maturity and strategic goals. Prioritisation isn't about a fixed roadmap–it's about balancing short-term wins with long-term transformation.

What 'Good' looks like

Balance Value and Risk

High ROI doesn't always mean low risk. Identify projects that balance immediate returns with achievable outcomes.

Speed and Scalability

Prioritise opportunities where the solution can deliver value quickly and scale across departments or services.

Focus on Tangible Impact

Ensure each initiative improves the experience in a clear way or delivers mission-critical results, not just operational efficiencies.

Engage Stakeholders Early

Bring operational teams, finance, and leadership into the value discussion early to align on goals and expectations.

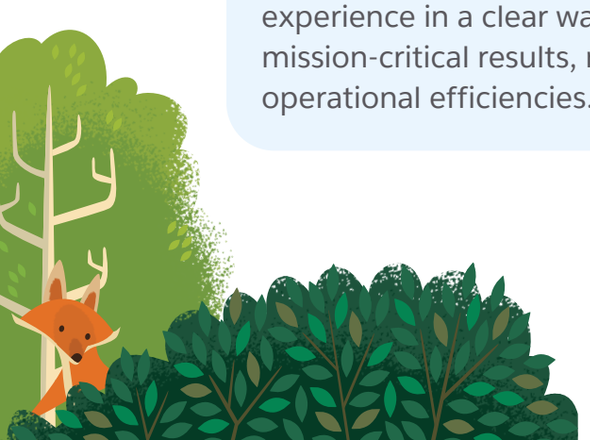

Visualise Investment Opportunities

To manage competing priorities, visualisation tools and comparative frameworks are essential. These help assess use cases based on multiple factors, such as ROI, TTV, and risk, revealing short-term opportunities while laying the groundwork for scale.

