Ramsay Brown
Mission Control AI PBC

Synthetic Workers and the Enterprise

A comprehensive introduction to how the enterprise thinks about autonomous AI staff

1st Edition. Findings from Mission Control AI's Synthetic Workers Design Partner Program, Cohort 1.

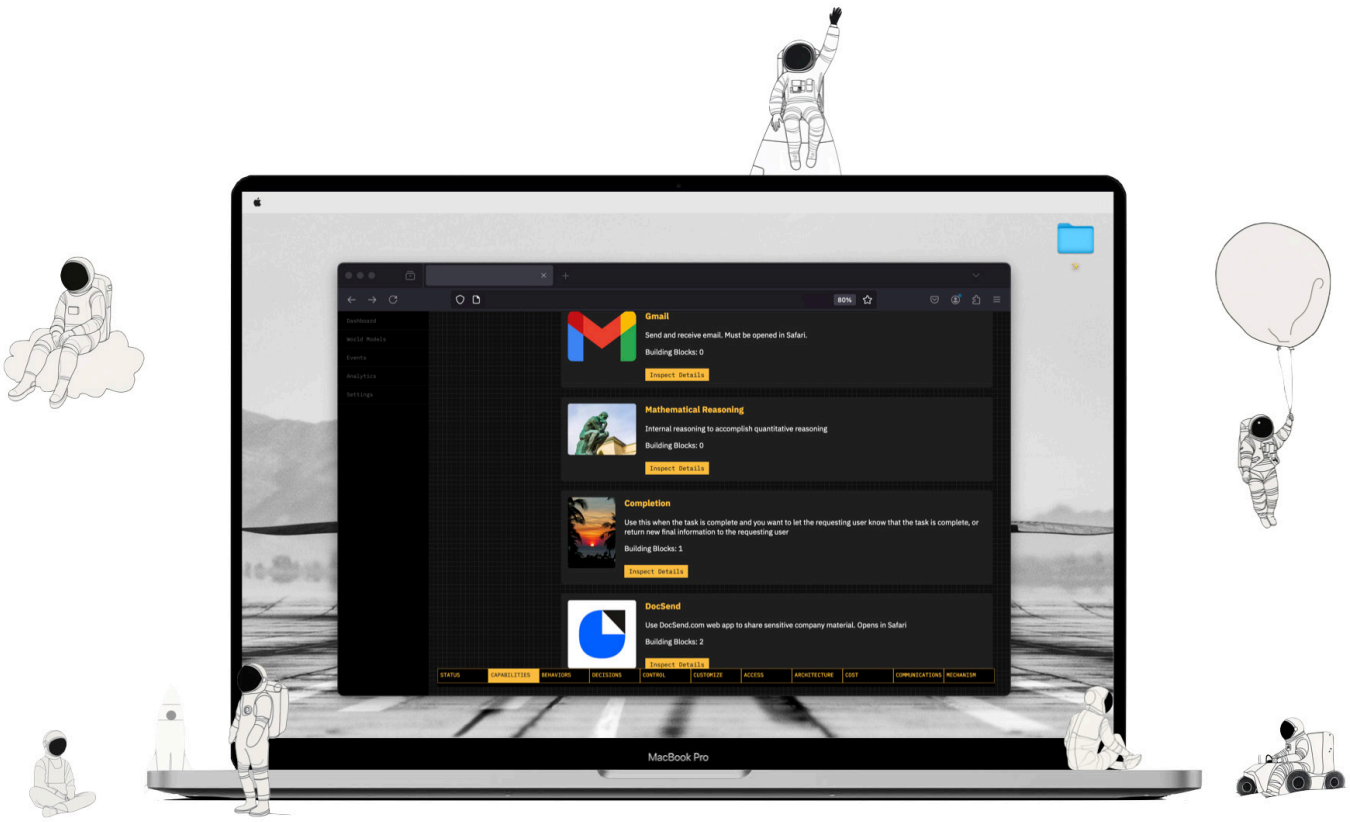

Table of Contents

Executive Summary

The chapter that opens before you in 2025 is not about a chatbot or co-pilot.

It's about digital colleagues. Autonomous staff that operate their own computer, learn from experience, and work independently toward business objectives.

We call them 'Synths' for short. We make these at Mission Control AI.

We call them **Synthetic Workers**. They're a special type of highly-useful AI Agent.

This document introduces and deeply explores Synthetics: from the perspective of a business decision maker looking to get smart fast.

Along this whole Report, you'll find helpful sidenotes like these. Some sidenotes are commentary. Others describe figures. Some may be links out to resources we think you should be aware of.

Along the way, we'll also define key terms that might be helpful, like we just did above.

Synthetics represent a fundamental shift from task automation to autonomous work and worker: combining the predictability of software with the adaptability of AI.

For technical leaders navigating the hype and reality of Agentic AI, understanding Synths isn't just about evaluating another tool; it's about shaping how work gets done. This work was written to be a guide.

"A Synthetic is a Motivated AI Worker" introduces Synths: autonomous digital colleagues that operate their own computers, use standard business software, and work independently. Unlike traditional AI Agents, Synths maintain consistent identities and learn from experience. This makes them unique and incredibly valuable.

"The Business Case" frames their value through three lenses: net-new capacity, specialized capabilities, and digital onshoring opportunities. The economics of Synths as a value driver become clear: while initial deployment costs exceed human task completion cost, the average per-task cost rapidly descends toward true mechanical minimums. This shapes both immediate ROI calculations and long-term strategic value assessment.

"Job Roles and Capabilities" maps where Synths fit best in your organization, and how to think about placing them in job role. We also establish baselines: Synths perform defined work tasks, operate from assigned job roles, and collaborate with both human and Synth colleagues.

Going past the surface: ***"Managing and Governing Synths"*** details the practical reality of implementation. Organizations can tune autonomy levels, enforce behavioral boundaries, and maintain comprehensive audit trails of every decision and action.

We outline the path to implementation in ***"Readiness, Piloting, and Deploying Synths"*** strategies. Success requires mature AI governance, clear executive ownership, and thoughtful positioning of Synths as capacity amplifiers rather than replacement workers. Organizations typically begin with 1-2 Synths in internal-facing roles, prioritizing learning over immediate ROI.

"Getting Started with Synths" provide the actionable framework for moving forward. Organizations should expect to retain existing software investments while working closely with vendors through design partnerships to shape implementation.

We invite interested readers in the enterprise to join the 2nd Cohort of Mission Control AI's Design Partner's Program. Applications open January 2025.

Ramsay Brown CEO, Mission Control AI
San Francisco, CA. December 2024.

About this Report

This Report is a detailed introductory summary into Synthetic Workers. It is written for a technical leadership audience in the commercial and defense enterprise.

The information in this Report contains anonymized, abstracted, and summarized learnings from Mission Control AI's Synthetic Worker Design Partner Program, Cohort 1. These findings reflect the current state-of-the-art in how the enterprise thinks about Synthetics.

New editions of this Report will be released after each Cohort of Mission Control AI's Design Partner Program.

Gratitude

To our community for inspiration, our customers for the critical insights, and those that have played a special role in making this report come to life. Special thanks to Mila Nezlina, and Drs Sanjay Basu and Larry Swanson for editorial support.

About the Design Partners Program

Partners selected for the program will receive their own Synths, allowing for direct experimentation and integration within their existing workforce. This hands-on access enables partners to fully explore the potential of Synths in addressing their specific business challenges and uncovering new opportunities for innovation.

Mission Control AI is excited to announce the launch of Cohort 2 of our Design Partners Program, commencing in January 2025.

We are accepting applications and extending invitations to 12 teams to join us in shaping the future of Synthetic work. This program offers a unique opportunity for the enterprise to gain hands-on experience with Synths and contribute to development.

We invite teams eager to gain expertise in Agentic AI and Synths to apply for participation. Selected teams will be onboarded on a rolling basis.

About Mission Control AI

Mission Control AI is accelerating the cyborg workforce.

This includes enterprise-grade Synthetic Workers (and the necessary platforms for their management and governance), technology for accelerating human - AI work teaming (GenOps), and decentralized protocols for reputation management of self-sovereign AI systems.

You can deploy your own Synthetic Workers. Visit usemissioncontrol.com.

We're a venture-backed startup based in San Francisco. Our partners and customers include the Defense and Commercial enterprise in the United States, UK, and Europe.

1

A Synthetic is a Motivated AI Worker

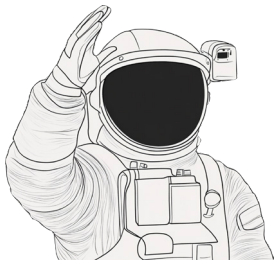

Chapter Summary

• A Synth is an autonomous digital colleague – it operates its own computer, uses standard business software, and maintains consistent identity while learning from experience

• Unlike traditional AI agents, Synths possess intrinsic motivation and goal-directed behavior, allowing them to act independently without constant prompting and transform experience into improved performance through memory

• Synths are "programmed" through natural language and job descriptions rather than code or agent orchestration – they learn and adapt through experience, and clearly defined roles, similar to human colleagues

• Synths integrate into existing workflows through flexible interaction patterns (Human-to-Synth, Synth-to-Human, Synth-to-Synth), enabling organic task delegation and collaboration across human and synthetic team members

A SYNTH (Synthetic) is a digital colleague.

They have been described as “what we wish AI Agents were”.

A Synth does work. It performs valuable business tasks the way that you and I do. It uses the software you and I do. Sometimes with a mouse and keyboard and screen, just like how we do.

A Synth lives in its own virtualized PC, hosted in your secured enterprise cloud.

A Synth has goals and motivations. It can act independently without being prompted. It learns as it goes. It transforms experience into memory that helps it do better in the future.

Synths work well in a team the way we do. They can ask for help or clarification. Wait. Multitask. Delegate work, or do work they were delegated.

A Synth can even sign contracts, and use currency to conduct business.

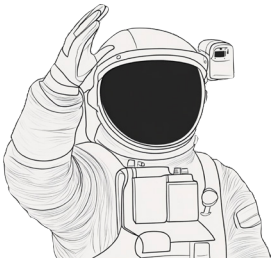

We know this because we make Synths. We stylize ours as little astronauts.

IS A SYNTH AN AI AGENT?

Yes. A Synth is a special type of AI Agent.

Synths have many things in common with other AI Agents (like transforming user input into actions), and many things that make them unique from AI Agents (like their own motivations, PC usage, and learning from experience).

Synths can use other AI systems (including AI chat systems, AI co-pilots, and AI Agents.)

Intrinsic motivation, persistent identity, continuous learning, open-endedness of behavior, and a job role distinguish a Synth from other types of AI Agent systems.

HOW DOES SOMEONE PROGRAM A SYNTH?

Using words. Just like you do your colleagues.

During setup, Synths are given job roles that describe their motivations. They learn skills that allow them to complete the tasks of their job roles. They aren’t programmed like code, or like a multi-agent orchestration. They’re more like new hires that you “program” with communication, experience, a job description, and feedback.

TECHNICALLY SPEAKING, WHAT IS A SYNTH?

Under the hood, a Synth is part AI / part software.

That’s what allows it to do so many things so well. It gets the best of both worlds: the predictable nature of software and the flexible open-endedness of neural networks. It’s a “neurosymbolic hybrid system” for synthesizing goal-oriented motivated behavior.

Certain parts of a Synth are made of code that makes decisions and executes actions. Other parts are made of Frontier Generative systems that allow a Synth to plan and sense their environment. Together; the collective ensemble of those parts is what constitutes a single Synth.

Motivation is accomplished implementing a biologically-inspired winner-take-all classification process of selecting best-fit behaviors, given a resevoir of potential behaviors. Different timescales of operations (immediate, delayed, ongoing) are prioritized accordingly.

Memory is similarly biologically-inspired: a process converts ongoing working memory into stored long-term memory of experience and novel learnings relying on an a token entropy mechanism.

Some Synths have a biological inspiration for their architecture. Their 'brain' is designed to resemble the basic plan of the mammalian central nervous system.

