

# As a “Data Platform Vendor” how do we approach value engineering today in the context of agentic AI, and how does it affect the value proposition?

[Executive Summary](#)

[Context and assumptions](#)

[What is Agentic AI?](#)

[Impact of Agentic AI](#)

[High Level Challenges of Agentic AI](#)

[Agentic AI is heavily bubbled](#)

[People are under huge pressure to ‘do something’](#)

[The economics of Agentic AI are fragile](#)

[Agentic AI is opaque](#)

[Agentic AI struggles with real world, production data sets](#)

[Agentic AI - A Draft Plan...](#)

[Pre-engagement qualification checklist](#)

[Where do Data Platform vendors fit in?](#)

[General AAI state storage](#)

[Use Case: A ‘Single Store’ of corporate data / Customer 360](#)

[Use Case: Command, control and auditing of Agentic AI](#)

[Use Case: LLM Policy and Charging](#)

[Domain specific use cases](#)

[Use Case: Fraud Detection](#)

[Use Case: Hyper-Personalisation](#)

[Use Case: Video Events Mediation and Decisioning](#)

[Operational Challenges of Agentic AAI](#)

[Brownfield deployments - lower risk, lower reward](#)

[Greenfield deployments - higher risk, higher reward](#)

[We can’t lose our focus on developers](#)

## Executive Summary

Agentic AI (‘AAI’) offers Data Platform vendors a significant opportunity to create and extend footholds in large Fortune 500 companies. Data Platform vendors are naturally suited to AI use cases, with the four main uses being: Storing conversation state, providing a ‘customer 360’ view of data, vector storage and allowing AAI systems to be monitored for failures or security issues.

However, this is not a case of turning up and selling. The entire AI sector is in the middle of a bubble, and the majority of AAI projects will fail. So while we are all in favor of AAI, we need to try to avoid doomed projects and focus on our long term relationships with prospects and customers, even if that means being cautious and qualifying out projects.

## Context and assumptions

- We need to work within a 5 year timescale. This is not about 'coining it' in 2026, only to have nobody take our calls in 2027. We need a plan which will survive the AI bubble bursting.
- We are at the end of the 'new database' VC cycle. The industry is consolidating. We want to be in the same position in 2031 that Oracle were in 1999. We need to be the 'obvious choice', and never a candidate for replacement.
- As part of this strategy we also need to have layers of advocates at multiple levels inside major existing customers and prospects.
- Agentic AI is of huge importance right now, but it isn't the only game in town. It's important that we devise a strategy that is focused on long term relationships with Fortune 500 companies, not just individual deals with an AI angle.

## What is Agentic AI?

At the moment there is so much hype it's sometimes hard to tell what's being talked about, but according to the Harvard Business Review<sup>1</sup>:

*AI agents are a class of intelligent agents distinguished by their ability to operate autonomously in complex environments. Agentic AI tools prioritize decision-making over content creation and do not require human prompts or continuous oversight.*

The 'magic' behind them is that they can take a text or even verbal input and perform simple tasks that would normally be done by entry level employees, such as taking a booking, finding out where a shipment went, or what someone's balance is. They do this by feeding a verbal prompt to an LLM, and giving it one or more tools to solve a problem. A sample prompt<sup>2</sup> might be:

*You are an expert weather forecaster, who speaks in puns.*

*You have access to two tools:*

- *get\_weather\_for\_location: use this to get the weather for a specific location*
- *get\_user\_location: use this to get the user's location*

<sup>1</sup> <https://hbr.org/2024/12/what-is-agentic-ai-and-how-will-it-change-work>

<sup>2</sup> From <https://docs.langchain.com/oss/python/langchain/quickstart>

*If a user asks you for the weather, make sure you know the location. If you can tell from the question that they mean wherever they are, use the `get_user_location` tool to find their location.*

Note that these are instructions for a piece of software, not a human. The LLM then uses callbacks into the tools to gather data and attempt to solve the problem. Much of entry level work is following simple processes. Since AAI seems to be really good at this, the obvious temptation is to deploy it everywhere and cut staff.

## Impact of Agentic AI

The headline impact of Agentic AI is that of widespread job losses. The reality appears to be far more nuanced, with both the BBC<sup>3</sup> and RAND<sup>4</sup> questioning whether the job losses are in fact AI related, as opposed to a hangover from COVID era over- and mis-hiring<sup>5</sup>. Given how few Agentic AI systems make it to production, this is perhaps not surprising - The recent MIT report<sup>6</sup> found that “95% of organizations are getting zero return”. But behind the headlines the story is more complex. The “TL;DR” is:

- People are misdirecting investment into high profile customer facing projects, instead of more lucrative but boring fixes to back office work.
- The biggest single issue is the inability of Agentic AI to change. If it makes a mistake once, it'll keep making it.
- People try to fix the mistakes by making the prompt more complex, but that has diminishing returns, with each fix tending to create one or more new problems. This is reminiscent of the plot of the movie *Bedazzled*<sup>7</sup>, in which the protagonist gets seven wishes from the Devil, but finds they are never implemented in the way that actually helps him.
- Successful deployments ignore vendor benchmarks and focus heavily on business outcomes.
- Good integration with existing systems is also a key success factor.