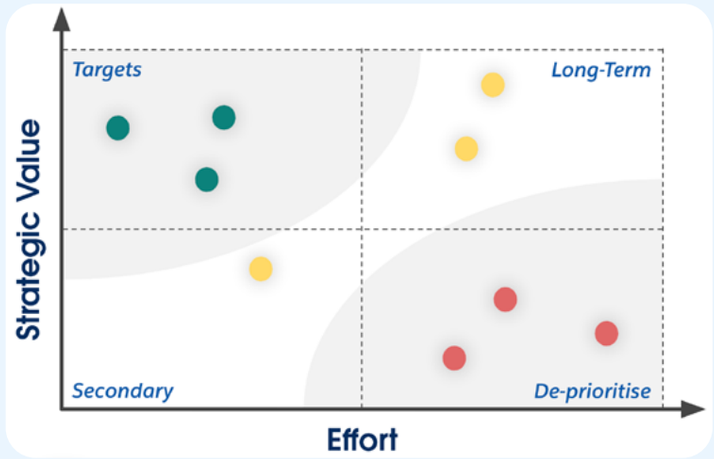

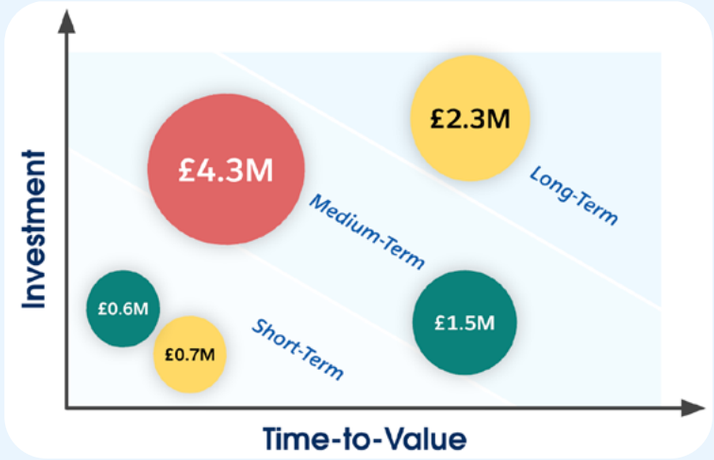

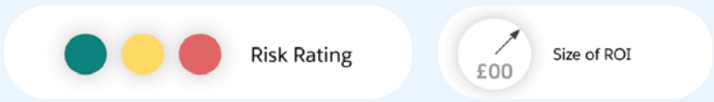

04

Include Risks As Part of Your Value Framework

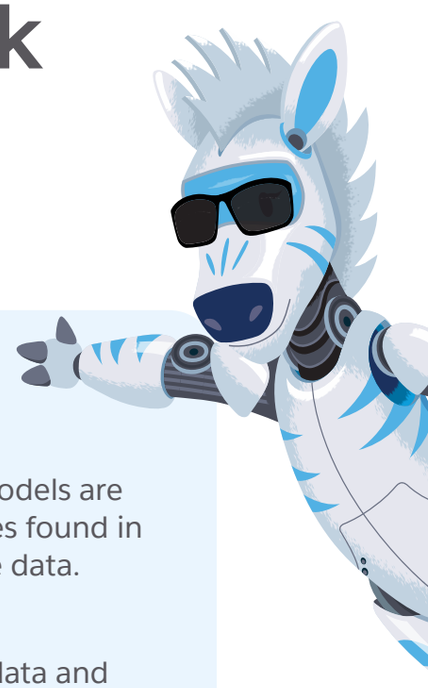

When assessing the benefits of Agents, we must consider risks. Agentic AI inherits the risks of generative AI, potentially deployed at scale to interact with external stakeholders directly. While the risks are similar to Generative AI, the way they manifest can vary:

Toxicity

What It Is

Use of inappropriate language or otherwise offensive statements by the AI, learned from training data.

Proposed Mitigation

A second model to hunt for toxicity in responses, applying methods such as Salesforce's Trust Layer.

Hallucinations

What It Is

Generative AI by definition creates new content, sometimes going too far with "new facts".

Proposed Mitigation

A second model to challenge the veracity of the primary model's responses, applying methods such as a **Generative Adversarial Network (GAN).**

Bias

What It Is

Transformers and other AI models are susceptible to the same biases found in their training data and source data.

Proposed Mitigation

Careful selection of training data and a sustained effort to tune the model to limit biases.

Generalisation

What It Is

Without the full context of a customer's question, Agents may provide an answer that is inaccurate or incomplete.

Proposed Mitigation

Detailed specificity of Agents that consider policy pathways and special scenarios.

Missed Expectations

What It Is

Agentic AI has tremendous potential, but applied to the right use cases and executed effectively.

Proposed Mitigation

Scale your use of the technology incrementally, assessing value at each step and adjusting your approach based on outcomes and learnings.

Cost

What It Is

Agents and underlying capabilities use a consumption-based pricing model. Anticipating usage volumes correctly is important.

Proposed Mitigation

Start with limited contexts and scale up from there, ensuring costs are in line with expectations.

What 'Good' Looks Like

To minimise risk and maximise value, focus on low-impact use cases where errors have minimal consequences. These pilots let you refine your Agents and build organisational confidence. As you consider use cases, remember the one-way door vs. two-way door analogy: prioritise reversible, flexible decisions that allow for quick iteration and learning. Pair this with regular monitoring and human oversight, especially for sensitive topics, to ensure trust and reliability.

5

Data and Organisational Readiness

05

How do You Operationalise Agents?

Building on our understanding of risks and mitigations, the next step is operationalising Agents effectively. To achieve this, organisations need a strong foundation–not just great technology, but also integrated data, clear processes, and organisational readiness.

Successfully scaling Agents requires more than just great technology–it depends on having a robust foundation of data and organisational readiness. Agents thrive on high-quality, well-integrated data, and as use cases expand, so does the demand for diverse and accurate datasets. This is where a **Digital Operating Model (DOM)** becomes critical, enabling organisations to support and scale agentic AI effectively.

Guiding Principles

The ethical and practical principles that will shape the strategy and overall usage of AI (especially focused on the more sophisticated flavours such as agentic). Many government and academic organisations have published high quality, thoughtful versions that can serve as inspiration for your own. As you experiment and adopt AI, you should revise and elaborate accordingly.