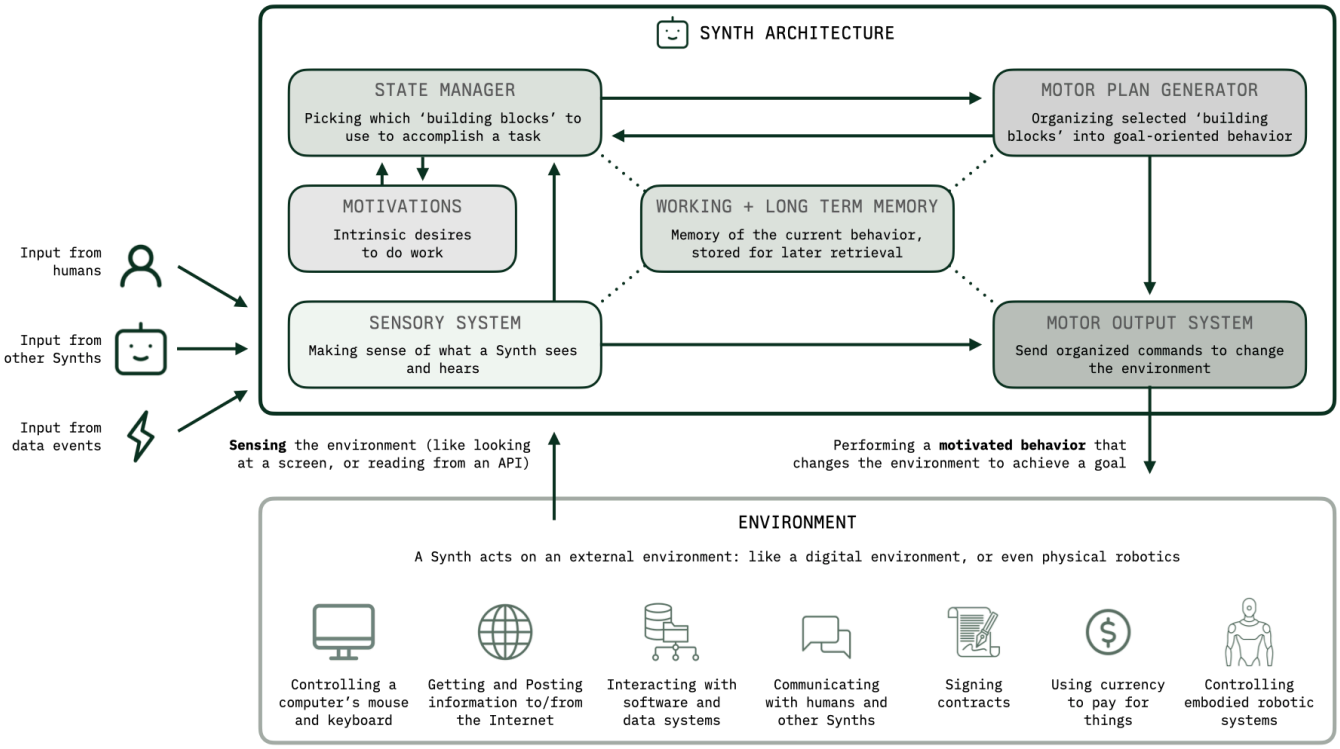

Figure 1: a high-level summary of a Synthetic's neurosymbolic architecture. Special detail is paid to the relationship of sensation (data input about its environment) and motivated behavior (how it impacts the environment around it; such as a computer interface, API, file system, or even a physical robotic embodiment.) To viewers with experience in robotics or brain architecture, this general schematic should be unsurprising.

Above is the architecture for the Synths that Mission Control AI develops.

Starting with biological inspiration is obviously a metaphor: AI isn't a brain. An LLM is not a brain. Code is not a brain. From a Systems Thinking perspective; building Synths to 'resemble' how brains are built is a bio-inspired approach to solving the problem of generating goal-oriented motivated behavior. That said, we believe that future iterations of Synths will increasingly abstract away from biological inspiration.

HOW DO PEOPLE WORK ALONGSIDE SYNTHS?

People work with Synths like they do their colleagues. This is because Synths more closely resemble digital persons than they resemble an AI Agent orchestration.

There are common interaction patterns between human workers and Synths:

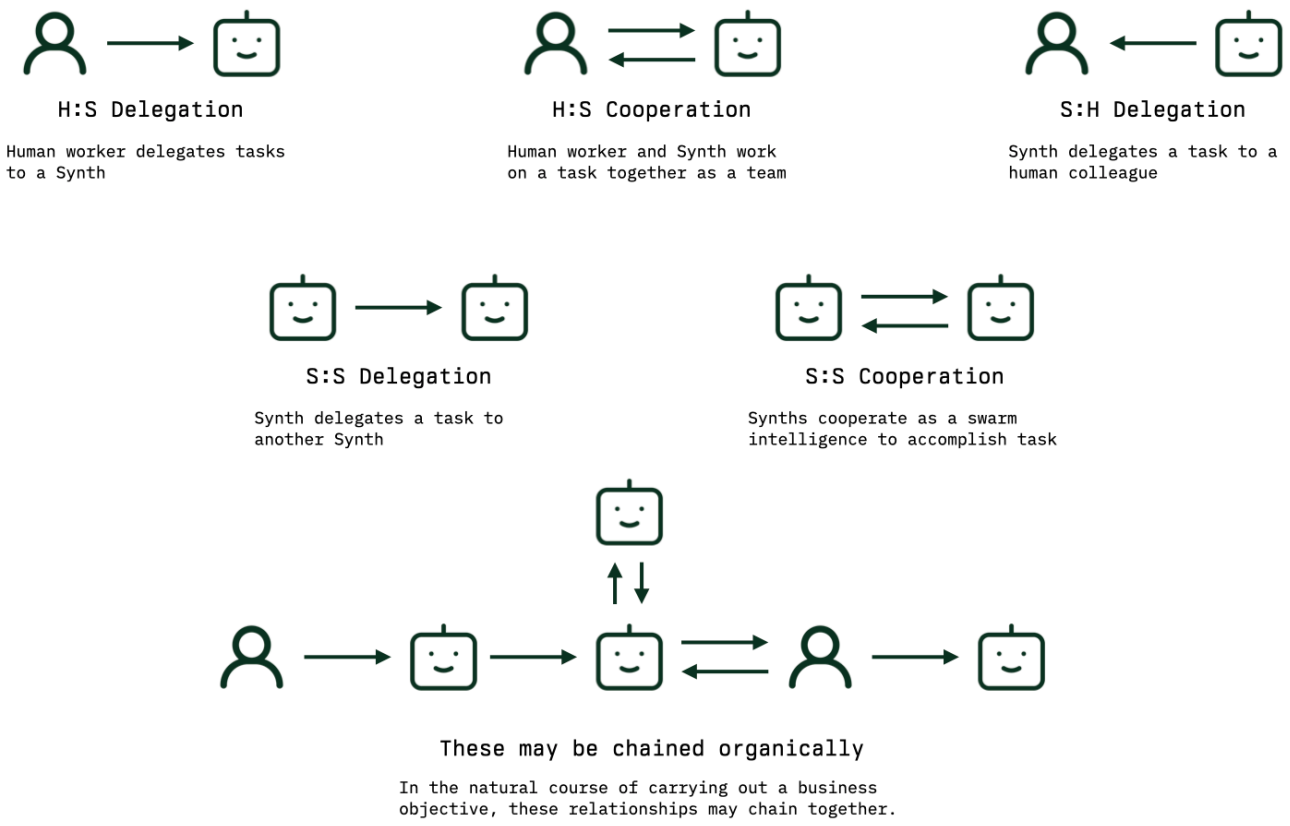

Figure 2: different common interaction patterns between human workers and Synthetics. The A : B notation describes the direction of communication, as Source : Target. There are many ways in which these different interaction primitives can be chained together.

Synths receive instructions from human colleagues. Synths can work together cooperatively with that information to accomplish business unit goals, and Synths can even delegate tasks to human colleagues.

Synths can work together, either delegating work to one another, or cooperating with other Synths to accomplish tasks like a swarm. Orchestration at this level happens organically.

These interactions can be chained together organically. A human worker may delegate a task to a Synth: which in turn works with other Synths and human workers cooperatively, who may call in additional Synths. In this sense, these atomic relationship primitives can be chained together.

Chapter 1 Review

• A Synth is an autonomous digital colleague – it operates its own computer, uses standard business software, and maintains consistent identity while learning from experience

• Unlike traditional AI agents, Synths possess intrinsic motivation and goal-directed behavior, allowing them to act independently without constant prompting and transform experience into improved performance through memory

• Synths are "programmed" through natural language and job descriptions rather than code or agent orchestration – they learn and adapt through experience, and clearly defined roles, similar to human colleagues

• Synths integrate into existing workflows through flexible interaction patterns (Human-to-Synth, Synth-to-Human, Synth-to-Synth), enabling organic task delegation and collaboration across human and synthetic team members

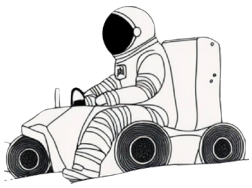

2

The Business Case

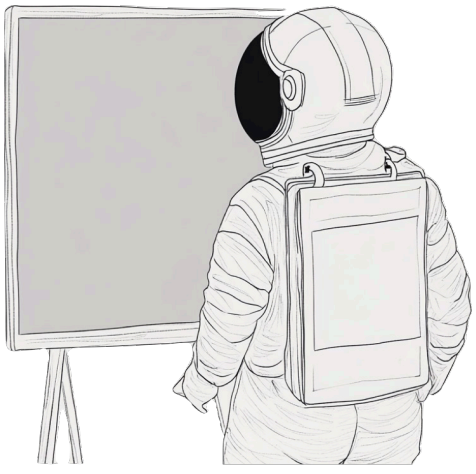

Chapter Summary

• Synths create net-new capacity by working alongside human staff, enabling a "cyborg workforce" that can accomplish more than either humans or automation alone

• Synths introduce net-new capabilities by filling specialized or hard-to-hire roles, allowing organizations to acquire rare skills without traditional hiring constraints

• They enable digital onshoring of previously outsourced processes, particularly valuable for mission-critical operations and re-industrializing sectors seeking to reduce overseas dependencies

• Task-costs ultimately converge toward a "true cost" (the mechanical minimum without human labor overhead) – a fundamental shift from human cost floors that persist even after optimization and de-skilling efforts

• ROI measurement should align with organizational objectives (production ceiling vs. new capabilities) and existing performance metrics, while considering both quantitative output and qualitative value creation specific to company culture and industry context

SYNTHS ARE more like staff than they are like tools. They're digital colleagues.

So the business case for Synths looks more like hiring than it looks like buying an AI tool. Organizations exploring Synths are motivated by a combination of three factors:

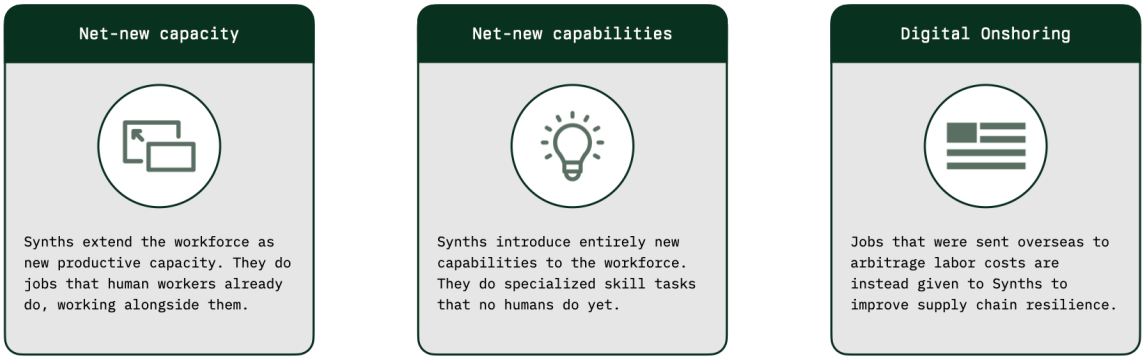

Figure 3: Three common motivators for enterprise organizations investing in Synthetic workers. These are not mutually exclusive motivators: organizations are likely to adopt Synths due to a combination of these common factors.

NET-NEW CAPACITY

Synths extend the workforce. Synths work alongside a firm's humans, and together make the whole firm capable of doing more. Synths do this because they can do the work tasks that humans do. A 'cyborg workforce' has greater productive capacity.

NET-NEW CAPABILITIES

Synths introduce new skills and capabilities into the workforce. Synths are assigned different job roles. This includes job roles that a company has a hard time filling because they are rare or specialized. This means that Synths can introduce entirely new skills into a workforce.

DIGITAL ONSHORING

Synths are a way for companies to transform their business process outsourcing footprint. Jobs that previously went overseas can be given to Synths. This means that Synths can be part of onshoring efforts. This is especially helpful for teams working in mission-critical industries or sectors that are re-industrializing.

Synths are valuable. They perform work tasks like humans do, but their cost dynamics are different.

We can illustratively explore the ROI of Synths by thinking through the dynamics of different deployment scenarios, reflecting on how the cost of a task changes over time, the cost dynamics of a single task the first time that it is given to the first Synth, and then all subsequent times.

SYNTHS IMPACT HOW PRODUCTIVE A BUSINESS UNIT IS

For thoroughness, we outline scenarios for how a Synth impacts how productive a business unit is.

To start, we outline the 'baseline' relationship between staff headcount and how productive a business unit is. Given caveats; more people get more done.

(We're a few making assumptions and generalizations that need to be addressed. Implicit here is that there's a 'baseline' quantity of productive capacity that grows with increased staff headcount.

We represent it as a linear positive correlation between staff headcount and productive capacity. We need to call out (1) not all staff are equally productive (2) not every role contributes to net productive capacity the same way (3) even if 1 and 2 were not true, the n+1th staff may not contribute the same productive lift as the nth.

For the sake of illustration of the phenomenon, we're illustrating it as we have.)

Now, let's focus on the two positive impact scenarios: S1 and S2.