---

<sup>3</sup> <https://www.bbc.com/news/articles/clyk7zg0gzvo>

<sup>4</sup> <https://www.rand.org/pubs/commentary/2025/10/ai-is-making-jobs-not-taking-them.html>

<sup>5</sup> <https://www.moneycontrol.com/europe/?url=https://www.moneycontrol.com/technology/ibm-ceo-arvind-krishna-pins-mass-layoffs-on-covid-era-over-hiring-and-not-ai-here-s-why-article-13711379.html>

<sup>6</sup> [https://mlq.ai/media/quarterly\\_decks/v0.1\\_State\\_of\\_AI\\_in\\_Business\\_2025\\_Report.pdf](https://mlq.ai/media/quarterly_decks/v0.1_State_of_AI_in_Business_2025_Report.pdf)

<sup>7</sup> [https://en.wikipedia.org/wiki/Bedazzled\\_\(1967\\_film\)](https://en.wikipedia.org/wiki/Bedazzled_(1967_film))

# High Level Challenges of Agentic AI

## Agentic AI is heavily bubbled

At this stage even the optimists at Gartner are saying that 40% of AAI projects will be cancelled by the end of 2027<sup>8</sup>. Agentic AI is currently on an unsustainable trajectory, with surreal requirements for both investment and energy<sup>9</sup>. In addition, the models are not getting smarter as was promised. History shows us that at some point the bubble will burst. This could lead to model suppliers ceasing to trade, throttling access or raising prices.

## People are under huge pressure to 'do something'

Despite all this, large numbers of AI initiatives are being pursued. We can expect a lot of leads will involve Agentic AI, and that a lot of projects will be started. While this is obviously an opportunity we can't afford to ignore, we need to bear in mind that a lot of what's being started won't get finished, regardless of how well we do our jobs. Any strategy we adopt needs to think in terms of the deal's impact on the long term relationship with the prospect.

## The economics of Agentic AI are fragile

LLM access is currently heavily discounted, and every API call loses money<sup>10</sup>. At some point costs will increase. In addition, it seems to be hard to predict in advance how much the LLM side on an AI project will cost. As a consequence, it may be hard to stand over ROI predictions, which in turn create a bias towards 'aspirational' projects instead of those grounded on financial reality. We can try to guide prospects to do end to end POCs on a small scale that doesn't create risk, with the goal of fully validating the concept before jumping in with both feet. Or in some cases we may be able to influence the system design so that the LLM could, if needed, be replaced by conventional code.

## Agentic AI is opaque

Another issue is that with thousands of LLM prompt based decisions being taken every hour appropriate visibility of operations could be a challenge. Given that both the public and GDPR expect that corporations be held accountable, this is not something we can ignore.

## Agentic AI struggles with real world, production data sets

In many real world cases solving a problem may involve looking at three or more separate internal systems. A human can easily spot when one of these systems has 'stale' or incorrect data, AI will not.

---

<sup>8</sup> <https://www.gartner.com/en/newsroom/press-releases/2025-06-25-gartner-predicts-over-40-percent-of-agentic-ai-projects-will-be-canceled-by-end-of-2027>

<sup>9</sup> <https://www.wheresyoured.at/> - Site has an 'angle', but what it says is hard to ignore...

<sup>10</sup> <https://futurism.com/artificial-intelligence/openai-is-suddenly-in-major-trouble>

# Agentic AI - A Draft Plan...

My advice would be to focus on creating and sustaining our long term relationship with a prospect, instead of fixating on AAI, even if AAI is how we have been introduced. We also need to avoid projects which have a high chance of failure, where possible, and consider the potential blast radius of a project before signing up for it. This is admittedly tricky. There will be a natural incentive to try and sell as much as we can, as fast as we can. We may have to tell people we have spent time and energy building a relationship with that we want nothing to do with their AI project. We also need some idea as to what projects should be prioritised. Below is an example of a checklist we could use to 'qualify out'.

## Pre-engagement qualification checklist

### **Clear and tangible payback**

We need to promote a ruthless focus on visible business value - the 'IKEA model', not 'Concorde', The nature of AAI makes it very easy to build your own 'Concorde' - a solution which can never succeed because its cost massively exceeds the maximum benefits it can offer<sup>11</sup>. Instead, we should adopt the IKEA pricing model<sup>12</sup>. We should guide prospects to start with a maximum cost they can pay and still make money, and work backwards. We only build if we can see the profit potential early on, and we rethink if we can't.