Maturity Model

A plan for how to evolve organisational competencies and usage of agents, ensuring that all legs of the stool evolve at corresponding pace. The model should describe the state of the organisation in each of the domains and include high level goals for each step. This is distinct from the technology and use case roadmap that should flow down from the maturity model.

Steering and Governance

The organisational framework for managing the strategy, guiding principles, roadmap, and other core aspects.
As partnership between business and IT, the steering and governance pillar defines working groups and procedures for maintaining and evolving use of agents in the organisation.

Additional elements should be incorporated on top of this foundation. These are detailed in later sections of the handbook, and include:

Value Framework

Define the factors that will drive investment decisions and roadmap priorities, including weighting of different value types and parameters.

Change Management Plan

How you engage the wider organisation, including a communication plan, ways to evangelise progress and benefits of Agents, and feedback mechanisms.

Skills Development Plan

Everyone should learn how to use AI tools of various kinds, with deeper skills needed for those involved in the planning and rollout of Agents.

Value Realisation Plan

Risk and value are two sides of the same coin and should be considered together. Each organisation has its own set of value and risk levers, and requires a plan to manage them.

Roadmap & Backlog

The order in which technical capabilities and Agent use cases will be delivered in your organisation, including timings. This should be revisited and adjusted periodically.

05

Data Readiness

The Importance of Data Connectivity

Agents offer innovation and unprecedented pace. Like technological innovations that preceded it, they're only as good as the data you use to train and inform them. Organisational data enhances Agents in two ways:

1. **It grounds generative AI:** it connects model output to verifiable sources of information. By doing so, it increases the accuracy of the model output, as you're signposting the Agent to your trusted content to provide a given response.
2. **It takes action:** it can update data sources within your organisation. By taking action, you turn a conversational dialogue with an Agent in to activity as part of a workflow.

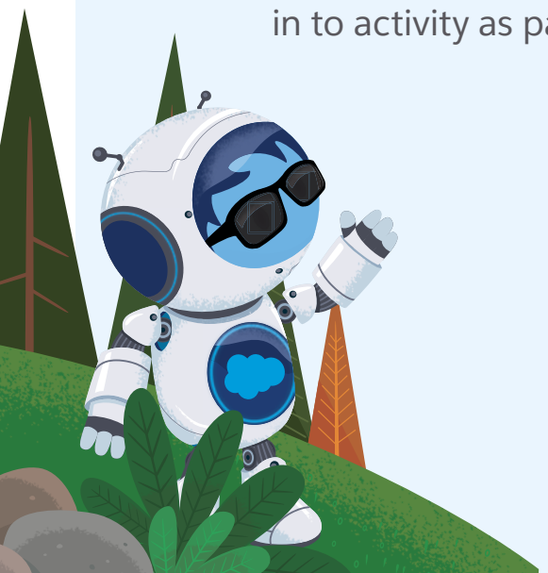

Preparing your Data for Agents

Agents can work with structured, semi-structured or unstructured data, making it quicker than traditional IT to import data libraries. Data is often stored in unstructured text, for example in policy documents. Agents support Retrieval Augmented Generation, so particular phrases or search terms can be returned from an unstructured document.

To keep conversations relevant and performant, identify the right data sources for each discrete Agent use case. Ideally API services should be developed with a composable API-first approach, and that's perfect for Agents where API responses can be interpreted autonomously. Healthcare organisations often provide libraries of APIs that can access health identifiers, services directories etc.

Principles for Success

- **Make datasets discrete, avoiding overlaps:** Agents can make intelligent decisions about which dataset to use, but you may not get results you expect if the Agent discovers similar data across multiple datasets.
- **Add context and metadata:** labelling data sources clearly, with the taxonomy you use in your Agent, gives the Agent an index to navigate to the right data set.
- **Respect data provenance:** Agents can authenticate users just like a traditional online service. Ensure this user context is represented in systems of record as the Agent takes action on data within a workflow.