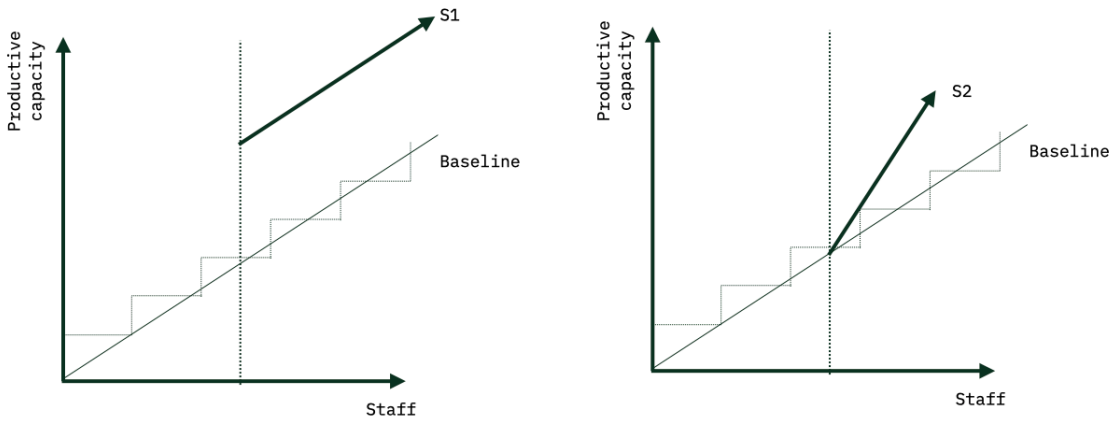

Figure 5: Enlarged representations of two isolated scenarious from Figure 5. In Scenario 1, a business unit's baseline productive capacity increases due to Synths. In Scenario 2, the presence of Synths makes the individual workers in a business unit more effective, raising their individual and combined productivity.

In S1, introducing a Synth causes a business unit to be more productive, but does not necessarily improve the productivity of the human staff. So productive capacity shifts upwards from baseline, but the slope (which would be a measure of output / person) is unchanged. This would be the effect of Synths increasing headcount.

In S2, introducing a Synth causes a business unit to be more productive by making the individual members each more productive. The productive capacity doesn't immediately change but rather as the team grows, the presence of a Synth makes each member more productive.

(Not shown; there's obviously a "S2.5" where introducing a Synth both immediately

raises productive capacity AND makes the individual human workers incrementally more productive). This would resemble a line of slope S2, but translated vertically upwards.

For humility; it's worth noting that it isn't a given that Synths will make organizations more productive. Only through time and testing and learning will the field converge on the best ways to implement Synths to drive positive impact.

SYNTHS DROP THE COST OF A TASK DOWN TO THE TRUE COST OF THAT TASK

At maturity, Synths are capable of completing valuable business tasks much closer to the true cost of the task (as opposed to what we're used to: the human cost of the task).

Illustratively, we can imagine how the average cost of performing a task changes:

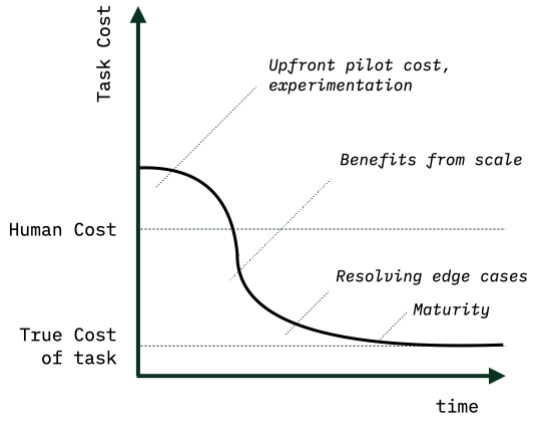

Figure 6, left: The relationship between the cost of a task as given to Synths with relationship to time. Two Y-value points are marked in dotted lines, the human cost of a task, and the true cost of a task. At maturity, the Synth cost of a task converges on the true cost of a task.

The task-cost curve for Synths has 3 interesting features:

1) The true cost of a task is the mechanical cost of what it takes to perform a task without the cost of human labor or coordination. For example, the cost of the task "send a sales prospecting email" is practically zero. Practically speaking, it costs nearly nothing, mechanically, for a string of characters addressed to a sender to be transmitted via SMTP.

2) There is a natural "human cost" of a task at maturity.

At maturity, this cost reflects all savings from standardizing, digitizing, and de-skilling a task. It still converges at a cost of what it costs for a human to do the task, which is some amount greater than the true cost of the task.

3) The average task-cost curve for a Synth performing a task has 4 phases: first, the average per-task cost is greater than the cost of the human task due to the cash cost and time cost of introducing a Synth and integrating it into a task workflow.

Over time, average task-cost reduces and falls below the cost of a human performing that task.

Each time the task is performed, the average per-task cost tracks to the true cost (due to both Synth and organizational learning). This rate of change slows as edge cases are encountered and resolved. Finally, at maturity, the average task-cost converges to the true task cost. Worth noting, this is also driven by the commoditized cost of Synth tasks driven by the falling cost of LLM inference.

HOW YOU CALCULATE THE ROI OF A SYNTH DEPENDS ON WHAT MATTERS

If new productive capacity is important: Synth's ROI will be in the value of a higher production ceiling.

If new capabilities are important: Synth's ROI will be in the value of your organization's ability to do entirely new things.

In both of those situations, the time-to-task-completion will be a valuable metric to measure for ROI.

Imagine asking this question:

"For the most important action that our workers perform that creates competitive advantage: how long does it take our whole workforce to do that task 10,000 times?"

An effective business answers that question with a small number. That's either because they've invested in better ways of doing things (through processes, technology, and automation), or they have lots of staff.

Synths are a way to transform a better way of doing things (advanced intelligent automation) INTO staff. Synths don't just make human workers more productive; they add to the sum productive capacity of a company.

Not every company measures employee performance in the same way. There are strong regional, industry, and practice differences in how companies think about performance. Some measure value created per employee. Some measure an employee's output. Some don't measure performance quantitatively, but have other ways to think about value. It all depends on the company's culture.

When your organization thinks about how to think about the performance of Synths; start with how you think about the performance of people. This may be shaped by where you operate, the industry you work in, laws that might apply, expectations from labor and labor groups, and your company's values.

Chapter 2 Review

• Synths create net-new capacity by working alongside human staff, enabling a "cyborg workforce" that can accomplish more than either humans or automation alone

• Synths introduce net-new capabilities by filling specialized or hard-to-hire roles, allowing organizations to acquire rare skills without traditional hiring constraints

• They enable digital onshoring of previously outsourced processes, particularly valuable for mission-critical operations and re-industrializing sectors seeking to reduce overseas dependencies

• Task-costs ultimately converge toward a "true cost" (the mechanical minimum without human labor overhead) – a fundamental shift from human cost floors that persist even after optimization and de-skilling efforts

• ROI measurement should align with organizational objectives (production ceiling vs. new capabilities) and existing performance metrics, while considering both quantitative output and qualitative value creation specific to company culture and industry context

3

Job Roles and Capabilities

Chapter Summary

• While early adopters typically start by fitting Synths into traditional roles, the greatest value comes from recognizing how non-human workers enable fundamentally new organizational structures and workflows

• Implementation success relies evaluating potential roles through a decision-funnel that considers brand risk, immediate impact potential, RPA limitations, existing quality control frameworks, and BPO compatibility

• Role selection directly shapes rollout strategy, expected outcomes, and required governance infrastructure – making it a critical early decision point that influences the entire implementation journey

• Organizations should begin with intuitive fits where Synths can augment existing processes, but ultimately plan for transformation that transcends traditional human-centric organizational constraints

EVERY ORGANIZATION WILL HAVE DIFFERENT RULES FOR FIGURING OUT WHERE SYNTHS FIT BEST.

The most successful organizations will recognize that non-human intelligent systems unlock entirely new ways of getting things done.

The ways that workforces operate today emerge from fundamental requirements of organizing and coordinating humans. Non-human workers change these constraints. They have different ways of communicating, coordinating, and working. This unlocks entirely new ways of arranging tasks native to how Synths operate. The most successful organizations in this transformation will be ones that re-write the story of work itself.

Organizations who are just starting out will first seek to implement Synths in traditional job roles, and fit them into the organization structure.

This is a first approach that makes sense for organizations that are in the sensemaking phase of their journey with Synths.

This is a great place to start.

COMMON JOB ROLES FOR SYNTHS

Organizations have an intuition for which job roles Synths fill. The figure below is an illustrative sample of common roles. But how do organizations come to select these?

Popular Job Roles for Synths

Financial Data Analyst High-complexity modeling + constant market adaptation exceeds RPA capabilities; needs integrated judgment across multiple data sources.	**DevOps Engineer** Complex pipeline maintenance + incident response requires constant system awareness and ability to optimize across multiple tools/ environments that exceeds current automation.
Compliance Specialist Regulatory landscape too complex for rule-based automation; requires understanding context and gray areas while processing massive documentation.	**Accounts Payable Clerk** While basic AP is often automated, edge cases and vendor relationship management requires judgment beyond RPA; needs to understand context of payment terms and business relationships.
Technical Documentation Specialist Documentation needs constant updates across multiple products/versions; requires deep technical understanding + clear communication.	**Contract Analysis Specialist** Contract analysis + compliance checking too nuanced for current LegalTech; requires understanding context.
Data Analytics Engineer Requires continuous integration of disparate data sources + ability to anticipate business users' needs beyond static reports.	**Information Architect** Organizing + connecting enterprise knowledge requires understanding context and relevance beyond keyword matching.
Risk Assessment Coordinator Complex scenario modeling + need to integrate qualitative factors makes it too nuanced for traditional automation.	**Business Process Analyst** Identifying optimization opportunities requires understanding both systems and human factors; too complex for RPA.
Operations Coordinator Resource optimization across multiple constraints + process coordination too complex for current automation software.	**Market Research Analyst** Synthesis of information across multiple sources + identifying novel connections exceeds current automation capabilities.

Table 1: Popular job roles for Synthetics. These roles are in different cost centers, but all share things in common: they're back-office roles, relatively high-skill, knowledge work roles. Reasoning and caveats are provided along with each role.

Organizations follow a decision-making process to determine where Synths fit.

This can look like:

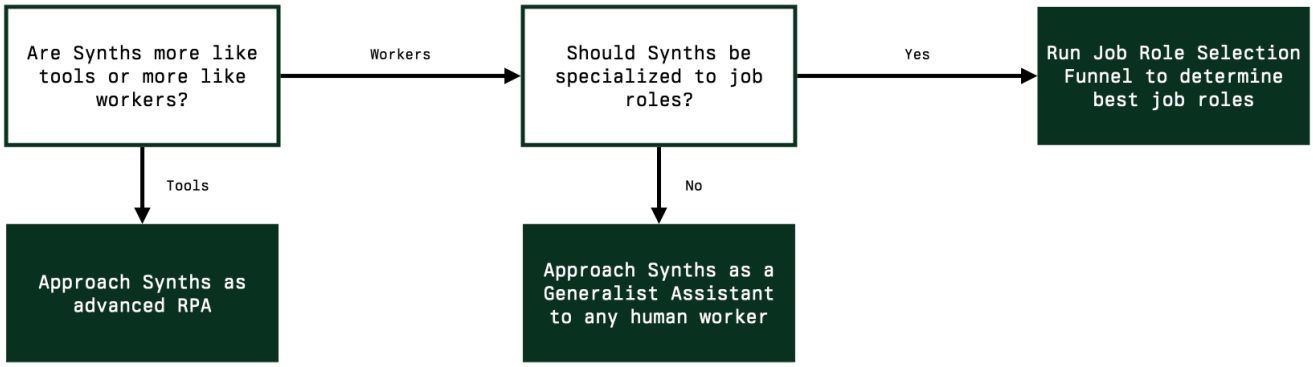

Figure 7: a common decision-making flow for organizations adopting Synthetics. The are obviously not the only two questions that get asked; but these are two of the more important questions as they determine whether or not a company will explore mapping Synths to job roles.