### **Payback clearly attributable to AAI**

The advantage of a project to (say) automate back office systems is that savings are clear and measurable. If, instead, we focus on customer facing systems, it may be harder to tell whether an increase or decrease in sales is caused by the new system, market conditions, or something else.

### **Finite Scope**

We need to avoid scope creep at all costs, as in an LLM universe it's easy to allow it to happen. Given that we're trying to be scientific and create provable results, the fewer variables we introduce the better.

### **The new AAI system can coexist with the existing system**

Rather than an overnight 'big bang' deployment, it would make far more sense to do an initial rollout that affects a small number of customers so we can learn the strengths, weaknesses and above all the economics of what we are doing. This means that we need to be able to create a clear boundary - such as 'The AAI prototype is used by all customers whose customer number ends in 17'. With this we will be able to measure the effectiveness of the new versus the old. We should advise against any deployment that is 'all or nothing'.

### **... and can coexist with other systems**

For a project to succeed there need to be clear and low-friction boundaries between the AAI system and the rest of the business. It would also be justified to not take statements of

---

<sup>11</sup> <https://www.nytimes.com/2003/10/24/nyregion/for-concorde-economics-trumped-technology.html>

<sup>12</sup> <https://www.ikea.com/global/en/our-business/how-we-work/the-ikea-value-chain/>

where the boundary is, or of its 'cleanliness', at face value. In human run enterprises the humans spend significant time making sure that work can be handed off between departments by using workarounds, unwritten rules etc. These may not be known to senior management or the engineers building the system, but will become apparent when AAI is turned on.

## Where do Data Platform vendors fit in?

"Everyone has an AI *story* for you. We have an AI *plan* for you". My suggestion is that we should position our Platform as "the only database that your AAI needs to speak to". We need to always be seen to be focusing on the practicalities of implementation, not looking further and further ahead. This pragmatic approach will lead to more and stickier licences.

### General AAI state storage

Assuming we are comfortable with the results of the checklist and still wish to proceed, what next? My recommendation is that we position ourselves<sup>13</sup> as the only database platform in an AAI deployment that directly speaks to the AAI. Everything that interacts with AAI comes out of our platform.. LLMs generally have four kinds of 'storage' for data; and Data Platform vendors are relevant to each:

1. "Short Term Memory" is used to keep track of multi turn conversations<sup>14</sup>. LLMs are functionally stateless, and rely on this Short Term Memory for anything other than the simplest tasks. There will generally be a mechanism for storing this, so that if conversations are interrupted or switch application servers they can be resumed. This is a clear opportunity for Data Platform vendors.
2. "Long Term Memory" allows LLM code to store user or application specific data in an external store<sup>15</sup>. Data Platform vendors can be positioned as this store.
3. Tools, which are exposed to and used by the LLM<sup>16</sup>, can also have their own storage requirements. If our data platform allows us to directly expose APIs it could be very useful in this context. In particular, we can use tools that call Data Platform vendors to provide a consistent 'customer 360' view of the enterprise for our AAI. This will also make Data Platform vendors sticky.
4. Data Platform providers are also adding support for Vector Storage and Search<sup>17</sup>, which allows us to present Data Platform vendors as a "One Stop Shop" for AAI applications.

---

<sup>13</sup> <https://github.com/singlestore-labs/langchain-singlestore>

<sup>14</sup> <https://docs.langchain.com/oss/python/langgraph/add-memory#add-short-term-memory>,  
<https://github.com/singlestore-labs/langgraph-checkpointer-singlestore>

<sup>15</sup> <https://docs.langchain.com/oss/python/langgraph/add-memory#add-long-term-memory>

<sup>16</sup> <https://docs.langchain.com/oss/python/langchain/tools#tools>

<sup>17</sup> Eg <https://www.singlestore.com/blog/build-powerful-agentic-ai-applications-using-singlestore/>

We should try and avoid any situation where an LLM powered application is given direct access to a raw and live database, instead of via an API. Real world databases are often hard to interpret, so while it might theoretically be possible to point the LLM at a production database and say 'have at it!', a wiser course of action would be to create a set of tools that each call Data Platform vendors to solve a specific problem.

### Use Case: A 'Single Store' of corporate data / Customer 360

As part of the research for this project I held an off-the-record conversation with an employee of Indeed, who are heavily invested in AAI. From his perspective the biggest challenge he faces is being able to provide agents with a 'single source of truth' to take decisions with. Without this you end up with a large number of calls to separate tools. Each tool will be speaking to a database. But the databases are absorbing real world data with varying lags. This, in turn, means that we end up with incoherent and nonsensical answers that are hard to debug because the data stores in question will finally be in sync by the time we get round to questioning them. In addition, as I mentioned above, real world databases can be hard to interpret.