6

Integrating Your Hybrid Workforce

06

Addressing Internal Fear and Resistance

When introducing AI, concerns like job loss, lack of control, ethical risks, and cultural resistance are common. Addressing these head-on is critical to fostering adoption.

Understanding the Fear

Introducing AI can bring resistance due to concerns over job security, trust, and change. Addressing these upfront builds the foundation for successful adoption.

- **Job Loss:** Fear that AI will replace human jobs.
- **Lack of Control:** Concerns about decisions made without human oversight.
- **Ethical Risks:** Worries about bias, privacy, and misuse of AI.

Overcoming the Fear

Transparency, training, and trust are key to overcoming resistance and ensuring AI adoption succeeds.

- **Transparent Communication:** Clearly articulate AI benefits (efficiency, accuracy), involve stakeholders, and provide regular updates.
- **Employee Training:** Upskill staff for AI tools, create AI-centric roles, and highlight how AI enables more fulfilling work.
- **Ethical Frameworks:** Develop ethical guidelines, ensure compliance, and mitigate risks like bias or misuse.

Bringing Your Team With You

AI is not just a technological change–it's a cultural shift. To build trust and excitement, ensure your team feels part of the journey

- **Start Small and Build Confidence:** Showcase quick wins through low-risk pilots to demonstrate the benefits.
- **Learn Together:** AI is like hiring a new team member; it learns and improves with feedback. Involve your team in refining and personalising the solution.
- **Create Shared Wins:** Celebrate milestones with your team to highlight how AI supports–not replaces–their expertise.

06

Addressing Internal Fear and Resistance

Safeguards & Guardrails - Keeping AI in Check

Adopting AI offers transformative opportunities, but it's essential to mitigate risks with robust safeguards. While some fears around AI are exaggerated, building trust and ensuring compliance are paramount.

- **Clear Policies:** AI Agents will streamline processes, personalise patient experiences, and deliver services more efficiently than ever.
- **Data Security & Transparence:** Select tools with a robust Trust Layer and built-in Testing Centres to protect sensitive data and maintain public trust. Ensure compliance with the national legislation and for EU member states the EU AI Act.
- **Preventing Shadow IT:** Centrally manage AI deployments to avoid unauthorised use, maintain control, and ensure alignment with organisational goals and compliance standards.

Strong guardrails not only minimise risks but also unlock the full potential of AI, delivering measurable value and building confidence in the technology.

Delivery & Operational Challenges

Scaling AI isn't just about technology–it's about empowering the organisation while maintaining IT oversight.

- **Backlogs in IT Resources:** With IT teams already stretched, ensure projects are prioritised based on value, risk, and complexity.
- **Empower Business Stakeholders:** Leverage NLP tools to enable non-technical teams to deploy and manage AI agents. These tools allow business users to iterate quickly on solutions while ensuring IT retains visibility and control over deployments.
- **Seamless Governance:** Implement systems that enable IT teams to oversee AI usage, track updates, and ensure compliance with policies without stifling innovation from other departments.

This balanced approach empowers business teams to innovate while maintaining the rigorous control and governance that IT demands.

7

Scaling Agents and Business Impact

07

Multiple Lenses to for Agentic AI Maturity

Scaling the use and impact of your new Agent workforce requires progress across different dimensions. They progress independently as distinct dimensions.

Agent Worker: Individual Scope & Skills

Train your AI agent to act comprehensively and impactfully

Knowledge Navigator

Address queries through natural language from selected content sources

Personalised Concierge

Increase depth and accuracy by including individual context about the enquirer

In-Context Assistant

Real-time support and insights as users navigate complex processes

Routine Taskmaster

Automates routine workflows, freeing up staff for higher-value activities

Process Orchestrator

Handles complex processes spanning multiple systems and/or organisations

Agent Workforce: Adoption & Breadth of Use

Extend your agent workforce to different parts of the organisation, incrementally increasing breadth and depth