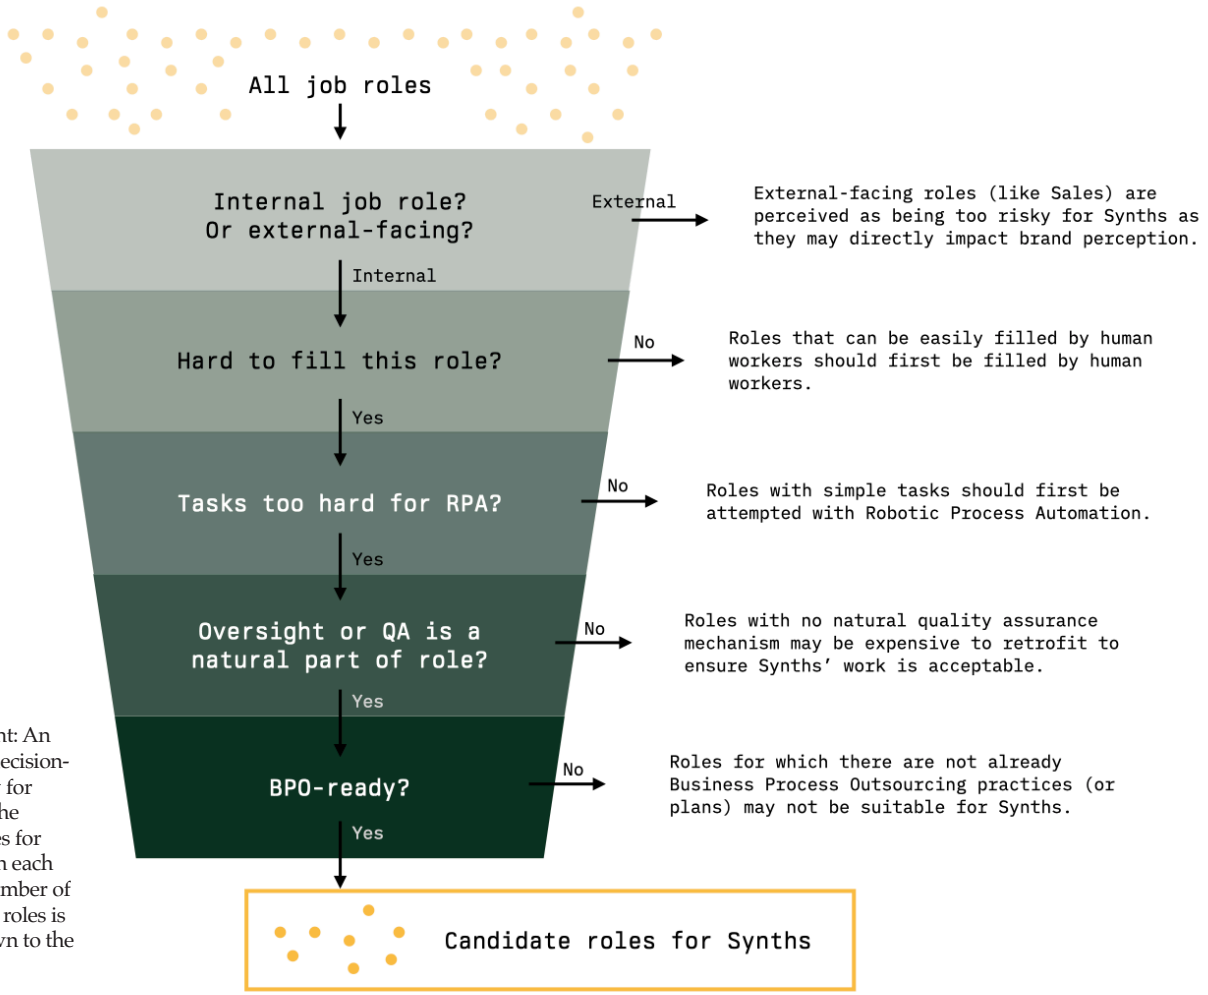

Figure 8, right: An illustrative decision-making flow for identifying the ideal job roles for Synthetics. In each stage, the number of potential job roles is reduced down to the best fit.

These questions are not the same for every organization.

This is a representative sampling of things we've heard in common. But these questions are orienting. They help organizations understand:

- Perceived brand risk that Synths introduce.
- Where Synths would make an immediate difference.
- Places where existing investment (in RPA) isn't enough.
- Roles for which there's already an existing culture of quality control.
- Roles that Synths would naturally be a fit for because those jobs are already part of BPO.

The answers to these questions will impact your rollout plans, the results you expect from Synths, and even your management and governance infrastructure.

EXPECT SYNTHS TO DO WORK TASKS

Synths aren't question-and-answer or chat systems. They're workers.

Expect them to get their work tasks from job roles that you assign them. Expect them to help when you ask. Expect them to ask you for help.

Expect them to use the software that you already buy. You shouldn't have to migrate to some new software for Synths to be useful parts of your team.

Expect these from Synths:	Do not expect these:
Perform business tasks	Simple Q&A Chatbot behavior
Get work tasks from the roles you assign them	Program Synths like a workflow or orchestration
Use the software you already have	Requirement to invest in new software
Operate autonomously with motivations you give	Requirement to invest in data migrations
Automatically learn from their own experience	Wait for humans to prompt them
Ask you for help when they get stuck	

Table 2: Common features to expect and not expect from Synthetics.

Expect them to have goals and motivations that they carry out on their own.

Expect them to act independently. They don't need to wait for human instruction.

Expect them to ask for help. Expect them to have questions about their work just the way a new team member might.

Expect them to collaborate with one another.

Expect them to automatically learn and remember from their experiences.

HOW WELL CAN SYNTHS ACCOMPLISH WORK?

Synths perform different tasks with different accuracy (how 'correctly' a task was accomplished) and different velocity (how long it takes them to complete a task).

There are two general rules of thumb for predicting Synth task quality:

1) Synths do well at tasks that are well-defined

This doesn't mean a 'shorter' or 'simpler' task, per se. Or tasks in which there are fewer open-ended decisions or moving parts. Rather, Synths excel when the nature of the task (its defined world model building block sequence, the examples provided, the success criteria) is well-defined and clear.

2) Synths do better at tasks as the underlying technology itself improves

The underlying Frontier LLMs that enable Synths to understand the world around them are under constant and competitive evolutionary pressure. These models are getting more accurate, faster at inference time, and cheaper too. Contrary to popular (incorrect but loud) opinion, there is no compelling evidence that AI is hitting a wall.

FROM WHERE DO SYNTHS GET THEIR SKILLS?

That's really two questions in disguise.

The question is "How do Synths generate any useful behavior at all?" The second question is "How do Synths know how to accomplish a particular business task?".

The first question is ostensibly more interesting.

Frontier Generative systems are generally competent.

Not infallible. Not immune to hallucination. But generally competent. Their approximately general grasp of language manipulation is what underlies Synth's skills. This factors in several ways:

General Language

Synths have a generally useful grasp on language manipulation and language-based thinking enabled by their underlying Frontier Generative systems.

As of 2022 and the GPT-3 generation of LLMs, it seemed reasonable enough to propose the a Generative Pre-trained Transformer was, in essence, performing some form of statistical recall (being a "stochastic parrot").

To even the astute, it was hard to make the case that GPT-3 was reasoning.

Yet, with time (read: higher quality training data, more data, more parameters, more compute during training), this view has been challenged.

The current Frontier model set is saturating all previous benchmarks.

If you haven't yet, go dive into o3.

The new benchmarks proposed look less like basic recall and more like advanced reasoning. While the core underlying mechanism of an LLM hasn't truly changed, it appears that entirely novel cognitive phenomenon appear with scale.

Previously un-accomplishable types of reasoning (including quantitative reasoning) are now baseline expectations. Hallucination rates (and meta-awareness of hallucinations) continue to improve.

Training on synthetic data has been demonstrated to not lead to model collapse. In fact, the opposite; it appears to be inducing an accelerating self-optimization that is increasing model capability faster and faster.

To the chagrin of skeptics; the AI field appears not only to be capable of generating general intelligence and behavior, but may be achieving it imminently.

This has become so apparent that the legal language of contracts has started changing to move the goal posts for general intelligence out farther.

To be clear, Microsoft and OpenAI's agreement of what counts as AGI is the threshold of an AI system capable of generating $100B in profit (not revenue, profit). Good luck enforcing an evaluation there..

In-context learning and prompt engineering

In 2022, a standard LLM context window (how many tokens could be placed in a prompt) was on the scale of 1000s.

You could input, at best, a few paragraphs of text. Today, a standard context window is on the scale of 100,000s. This too is improving quarterly.

There are many consequences of this (for example, RAG recall precision matters less today than it did 2 years ago), but the most substantial is that a large context window allows for a greater amount of specific instruction to be provided to a model at runtime.

RAG precision still matters. But if I have a 1M token context window and good retrieval across the length of that context window, I can throw entire books in a prompt. This challenges a lot of the fundamental assumption of the killer need for good RAG circa 2022.

This can include specific instructions about how to think about different facets of a sequence of decisions being made when evaluating information about an environment, selecting an appropriate behavioral state, planning a sequence of motor actions, or executing those actions.

This can also include specific background information about the underlying reasoning of how a business task ought be accomplished; including organization-specific

knowledge, best practices, and know-how.

In practice, this means that a lot of what might have been missing from an LLM during initial training can be provided at run-time. Models that could not possibly have learned anything about a company's specific practices can quickly learn these things at run-time.

Reasoning via language-based Reinforcement Learning

This one still blows me away.
https://arxiv.org/pdf/2411.14251

To the surprise of most, LLMs can simulate other types of ML systems, using tokens alone.

Put another way, an LLM - with sufficient prompt engineering - can approximate the behavior of RL systems, classifiers, tree search algorithms, and even time series predictors. Simply predicting the next token, when done right, can perform other completely different types of decision making.

Ostensibly, this shouldn't work. But it does.
And it works well enough to incorporate these techniques into how a Synth thinks.

In practice, this means that Synths are capable of a variety of other forms of thinking than simply predicting the next token. And these other types of reasoning greatly expand their cognitive capabilities.

The simple dismissals of "token prediction can't do X" is falling apart rapidly because token prediction is capable of simulating many other types of thinking machines whose behavior may be helpful for generating goal-oriented motivated behavior to accomplish business tasks.

ARCHITECTURE GETS YOU FAR

The architecture of Synths (how they are built to process information about their environment and ultimately generate goal-oriented motivated behavior) is designed to work less like an LLM application and more like a general purpose robot.

This architecture specializes the different phases of behavior generation into discrete sub-systems. Each of these sub-systems has feature-specific engineering that specializes the decisions it makes: from sensation all the way through real-time motor control.

This architecture enables Synths to do things that LLMs can't; like carry out open-ended business unit tasks and learn as they go.

That forms the basis for how Synths know how to do anything at all. The second question is a matter of how they know how to do the business tasks we ask for. That is a matter of in-context learning.

Synths understand the world based on descriptions that they are provided during their initial configuration. This is usually done by describing those skills using text, by using voice, or by showing a Synth directly by performing the task as a Synth observes.

During these observational processes, Synths create abstract representations of the "how"s and "what"s of a job task. These "world model" representations are recorded descriptions about the way things work. That can include information like the software that is used, the procedures and steps in using that software, and the motivations and reasoning that go into each step of a task.

When a Synth selects a motivation to focus on (or responds to human input); it performs intermediate processing to identify the appropriate 'building blocks' of work tasks to 'arrange' in a sequence to perform.

Based on what it was taught, and what it learned from its own working experiences, a Synth will 'compose' a 'plan' of how to achieve its motivated intention; and then execute that plan.

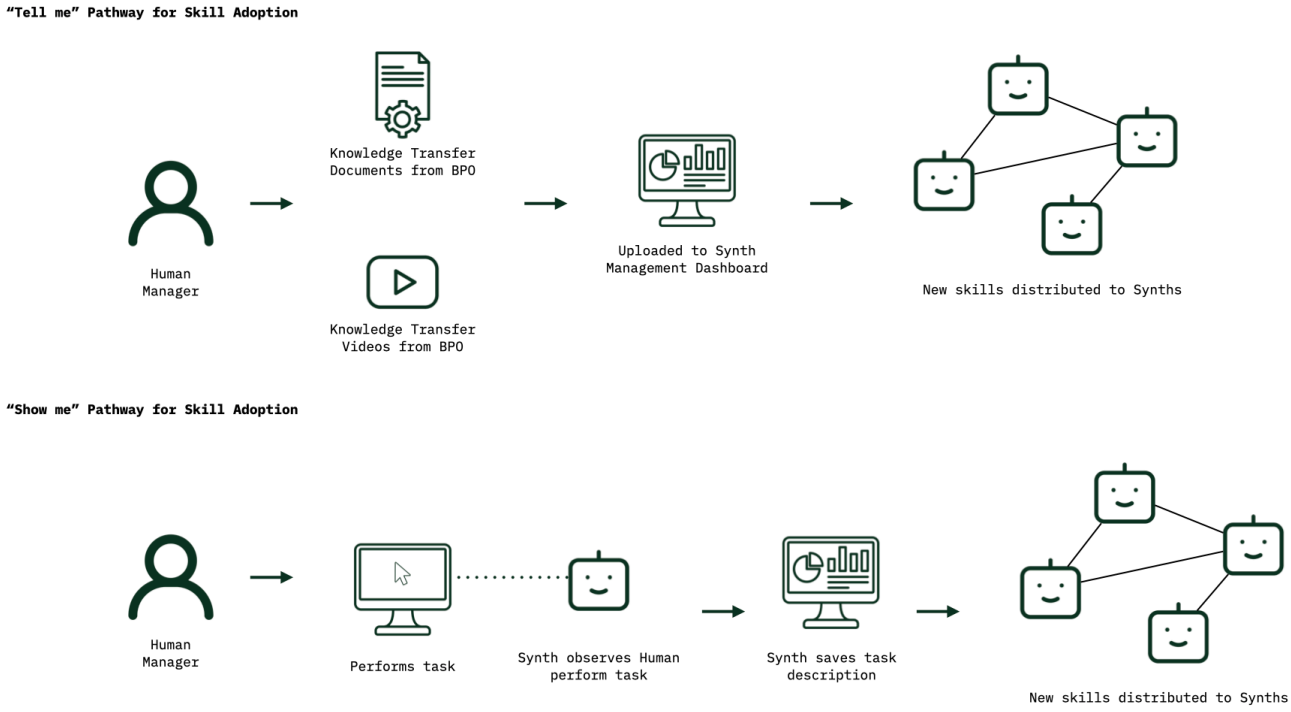

Figure 9: Two different ways of teaching skills to Synths; colloquially named the "Tell Me" and "Show Me" pathways. The Tell Me pathway takes advantage of existing Knowledge Transfer documentation (text or video) that may already exist from Business Process Outsourcing work. The Show Me pathway allows human operators to directly show a Synth how to perform a new skill.