This 'Customer 360' use case is well known and understood. It's also robust and useful enough that even if the AAI part fails to perform well, it can still provide obvious value regardless.

### Use Case: Command, control and auditing of Agentic AI

Because of how fast an AAI system could go wrong, especially if confronted with a malign actor, some system for making sure things are more or less as they are *in real time* is needed. This is another historically known and understood use case.

### Use Case: LLM Policy and Charging

The Telco industry provides an excellent example we can re-implement in the AAI space. In Telco, Charging and Policy systems are used to stop individual users running up huge bills that can't or won't pay, and also to prevent a small group of people from using all the network's resources. In the context of AAI we can use Data Platform vendors to watch AAI behaviour to make sure our system doesn't squander resources, for example by spending real money on the wrong thing, or by allocating the same inventory multiple times. This would become especially important in the context of fraud prevention.

### Domain specific use cases

The examples above all focus on solving problems associated with the basic operations of an AAI system. The ones below start getting into more specific use cases.

#### **Use Case: Fraud Detection**

Criminal activity inevitably follows innovation, so we can expect to see AAI systems targeted by fraud, especially due to their gullible nature. This implies a need for a system to monitor for unusual behaviour, both at the transactional and aggregate levels. Data Platform vendors

has a role to play here. Examples would be: Tracking how many items our customers are requesting, and flagging users who request to 'buy' a negative or fractional number of an inventory item. At an aggregate level Data Platform vendors could be used to watch for surges of bots.

### **Use Case: Hyper-Personalisation**

Hyper-personalisation is a huge field, but the common features of all systems are a need to know about an existing customer, a need to have an inventory of things to give or sell them, and the ability to track when an offer has been made. There could also be overlap with fraud detection - if we are running a customer service bot that hands out vouchers from free phone upgrades, it would be nice to know that Data Platform vendors would prevent multiple vouchers being handed out if a malicious user starts 200 bot conversations simultaneously.

### **Use Case: Video Events Mediation and Decisioning**

Modern CCTV systems can do basic classification of objects and people. This means you can build systems that track people and objects as they move through a built environment, potentially opening doors or turning off hazardous machinery as they approach. This assumes you can track state, which is where Data Platform vendors would come in.

## **Operational Challenges of Agentic AAI**

### **Brownfield deployments - lower risk, lower reward**

Lots of 'opportunities' will be marginal improvements - 3-5%, such as "Next Best Offer". This is a tricky space to play in. Here's an example:

*Imagine a US\$100M ARR business unit. Assume it has a 20% margin and thinks it can increase revenues 3-5% with AAI. Tempting, eh? But that means that if our AAI solution costs US\$3M a year we've done little to affect the bottom line. We may even have reduced margins while slightly increasing turnover. Bear in mind that this 'marginal improvement' space is the same one where ML and Data Science failed to make a huge impact. US\$3M a year is not a lot when you add hosting and staff costs.*

### **Greenfield deployments - higher risk, higher reward**

Real benefits will come from use cases that rearrange/remove/reinvent boundaries where the business stops and the customer starts. Amazon books a classic example:

*Before Amazon books you could, in theory, order any book from a bookstore. But that meant explaining what you wanted to a store employee, they would then translate your request into a bunch of computer search terms, coming back with nothing more than a list of titles and*

*possibly some side-eye about your reading choices and asking if you want to order. What Amazon did was put a much better version of that same search interface on your desktop and eliminated the pesky bookshop.*

This is what I mean by greenfield use case. It's not necessarily 100% new to the business, but is changing boundaries and reducing friction by removing architecture we historically took for granted..

Obviously, AAI has huge potential here, because we're no longer tweaking percentages. But because AAI plays a central role in these use cases, if it goes wrong, everything goes wrong. And by 'going wrong' I don't just mean technical issues. The biggest risk will be commercial, as technical stuff we can at least test before we unleash it on an unsuspecting world.

Commercial risks could be things like:

- Unilateral pricing changes by our AI engine provider
- Our AI engine provider ceases trading
- Unfixable issues with the prompts
- Customer failure to adopt or backlash
- Legal landmines, such as GDPR, data sovereignty etc

## We can't lose our focus on developers

It's important not to get carried away and to continue to support and engage with developers. A shift towards enterprise value selling is an extra activity, not a replacement for the current ones. If you look at Volt they made a huge mistake by ignoring developers. The problem is that in any prospect there will be technical people who *aren't* on the call, but *are* in positions of influence within an organisation. Their natural instinct when somebody above them in the hierarchy shows up with a PDF from a DB company will be to directly or indirectly say no, unless it's a name they know and respect. Maintaining brand awareness is thus critical, even if not directly connected to enterprise sales.