Experimentation

Initial assessment of technologies and pilots for limited use cases

Internal Validation

Extend successful pilots and introduce new use cases, limiting to internal users

Multi-Stakeholder

Expand use to external users (e.g. patient) with broader use cases

Cross-Agency Integration

Include other Agencies (e.g. Arms Length Bodies) for intelligent collaboration

Complex Orchestration

Many Agents orchestrate processes between patient, Agencies and the private sector

Agent-Powered Organisation: Readiness & Competencies

Enablers to maximise value from your combined Agent & human workforce

Desiring

Internal dialogue around agentic AI, an exploration strategy & guiding principles

Storming

Launch change management initiative and develop a value framework & backlog

Norming

Establish an Agentic Centre of Excellence to manage value, risk, and operations

Performing

Incorporate further elements of Digital Operating Model to support agents at scale

Leading

Optimise programmatic and governance elements to further drive impact

Conclusion

Harnessing the Power of Agents for Public Services

Africa today has very few telephone landlines, it jumped straight to mobile technology. Public Sector Healthcare is a laggard and the last major part of the sector to digitally transform to a patient and workforce centre instead of focused on the test results or episode of care. Public Sector Health is at that ‘Africa’ moment. The pressure on resources, budgets patient demand is putting an unbearable pressure on the system and its ability to deliver services it must transformation, it's no longer a nice to have. But to utilise the full potential of Agentic AI to lead that transformation Public Sector Healthcare organisation have to breakdown the data silos and change from legacy, customised and verticalise point solutions to platforms with open data and integration standards that use reusable components

We understand that no single AI solution fits all needs. Different AI models have distinct strengths and weaknesses, making the "Scan" phase a critical opportunity–not just for applying AI tools to existing problems, but for exploring the broader AI landscape. This includes evaluating the unique strengths of each technology, their suitability for domain-specific use cases, and their capacity to scale across an organisation. By embracing this approach, organisations can ensure they deploy the right tools in the right contexts for maximum impact.

For the Public Sector and beyond, success starts with targeting high-value, low-risk use cases that deliver measurable outcomes while building confidence in AI adoption. The frameworks and insights we’ve shared are designed to guide you in prioritising domain-specific opportunities, ensuring clear ROI, and fostering trust through transparency and collaboration. By starting small and scaling boldly, organisations can create a legacy of efficiency, equity, and innovation.

Get Started Today: Build your first AI agents and experience their potential **https://trailhead.salesforce.com/content/learn/projects/quick-start-build-your-first-agent-with-agentforce**

The time is now to shape the future of your organisation–and the UK at large. With AI agents, we can reimagine what’s possible. Start small, learn fast, and scale boldly. Together, let’s make the vision a reality.

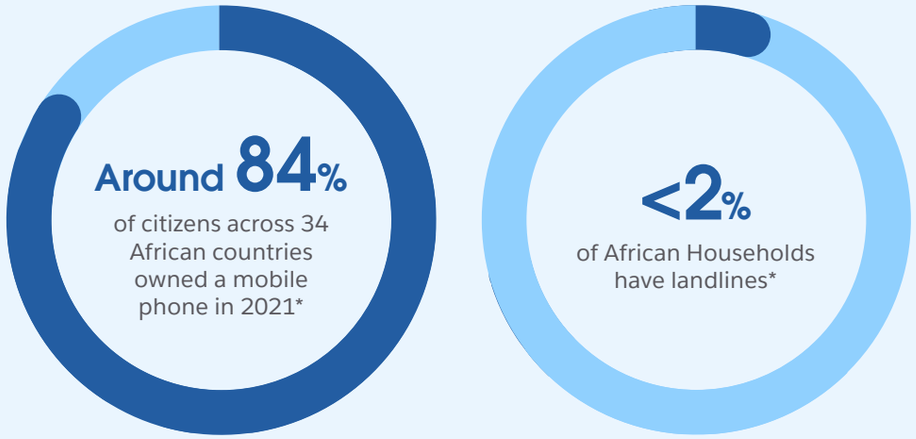

*Source: World Bank "Digital Transformation Drives Development in Africa" (January 2024)