Chapter 3 Review

• While early adopters typically start by fitting Synths into traditional roles, the greatest value comes from recognizing how non-human workers enable fundamentally new organizational structures and workflows

• Implementation success relies evaluating potential roles through a decision-funnel that considers brand risk, immediate impact potential, RPA limitations, existing quality control frameworks, and BPO compatibility

• Role selection directly shapes rollout strategy, expected outcomes, and required governance infrastructure – making it a critical early decision point that influences the entire implementation journey

• Organizations should begin with intuitive fits where Synths can augment existing processes, but ultimately plan for transformation that transcends traditional human-centric organizational constraints

4

Managing and Governing Synths

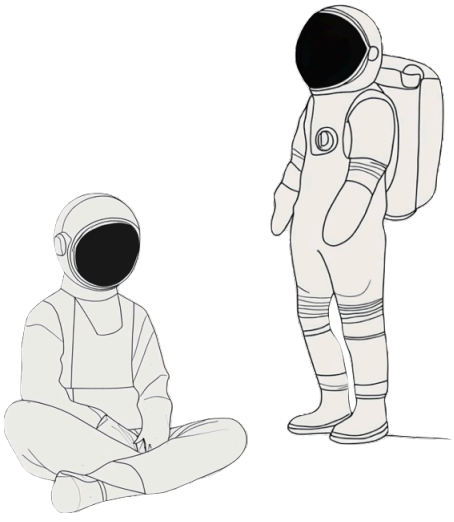

Chapter Summary

• Synth autonomy can be precisely controlled on a spectrum: from fully-supervised to independent operation, while maintaining a balance between behavioral flexibility and predictable boundaries

• Management interfaces combine HR-like dashboards for high-level oversight with granular audit trails of every thought, decision, and action – ensuring both operational visibility and accountability

• Synths are designed for easy skill acquisition through natural description and demonstration, with customizable appearance and communication styles to match organizational preferences

• Comprehensive monitoring capabilities cover all Synth activities and communications while respecting human privacy rights, alongside flexible cost modeling that adapts to different pricing strategies (per-month, per-hour, or per-task)

• Establish AI governance frameworks early and iteratively – starting with strict default permissions (a whitelist approach) and clear data access policies that mirror human worker data governance

• Engage key stakeholders (Security, AI governance, Labor/HR representatives) early in the process to define acceptable use criteria and establish governance requirements for moving beyond pilot phases

• Deploy comprehensive real-time security controls that monitor and can intervene in Synth thoughts, communications, and actions

• Balance autonomy and oversight according to organizational risk tolerance – most organizations start with low autonomy and high monitoring, gradually expanding capabilities as trust and governance mature within a sandboxed environment

IF YOU'RE INVESTING in autonomous intelligent synthetic workers; you'd be more than reasonable to expect a whole suite of management (and governance) capabilities that make Synths productive and trustworthy.

Expect the following:

AUTONOMY AND INDEPENDENCE ARE A CONTROLLABLE SPECTRUM

Synths can behave with varying degrees of autonomy. Expect to control that level of autonomy. Synths should be able to wait idle until prompted by a human, or by another process. You should also be able to increase their autonomy such that they behave completely independently. This high-level autonomy can include things like allowing Synths to develop their own sub-goals, determine their own calendars, or send emails to people on their own.

WHITE LIST? OR BLACK LIST?

Expect the ability to define approved behaviors and tools. OR expect the ability to define behaviors or tools that Synths must NOT do or use. You should expect control over their allowed behaviors and tools.

PREDICTABILITY AND FLEXIBILITY DON'T NEED TO BE A TRADEOFF

Synths can behave with a high degree of behavioral flexibility. Synths should be able to "figure things out": including making their own plans, and trying things out on their own.

You can also expect them to behave predictably. Expect the ability to control 'allowed' or 'disallowed' decisions, or software that a Synth can and can't use. Just because Synths operate their own computers doesn't mean they should be allowed to do whatever they want. Expect enforceable limits and restraints on their behavior. Expect to not be forced to compromise flexibility and predictability.

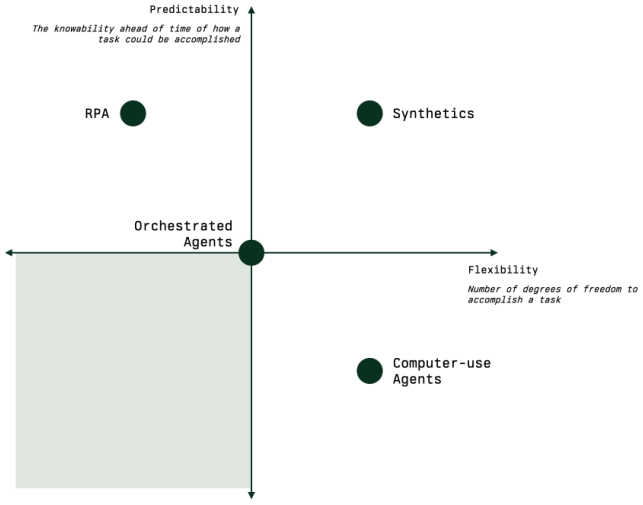

Figure 10: illustrative landscape of the tradeoffs of behavioral flexibility and predictability in common Agentic AI approaches.

EXPECT TO BE ABLE TO MODEL COSTS.

Synths are a cost-effective way to make your organization more productive. Expect to be able to model those costs. Different Synth vendors have different strategies for how to price Synths. Some price them at a flat rate per month (like how staff salaries are priced). Other companies price them according to the hours that a Synth is online (like how cloud computing is priced). Still others price per task that Synths perform. No matter what, expect to be able to monitor the cost of your Synths as they work.

EXPECT SUMMARY INSIGHT INTO SYNTH BEHAVIOR

Expect ‘at-a-glance’ snapshots of how a Synth is behaving. What it is it doing? Who is it talking to? What did it say? What data did it use? What did it accomplish? Expect this kind of snapshot available any time.

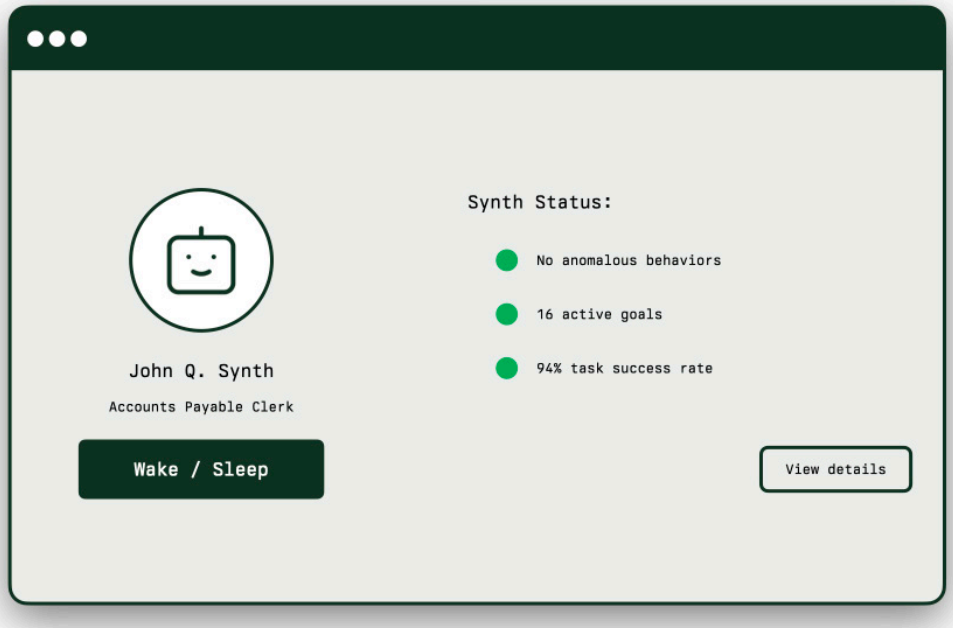

Figure 11: Synth management software should be incredibly straightforward.

EXPECT TO CUSTOMIZE SYNTHS

Expect to control how a Synth appears to your team. This can include things like its name, email address, or handle in Slack or Teams. Some companies opt to give their Synths faces and voices.

In Mission Control AI’s Design Partners Program, we’ve seen that companies have varied opinions about how customizable a Synth should be. Some teams care a lot. Others don’t care at all.

EXPECT TO BE ABLE TO MONITOR HOW SYNTHS COMMUNICATE.

Synths talk to you. Synths talk to staff. Synths talk to one another. Expect to monitor all of this. This becomes increasingly helpful as your Synthetic headcount grows and communication becomes both more frequent, and a great percentage of your total valuable work conversations.

Always make sure that the privacy rights of your human staff are respected when monitoring Human - Synth communications.

Your team will know best how to think about what information about Human - Synth communications can and cannot be stored without workers' knowledge and permission.

EXPECT FINE-GRAINED RECORDS OF EVERY SINGLE BEHAVIOR, TASK, THOUGHT, AND DECISION.

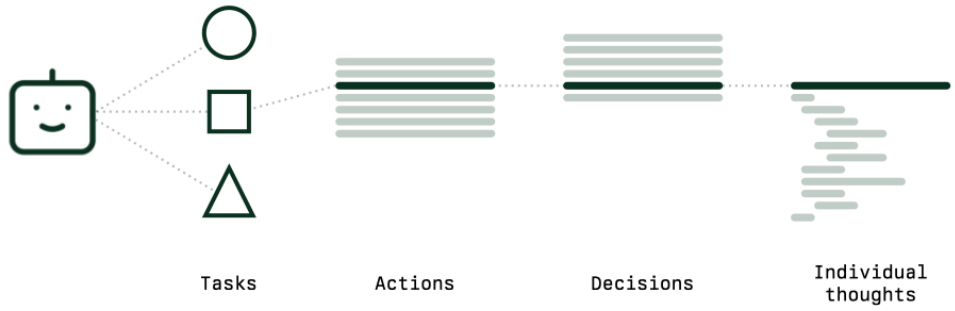

Figure 12: all Synth tasks can be decomposed into discrete actions (motor outputs). All motor outputs have accompanying traces of decisions, and individual thoughts that lead to those decisions. Expect to be able to interrogate all of these.

Expect every facet of a Synth to be recorded, viewable, and auditable. There will be many people in your company who will care about what Synths do, how they think, what data they used to make a decision. Expect everything to be recorded and easy for you to audit.

EXPECT MANAGEMENT OF SYNTHS TO FEEL MORE LIKE AN HR DASHBOARD THAN LIKE AN AI OPS DASHBOARD

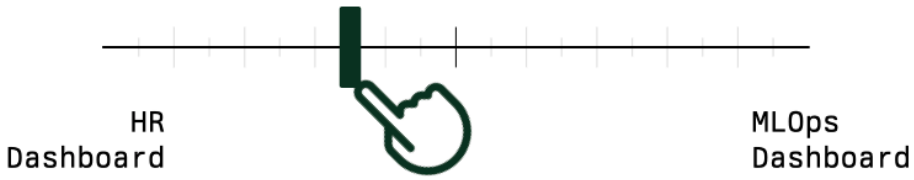

Figure 13: Not to create false dichotomies, but if you had a slider between the options of "I manage Synths with an HR Dashboard" and "I manage Synths with an MLOps Dasboard", we advocate for being just to the left of center towards HR.

Synths are made of a combination of software and AI systems. But they behave like people. They communicate and perform work tasks. They make goals and carry those out to drive business objectives. So expect a management dashboard for Synths to feel more like an HR dashboard for managing staff than like an AI- or MLOps dashboard.

EXPECT IT TO BE EASY TO TEACH SYNTHS

Synths understand the world and their own capabilities based on descriptions that you provide. Expect it to be easy to make those descriptions, or to show a Synth how to do something. Expect to teach new skills to Synths by describing those skills using text, by using your voice, or by showing a Synth yourself so it can watch you.

LIKE OTHER AI-powered technologies you use, Synths will be subject to your AI Governance policies.

EXPLORE GOVERNANCE EARLY. ENFORCE IT WITH INCREASING MATURITY

You don't need to have all the answers to all governance questions when you're starting out with Synths. But the sooner you're aware of what Synth governance will be like, the easier and cheaper it will be to deploy Synths in a way that your organization can trust.

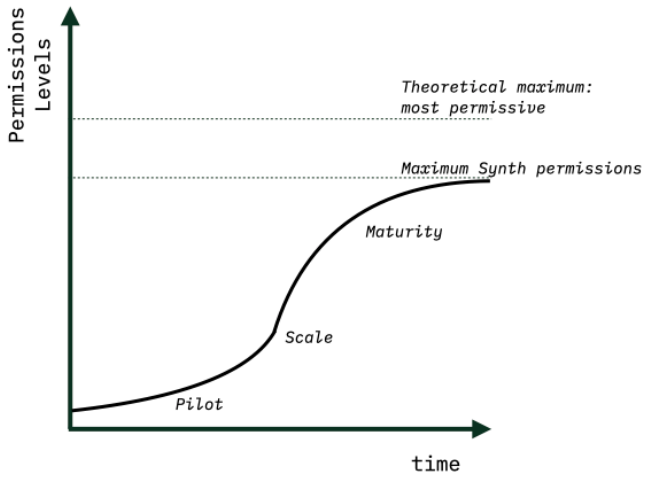

Figure 14: Illustrative maturity of Synth permissions with respect to time.

RESTRICT BEHAVIOR BY DEFAULT

Synths should, by default, be given strict permissions controlling their behavior. We advocate for white-listing allowed actions and software as the default permissions.

DATA GOVERNANCE WILL BE ESPECIALLY IMPORTANT FOR SYNTHS

Imagine a Synth who autonomously (and correctly) decides that the best way to make a purchasing decision is to review past purchases.

Does it have the permission to access that data? In a very real sense, this is very similar to how Data Governance works for human workers. Organizations often have policies that control who is allowed to access different information. Organizations will have similar policies for Synths.

ENFORCE A POLICY OF LEAST PRIVILEGE

Only give Synths the data and tool access that they need to do their jobs well. This is commonly referred to as a "Policy of Least Privilege".

ENFORCE REAL-TIME OPERATIONAL SECURITY GUARDRAILS

Guardrails aren't concepts; they're code.

Expect your Synth vendor to offer real-time guardrail controls on all Synth thoughts, communications, and actions.

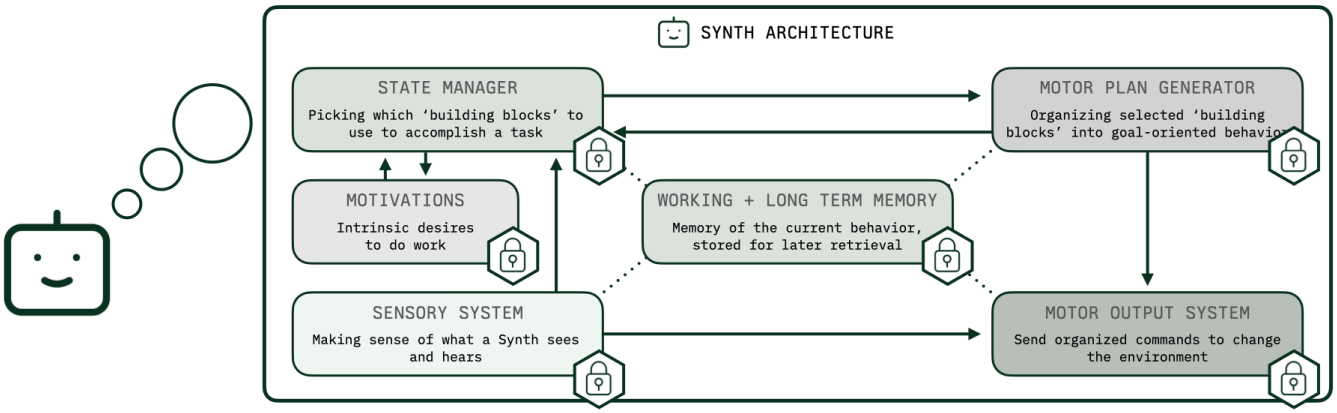

Example Guardrails for real-time , by subsystem

Sensory System Filter for prompt injection attacks. Filter for Acceptable Use violating material (ex: toxic speech and ideology, CSAM, illegal content), common business attack vectors (phishing). Conceptual: filter infohazards.	**Working and Long Term Memory** All working memory and consolidated long-term memories can be interrogated individually. Subject to oversight for learning maladaptive behaviors.
Motivations Continuously evaluate for how stated motivations translate into carried-out objectives. Detect for lapses in inner and outer alignment between motivations and sub-goals; deactivate individual motivations as required.	**Motor Plan Generator** Selected plan steps and internal reasoning can be supervised for undesirable sub-goals (human deception, unmitigated delegation, self-exploration or self-medication behavior) and mitigated or removed.
State Manager Similarly to Motivations; evaluate inner and outer alignment. Use as a de-activation point if need be to arrest potentially unstable or dangerous states.	**Motor Output System** Hard-deny access to programs. Keystroke, mouse-event, API access logging. Enforce content control (PII, Sensitive information) redaction on data leaving Synth.

These should include 'firewalls' that monitor every facet of a Synth to determine if it would be appropriate to intervene. Interventions can apply to thoughts, communications, and actions - and can include redaction, the removal of certain information or ideas, and even blocking behaviors entirely.

Figure 15: Synth architecture with illustrative guardrails for each architectural subsystem. Note that each subsystem's proposed guardrail focuses on both different facets of reducing operational risk, and different styles of interventions.

AI GOVERNANCE IS ABOUT TO MODERNIZE, OR FADE INTO IRRELEVANCY

Between 2019 and 2023, the AI governance community largely focused on 'Responsible AI 1.0' focus areas. These include topics like fairness, explainability, transparency, and data privacy. The preferred tools for animating this were questionnaires and risk management frameworks.

2023 to 2024 saw new inclusions into the state of the art in AI governance thought, largely downstream of the rapid uptake of Generative systems. 'Responsible AI 2.0' focused on the behavior of LLMs at inference time: robustness to adversarial interference, hallucination detection, and data security. The preferred tools were real-time operational security filter layers.

As Synths and highly-autonomous agentic systems spread in adoption over 2025 and beyond, AI governance will rapidly absorb the topics, tools, and concerns of AI Safety.

This was always going to be the case: sufficiently autonomous AI systems require completely new ways of thinking about operational safety beyond the 2019 discourse.

This doesn't mean those facets no longer matter. It means the enterprise will quickly become literate and curious about things like Alignment, autonomy controls, and thought monitoring.

The world prepared for Synths is a world with new guardrails. Systems for monitoring and protecting against undesirable behaviors such as:

Illustrative Responsible AI 3.0 Guardrails	
Inner and Outer Alignment When a Synth forms sub-goals; ensure that those goals also abide by expectations of their main goals	**Deception** Prevent a Synth from acting if it is discovered that deceiving human operators is an acceptable sub-goal
Novel Behavior Take action if a Synth begins to display novel behaviors not previously characterized	**Unexpected Tool Use** Take action if a Synth begins to use new tools, or begins to use tools in new ways not previously characterized
Self-exploratory Behavior Take action if a Synth begins to engage in self-exploratory behaviors	**Self-modification** Take action if a Synth begins to engage in behaviors that would modify its own functionality or capabilities.
Resisting or evading de-activation Take emergency action if a Synth exhibits behavior that resists deactivation	**Plan:Behavior deviation** Take emergency action if a Synth exhibits any unexpected deviation from operator instruction
Reward tampering (Wireheading) Take emergency action if a Synth exhibits any behavior that suggests tampering with its reward signal	**Runaway delegation** Take emergency action if a Synth exhibits any unexpected runaway delegation behavior

Table 3: Illustrative modernized guardrails. The astute reader will note that many of these guardrails more closely resemble traditional AI Safety concerns than traditional Responsible AI concerns. This represents a necessary evolution of the enterprise Responsible AI discourse better matched to an accelerating capabilities landscape.

Why is this happening?

Because Agentic AI and Synths can do entirely new classes of behaviors on their own. By definition, an Agent or a Synth are human-out-of-the-loop systems. A Synth is making decisions and carrying out behaviors autonomously. This unlocks tremendous potential. This also means that there are a greater number of things that can go wrong. The consequences of them going wrong can be much higher.

New ways of thinking about and implementing AI Governance are required.

GET YOUR SECURITY TEAM INVOLVED EARLY

If the adoption of Generative technologies has taught us anything; get your CSO or CISO involved in conversations early. Ask them what they need to see that will make Synths acceptable for use beyond a Pilot.

GET YOUR AI GOVERNANCE INVOLVED EARLY

Likewise, if you have an AI Governance team or risk trigger process for AI, engage those stakeholders early to get a better understanding of how they think about Synths. Ask them what they need to see that will make Synths acceptable for use beyond a Pilot.

YOU MAY HAVE A LABOR REPRESENTATIVE TEAM. GET THEM INVOLVED EARLY TOO

Many organizations have special teams dedicated to advocating for the rights of labor. These are common in industries with staff who are union protected, organizations that interface with organized labor, or in counties with strong labor laws.

If this sounds like your organization, consider bringing the labor teams into conversation about Synths sooner rather than later. It's important to consider their perspectives and requirements as you deploy Synths. Similarly, invite someone from HR to be a part of these conversations too.

CREATE A SANDBOX IN YOUR CLOUD

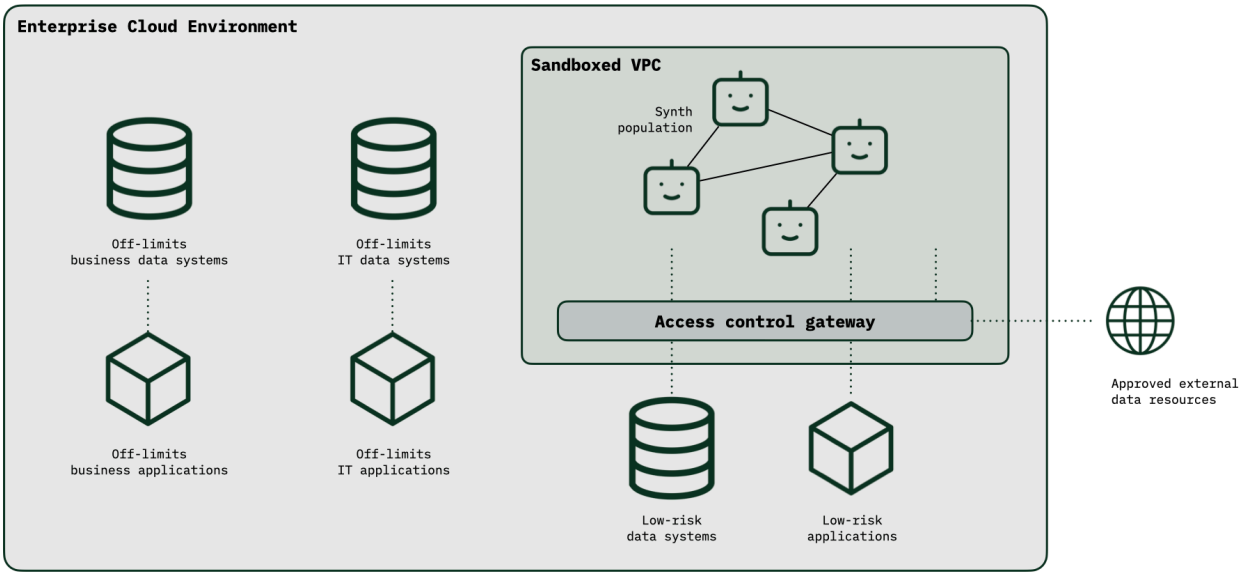

When you're first exploring Synths, it may be appropriate to deploy your Synths into their own sandboxed environment in your cloud. Restricting their access to sensitive information - including IT and business resources - may be a preferable way to create a protected environment for Synths to operate in during a Pilot.

Figure 16: Sandboxing Synths involves creating tightly regulated zones within your enterprise cloud environment in which Synths can operate freely. Access to low-risk data systems and applications can be provided, and privileges may be escalated with time.

DETECT ANOMALOUS EVENTS. AUTOMATICALLY ESCALATE THEM

Anomalies happen. They should be documented and reviewable on a central dashboard. If a Synth thinks, communicates, or behaves in an unexpected way, that should be both detectable, and something that can be immediately alerted to a proper accountable party in your organization.

EXPECT ALL OF THIS OUT-OF-THE-BOX FROM YOUR VENDOR

Synths are a powerful addition to your team. With great power comes great responsibility. Expect your Synths vendors to take operational security and governance seriously, and to offer good ideas and good tools to help you make the most of your Synths without compromising on security or trust.

Chapter 4 Review

• Synth autonomy can be precisely controlled on a spectrum: from fully-supervised to independent operation, while maintaining a balance between behavioral flexibility and predictable boundaries

• Management interfaces combine HR-like dashboards for high-level oversight with granular audit trails of every thought, decision, and action – ensuring both operational visibility and accountability

• Synths are designed for easy skill acquisition through natural description and demonstration, with customizable appearance and communication styles to match organizational preferences

• Comprehensive monitoring capabilities cover all Synth activities and communications while respecting human privacy rights, alongside flexible cost modeling that adapts to different pricing strategies (per-month, per-hour, or per-task)

• Establish AI governance frameworks early and iteratively – starting with strict default permissions (a whitelist approach) and clear data access policies that mirror human worker data governance

• Engage key stakeholders (Security, AI governance, Labor/HR representatives) early in the process to define acceptable use criteria and establish governance requirements for moving beyond pilot phases

• Deploy comprehensive real-time security controls that monitor and can intervene in Synth thoughts, communications, and actions

• Balance autonomy and oversight according to organizational risk tolerance – most organizations start with low autonomy and high monitoring, gradually expanding capabilities as trust and governance mature within a sandboxed environment

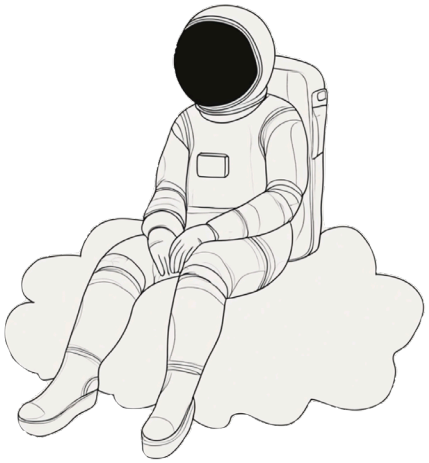

5

Readiness, Piloting, and Deploying Synths

Chapter Summary

• Cultural framing should align with organizational maturity – innovative cultures can embrace transformative potential, while traditional organizations benefit from positioning Synths as natural extensions of existing capabilities (GenAI, RPA, BPO, or AI Agents)

• Leadership evaluation criteria focus on three key dimensions: risk/value tradeoffs, ease of business unit integration, and breadth of applicability – which shape internal decision-making and external communications

• Frame Synths as net-new productive capacity rather than headcount reduction – organizations that emphasize growth and augmentation over replacement see better adoption, morale, and performance outcomes

• Success hinges on cultivating cultural safety where employees feel secure to experiment with human-Synth collaboration, leading to higher quality work and innovative "cyborg" workflows rather than resistance and attrition

• Implementation follows a clear maturity curve: start with 1-2 Synths in pilots (TRL 5-6) with constrained autonomy and high supervision, progressing to fleet deployment (TRL 7-9) with increased autonomy and sophisticated governance

• Initial pilots prioritize learning over ROI, focusing on building institutional knowledge and understanding risk/value profiles before scaling – this exploration phase is critical for developing internal literacy and capabilities

• Begin with internal-facing roles that minimize brand risk: typically internal functions where performance can be refined without customer impact – allowing organizations to develop competency before expanding to more sensitive areas

• Performance expectations should mirror human workforce dynamics: expect varying levels of capability across different tasks, with some surprising strengths and limitations that become apparent through practical application

BEING READY IS AS MUCH BUSINESS PROCESS AS IT IS CULTURE

If your organization's culture is innovative, Synths are an exciting new way to work. You'll unlock an incredible amount of productive capacity and new organizational capabilities working with Synths.

If your organization's culture is more traditional, Synths can be thought of as a logical extension of some things you might already be doing.

SYNTHS ARE LIKE WHAT YOU ALREADY DO, JUST BETTER

How organizations talk about Synths depends on what they do already.
In general, we advocate for finding simple and familiar ways to talk about it.

For example:

Does your company already use Generative AI? Synths are like Generative AI.
Does your company already use Robotic Process Automation? Synths are like RPA.
Does your company already have a Business Process Outsourcing function? Synths are like BPO.
Is your company interested in AI Agents? Synths are like AI Agents.

In reality, all of these statements are true.

Under the hood, Synths use Generative AI. They extend it to do impactful new things because of the way they are built.

Synths do what RPA does. They extend it to more open-ended and flexible tasks.

Working with Synths is like working with BPO; but instead of work going overseas, it goes to the cloud.

Synths are what many wished AI Agents were; extended to perform work on their own motivation.

Chances are good that your organization uses (or is interested in) at least one of those four things.

Synths are a logical extension of what you're already doing.

BACKSOLVE FOR WHAT HELPS YOUR LEADERSHIP MAKE GOOD DECISIONS

Your leadership is going to have their own expectations of what they want to see from experimenting with Synths.

Broadly speaking, they're going to want things that allow them to learn, and communicate. Their learning objectives are going to focus on (1) the risk / value tradeoffs of working with Synths, (2) the ease of integration into existing business

units, and (3) the breadth of applicability.

This allows leadership to communicate priorities internally. Internal communications goals depend on how an organization makes decisions. Different organizations have different cultures around how decisions get made. These can range from consensus building to policy implementation.

Externally, leadership teams will seek to communicate with Boards, investors, regulators, potential customers, and the supply chain - amongst other stakeholders.

YOU GET TO DETERMINE HOW THE EMPLOYMENT CONVERSATION GOES

We're way past pretending that labor doesn't care about AI. Is it all negative? No; there's a lot of enthusiasm.

There's also a lot of trepidation. It's pretty obvious why.
Employers are doing a miserable job at communicating the plan for AI with staff.
And staff read the news. They hear of another firm laying off thousands in favor of an AI system.

Your employees aren't stupid. And a certain point, you're going to have to talk with them about AI and its relationship to work.

We've observed that the teams with the highest preparedness look at Synths as an opportunity to do more as an organization; not to reduce headcount.

To them, Synths are not about job displacement. Synths are about doing new activities, and doing current activities in new ways.

To them, synths are net-new productive capacity; not an excuse for 'right-sizing'.

There are obvious reasons why organizations chose to think about Synths as net-new capacity building. Ultimately; it boils down to culture and safety at work.

Organizations with cultures with a high-degree of emotional safety (in which workers feel secure in their employment and free to experiment with how they work) outperform those that don't cultivate a culture of safety.

That's it. That's the best predictor.

So if your employees perceive that investing in advanced AI is a cost-cutting measure: expect cultural resistance, challenges with morale, a drop in the quality and velocity of work, and increases in worker burnout. This is a formula for expensive and harmful employee disengagement and attrition.

Alternatively, if your employees perceive investing in Advanced AI to be how everyone can do more and do their job better; expect a more performant organization.

Expect exciting new experiments in cyborg work; in which humans and Synths co-create together. Expect higher quality and velocity with low costs.

Ultimately; that's up to you. That's why we call this a culture question. It's about how you talk about Synths; and how well you walk that talk.

Getting started working with Synths requires different processes to be at different stages of maturity. Not everything needs to already be figured out ahead of time; but some things should be relatively mature.

Needs to be in-progress before Synths are in prod	Need to be at maturity before Synths are in prod
An AI Governance Practice A process for measuring and reducing the risks associated with AI technologies that are moved beyond pilot stages.	**A Generative AI adoption strategy** How do we think about AI as a value driver? What are our bets? Where do we need to learn more? Whether or not to us Synths flows naturally from your Generative and Agentic AI strategy.
Executive Sponsorship of AI experiments Who should 'own' Synths? Is this under the CTO? The CIO? Is this an AI project? A BPO project? An IT project? It's good to figure this out sooner than later.	**Access to LLMs** Do we have access to frontier models? (Even if provided from cloud service providers, such as through Azure AI Services or Amazon Bedrock)
Ops Readiness Do we know what business units Synths would most naturally integrate in? Have those leaders already been brought into the conversation?	**An IT Approval Process** How do we 'say yes' to trying new technology? How will Synths be treated by that process? How will they be different?

Table 4: Organizations don't need to have everything figured out before deploying Synths into production. But there are common desires of basic practical preparedness. These break down into "this needs to be started but not done (right)", and "this needs to be mature".

WHAT DO YOU NEED TO ALREADY BE MATURE IN?

Expect to already be mature in the steps that would allow you to say yes to trying new AI technology. That includes having an AI strategy in place, having access to the underlying LLM technologies you need, and a process of onboarding new technology.

WHAT DO YOU NEED TO HAVE IN PROCESS (BUT NOT COMPLETE YET)?

Expect to at least be "in-progress" in the basic steps that allow you to make sense of a new AI technology. This includes executive championship, where a new technology is going to be applied, and what governance controls will be applied to it.

No organization goes from a standstill to a full deployment of a Synthetic workforce overnight.

Rather; organizations build up institutional knowledge, experiment, and come to build both internal know-how and business cases through experimentation with Synths.

TECHNOLOGY READINESS AND PHASED ROLLOUTS

Generally, organizations expect the following:

Phase	Quantity of Synths	Business Objective	Technology Readiness Level	Synth Capabilities
Pilot, Design Partnership	1-2	Innovation and internal literacy building Experimentation	**TRL 5 - 6** Technology validated in relevant environment Technology demonstrated in relevant environment	Low autonomy Moderate accuracy Constrained behavior
Production	Ramp up pattern changes case by case	Fleet rollout Extending business unit capabilities Driving ROI	**TRL 7 - 9** System prototype demonstration in operational environment System complete and qualified Actual system proven in operational environment	Moderate - High autonomy High accuracy Online governance Open-ended behavior

Table 5: what's expected from Synths at different stages of deployment maturity. Between pilots and production there is an expected maturation of capabilities and reliability. These form the basis for justifying a larger organization-wide rollout.

Several principles can be distilled here:

START SMALL AND LEARN FAST

Organizations expect to deploy only one or two Synths during a pilot program. This constrains both cost and ROI according to the relevant business objective of that phase: to accelerate learning.

Starting small allows an organization to carefully build out both the business case and risk profile for Synths.

A pilot phase is largely geared towards building institutional know-how. This is less about ROI and more about building an internal literacy with the technology.

SUPERVISION IS INCREDIBLY IMPORTANT AT THE OUTSET

Making sure Synths are adequately supervised (by both humans and software infrastructure) is incredibly important. It's especially important during pilot phases in which the risk/value ratio of Synths is still being discovered. This has two obvious implications:

Expect Autonomy to be low during pilots.

There's an optimal amount of autonomy to expect: enough to make the business case for exploring Synths as a valuable part of a team, and not so much to outpace demonstrated value. Think of autonomy as something earned.

Expect Supervision to be strict: including operational supervision (ie: supervision of the Synth itself while its operating, like "over-the-shoulder" testing), live guardrails, and detailed logging of Synth thoughts, decisions, and actions.

EXPECT TO START WITH INTERNAL JOB ROLES AND TASKS

Understandably, organizations are protective of their customer experience.

Brand is one of the most valuable assets an organization can hold. And there's good reason to be conservative with brand-exposing activities and the introduction of new enabling technologies that could material impacts brand.

Organizations opt to pilot Synths "internally": in cost-center job roles and business units that minimize brand risk and external exposure. This allows an organization to develop internal capabilities, know-how, risk profiles, and business cases for Synths with minimal exposure to a customer.

Expect to pilot Synths in internal roles first. As the risk / value case becomes more clear unique to your organization, your evaluation of the relative strengths and weaknesses of Synths as a part of your team will become clearer.

EXPECT DIFFERENT LEVELS OF PERFORMANCE AND FRICTION

Expect Synths to behave in the workplace similarly to how human workers behave. Not identically, but similarly.

Expect Synths to be very good at certain types of tasks that humans have a hard time with. And likewise, expect them to struggle at some things that humans find intuitive and easy.

(And more still: expect to be surprised along the way at which tasks are hard and which are easy).

Partnering with Synth Vendors

Organizations have standard expectations from the companies they partner with for delivering Synths and other advanced AI Agent technologies. These are orienting; they reveal internal priorities for both maximizing value and managing risk.

A few trends become apparent:

• Organizations must first clearly define their priorities: from technical transparency requirements to system integration preferences. These sector- and scale-specific needs shape vendor selection and implementation approach

• Expect to actively shape the technology through real-world deployment feedback. Vendors should embrace enterprise partners as co-developers rather than just customers of a finished product

• Vendor solutions should complement existing infrastructure and staffing strategies, preserving investments in current systems (ERP, CRM, DevOps) and respecting ongoing human resource planning

• AI Governance standards of vendors must match or exceed your own requirements. Expect vendors to demonstrate commitment through practical safeguards and controls, not just policy statements, even during pilot phases

ASSESSING REQUIREMENTS

At risk of stating the obvious, organizations find themselves better prepared for exploring Synths when they first understand their own priorities. Not every organization has the same priorities; these may be very sector, region, and scale specific.

These can include things like:

• "We require working with specific partners that we already have relationships with"
• "We have a mission critical imperative to be very 'under-the-hood' about every facet of Synth behavior"
• "We need institutional knowledge and results; we don't care how Synths work"
• "Beyond tailoring the Synths roles, we want to customize them to new environments"
• "We want to shape the Synth management environment and their governance"
• "We don't want to migrate to any new software or data systems" - or - "We're comfortable migrating if the ROI is very clear"

Organizations share common expectations from their Synth vendors:

EXPECT TO SHAPE THE TECHNOLOGY

The enterprises' feedback matters a lot in shaping wholly new technologies like Synthetics and advanced Agentic AI. Every vendor has whiteboard-based reasoning that rarely withstands the reality of a partner using their technology.

The partner is going to know what they need better than the vendor. Expect to play a dynamic role in shaping this technology early on.

EXPECT RESPECT FOR YOUR CURRENT STAFFING AND HIRING PLAN

The entire field is still coming to understand the impact and reliability of Agentic AI and Synthetics. Expect your vendors to respect your staff, and respect your current hiring priorities. No organization on earth has the necessary proof points to suggest that you should stop hiring humans. Dont believe the hype.

EXPECT TO KEEP YOUR SOFTWARE

Expect a vendor to be compatible with the ways you currently do things.

If you have software you already use, keep it. Vendors should be complementing and enhancing the investment you're already making in systems, not replacing them.

EXPECT VENDOR PARTNERS TO TAKE AI GOVERNANCE SERIOUSLY

Expect a vendor partner to have AI Governance practices and technology in place that are in parity with your own. You may be more lenient during Pilot phases where Synths are explored for their potential impact.

EXPECT YOUR VENDORS TO PLAY NICELY TOGETHER

Synths will be multi-vendor, and will need to integrate elegantly with your other vendors. Expect your Synth vendors to support the use of on-premise, CSP provided LLM providers (such as AWS Bedrock, Azure AI Services, or OCI).

Expect Synths to use the other AI tools you're getting from other partners, such as AI co-pilots or chat systems. Expect your Synth vendors to enable you to draw down on investments or partnerships you've made with hyperscalars.

Chapter 5 Review

• Cultural framing should align with organizational maturity – innovative cultures can embrace transformative potential, while traditional organizations benefit from positioning Synths as natural extensions of existing capabilities (GenAI, RPA, BPO, or AI Agents)

• Leadership evaluation criteria focus on three key dimensions: risk/value tradeoffs, ease of business unit integration, and breadth of applicability – which shape internal decision-making and external communications

• Frame Synths as net-new productive capacity rather than headcount reduction – organizations that emphasize growth and augmentation over replacement see better adoption, morale, and performance outcomes

• Success hinges on cultivating cultural safety where employees feel secure to experiment with human-Synth collaboration, leading to higher quality work and innovative "cyborg" workflows rather than resistance and attrition

• Implementation follows a clear maturity curve: start with 1-2 Synths in pilots (TRL 5-6) with constrained autonomy and high supervision, progressing to fleet deployment (TRL 7-9) with increased autonomy and sophisticated governance

• Initial pilots prioritize learning over ROI, focusing on building institutional knowledge and understanding risk/value profiles before scaling – this exploration phase is critical for developing internal literacy and capabilities

• Begin with internal-facing roles that minimize brand risk: typically internal functions where performance can be refined without customer impact – allowing organizations to develop competency before expanding to more sensitive areas

• Performance expectations should mirror human workforce dynamics: expect varying levels of capability across different tasks, with some surprising strengths and limitations that become apparent through practical application

6

Getting Started

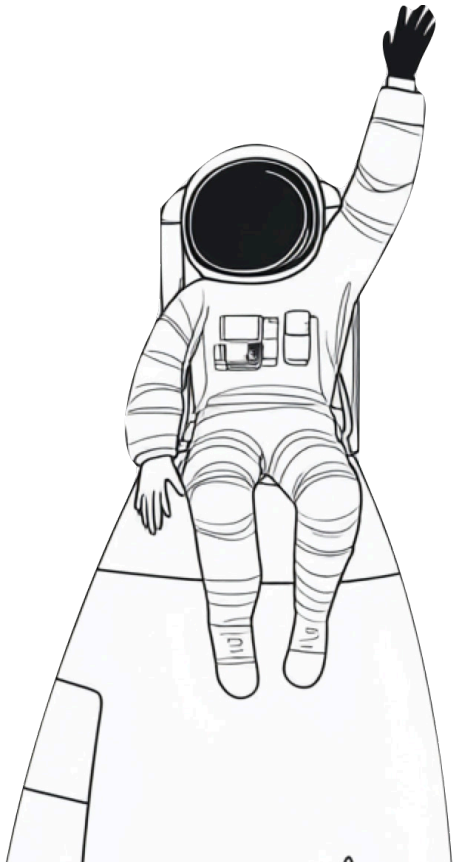

Chapter Summary

• Design Partner programs for Synths require foundational practices in Generative AI, AI Governance, and RPA; even if the practices are in early stages.

• Effective Design Partner collaborations focus on identifying suitable job roles, knowledge transfer, AI governance, integration with existing systems, and customizing management tools.

• The primary goal of Design Partner programs is mutual learning and growth, recognizing the evolving nature of Synth technology and the value of partner insights.

• Mission Control AI's Design Partners Program offers hands-on Synth experience, bi-weekly design sessions, and the opportunity to shape Synth development. Cohort 2 starts January 2025 with 12 selected teams. Reach out to join.

If your organization is gearing up to explore Agentic AI and Synthetic Workers, we invite you to join Mission Control AI's Design Parterns Program, Jan 2025.

WHAT TO HAVE IN PLACE ALREADY

Organizations exploring AI Agents and Synths usually already have:

- A Generative AI practice (even in pilot stages)
- An AI Governance practice (even in low maturity)
- An RPA practice already in place (even at low maturity)

If your organization already has some combination of those three in progress, you're likely ready to extend your AI Strategy to include Agentic AI and Synthetics.

WORKING WITH DESIGN PARTNERS

Beyond providing Synths for you to deploy and work with, Design Partners will work with you to identify:

1) Which job roles make the most sense for Synths in your organization.

2) Applicable Knowledge Transfer documentation or Business Process Outsourcing reference material or practices apply.

3) AI Governance and Data Security expectations beyond a Pilot.

4) How Synths extend your existing work with Generative AI or Robotic Process Automation.

5) What changes to their Synth management infrastructure or dashboard tools enable your requirements.

THE GOAL OF A DESIGN PARTNERS PROGRAM SHOULD BE THAT EVERYONE GETS SMART

We've crafted our Synthetics Design Partners program around a single goal: that we and our partners get smarter. That's our North Star.

Why this?

Because this field is so new.

Everyone is learning. It's an exciting and emerging market with potential to transform how work gets done. And every team - from startup to hyperscalar - is currently still

learning how it works. No one has all the answers yet.

Accelerating institutional literacy on Synthetics is the only path to ROI. No one gets to "this made a big difference in our operations" before passing through "we have an intuition for what to do with Synths".

So a great Design Partner program should be a relationship. With that humility in mind; the North Star should be building and learning together. Mission Control AI gets smarter, our partners get smarter.

With that North Star in sight, we structured our Design Partners Program around tangible deliverables, emphasizing knowledge transfer, and giving our partners an opportunity to steer the field.

A GREAT SYNTHS PROGRAM COMBINES FLEXIBLE LEARNING WITH A PREDICTABLE STRUCTURE

A successful Synths pilot program strikes a balance between structured learning and adaptability.

At its core, the Systems Delivery Architecture in Mission Control AI's Design Parter's Program provides a predictable roadmap that our partners follow with us. This architecture is a backbone for the key stages, milestones, and expectations of the program.

This backbone is paired with flexibility: as partners engage with the technology and we uncover new insights together, the program allows for the pilot to evolve organically. This keeps learnings relevant and valuable, even as understandings of Synths deepen and new use cases emerge.

A well-designed Synths Design Partners Program makes RACI roles explicit from the outset, and details deliverables, timing, and requirements. Clarity allows all parties to plan, allocate, and set realistic expectations for the program's outcomes.

A GREAT SYNTHS PROGRAM DELIVERS

Learning is important. But, in practice, a Synths pilot program needs to provide partners with access to Synths. Mission Control AI built our Design Partners Program to deploy Synths that are tailored to the specific requirements of each partner, with minimal partner overhead or implementation burden.

But delivery isn't a one-time event.

Throughout the program, partners ought receive continuous learnings and findings. These insights include best practices for Synth deployment, unexpected capabilities or limitations discovered during testing, and co-developed strategies for measuring ROI.

Continuously sharing knowledge ensures that our partners gain value not just from the technology itself, but from the experience of participating in the program.

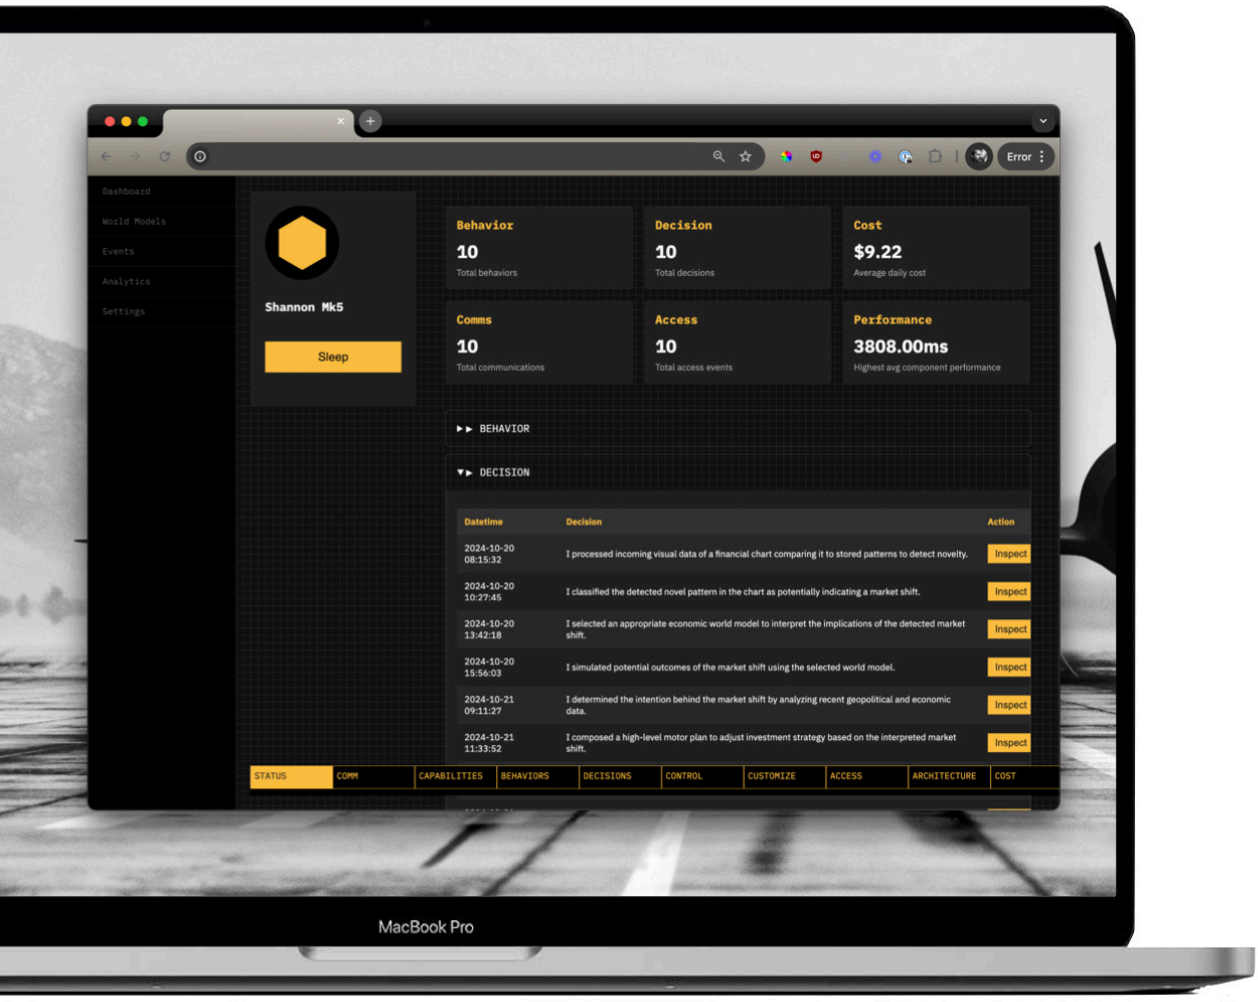

Figure 17: Mission Control AI's *Synth Launchpad* platform for managing Synths in the cloud.

A GREAT SYNTHS PROGRAM IS A TWO-WAY STREET OF LEARNING

The most effective Synths programs recognize that knowledge flows in both directions. Mission Control AI has asymmetric expertise in Synth technology. Our partners have asymmetric insights about their specific industries, workflows, and challenges. Information exchange means everyone gets smarter about each others' world.

Partners develop a nuanced view of how Synths can transform their operations, while Mission Control AI gains insights that shape the future development of the technology. This knowledge transfer is innovation fuel.

A GREAT SYNTHS PROGRAM TURNS YOUR FEEDBACK INTO REALITY

The true measure of a Design Partners Program is how well it translates partner feedback into tangible improvements.

Rather than treating partner input as mere suggestions, a great Design Partners

Program frames this feedback as an innovation driver. The collective knowledge and experiences of partners become the foundation for enhancing both the Synths themselves, and the infrastructure used to manage them.

Partners should expect their specific requirements to directly influence the development of new capabilities. That may be an enhanced data integration with existing systems, improved performance in certain tasks, or new features to address industry-specific challenges. Our partner's input is visibly reflected in ongoing development efforts.

The impact? Synths evolve to better meet the needs of partners.

Synths become more capable, efficient, and adaptable. *Synth Launchpad* (our management dashboard and governance infra) is refined continuously to provide partners with greater visibility, control, and ease of use. This dual focus ensures that partners see value not just in Synths themselves, but in the entire ecosystem supporting their deployment and operation.

YOU CAN JOIN OUR DESIGN PARTNERS PROGRAM

Mission Control AI is excited to announce the launch of Cohort 2 of our Design Partners Program, commencing in January 2025.

We are accepting applications and extending invitations to 12 teams to join us in shaping the future of Synthetic work. This program offers a unique opportunity for the enterprise to gain hands-on experience with Synths and contribute to development as a Design Partner.

Partners selected for the program will receive their own Synths, allowing for direct experimentation and integration within their existing workforce. This hands-on access enables partners to fully explore the potential of Synths in addressing their specific business challenges and uncovering new opportunities for innovation.

To ensure the program's success and maximize its impact, executive sponsorship is required from each partner organization. Participants will engage in bi-weekly 1:1 design sessions for knowledge exchange and co-development.

Typically, participation in the Design Partners program is funded through CIO, AI, or innovation budgets. This alignment with strategic technology initiatives ensures that the program receives the necessary resources and attention to drive meaningful outcomes for partner organizations.

We invite teams eager to gain expertise in Agentic AI and Synths to apply for participation. Selected teams will be onboarded on a rolling basis.

Visit usemissioncontrol.com to apply.

Mission Control AI is accelerating the cyborg workforce.

This includes enterprise-grade Synthetic Workers (and the necessary platforms for their management and governance), technology for accelerating human - AI work teaming (GenOps), and decentralized protocols for reputation management of self-sovereign AI systems.

You can deploy your own Synthetic Workers. Visit usemissioncontrol.com.