



SARAH HEFFRON ADVISORY

THE LIVING ARCH

A human-centric, AI-enabled, continuously adaptive approach to operating model design

Sarah Heffron

sarah@sheffronadvisory.com | sheffronadvisory.com

WHAT THE LIVING ARCH™ FRAMEWORK IS

The Living Arch framework is an approach to operating model design for organisations adopting AI. It is the design response to the diagnostic: where the Seven Signs™ diagnostic establishes where an organisation is, the Living Arch framework guides where it goes next and how it gets there.

It rests on three principles. The operating model should be human-centric, placing people at the centre of organisational value. It should be AI-enabled, applying AI deliberately to amplify what people can do and to solve structural problems that were previously too hard to fix. And it should be continuously adaptive, designed to keep changing as technology and conditions change, without requiring a transformation programme each time.

These are not three separate ideas that happen to sit together. They are an interlocking system, and the framework's central argument is that they only work as a system. An operating model that is AI-enabled but not human-centric destroys the value it was meant to create. One that is human-centric but not continuously adaptive cannot survive the pace of change it now faces. The principles protect each other. That is what makes this a framework rather than a list of good intentions.

WHY LIVING ARCH

An arch is one of the oldest load-bearing structures in engineering. It stands because every stone is held in place by the ones around it, and because the keystone at its crown locks the whole structure together. The arch can span further and carry more than any single element could support alone. And a living arch is one designed to flex under changing load rather than crack when conditions shift.

In the Living Arch framework, the operating model is that arch. The human is the keystone: the load-bearing element at the crown that the entire structure depends on. AI is what gives the arch its reach and its capability, allowing it to span further and carry more than it could unaided. And the continuously adaptive design is what allows the arch to flex under changing load rather than cracking when conditions shift.

The mapping is precise rather than decorative. It captures the framework's core claim: that the human is not one component among many but the element the entire structure depends on. An operating model can have sophisticated AI, elegant processes, and clear governance, and still collapse if it removes the people who were holding it together. The keystone is load-bearing. So are the people.

WHAT LIVING ARCH IS A RESPONSE TO

The Living Arch framework is a deliberate counter to AI-first operating model design.

The AI-first school places the technology at the centre of the operating model and arranges everything else around it. In this view, people are the periphery: users of the system, risk factors to be managed, change-management subjects to be moved from a current state to a future one. The design question is how much human activity can be automated, and the measure of success is how much cost can be removed.

This approach is not just philosophically uncomfortable. It produces worse outcomes, including worse financial ones, and it does so predictably.

It mistakes cost reduction for value creation. Removing people removes cost, and on a short time horizon that looks like success. But the relationships, judgement, experience, and undocumented knowledge that people hold are the source of organisational value. Automating them away does not transfer that value into the system. It destroys it and replaces it with savings that do not account for what was lost.

And it generates risk that the savings do not offset. Reputational risk, when over-automation is experienced by customers or the public as degrading. Regulatory risk, particularly acute in financial services and other regulated environments where AI use is subject to scrutiny and accountability requirements. Customer relationship and retention risk, because the relationships that drive loyalty and lifetime value are human ones. Institutional knowledge risk, because the undocumented knowledge that lives in people leaves with them and cannot be reconstructed. Operational risk, when AI is deployed into processes it does not fully understand, without the human judgement that used to catch the exceptions. And a future-looking risk that is the most strategically serious of all.

Maximal automation is a bet that you know where AI is going and what the future will require. Nobody knows that. Committing fully to one configuration assumes a certainty about the future that does not exist, and the bet is asymmetric. If you are right, you have saved some cost. If you are wrong, and you need human capability back, rebuilding it is slow and expensive, and several organisations are already discovering exactly how hard it is to reverse. The downside is far larger than the upside.

The Living Arch framework is built on the opposite premise. People are the centre. AI is applied in service of them. And the operating model is designed to keep its options open, because the future is uncertain and the organisations that thrive will be the ones that can adapt rather than the ones that bet everything on a single view of what is coming.

THE THREE PRINCIPLES

Principle One: Human-Centric

Human-centric means that people are the most important element in the organisation, and the operating model is designed around that fact rather than in spite of it. It rests on four ideas.

Primacy. People are the most important element in the organisation. This is the starting axiom, not a sentiment added at the end. Everything in the operating model design proceeds from it. If people are the source of organisational value, the model must be built to protect and amplify that value, not to minimise it.

Knowledge. People know best how things work today. They hold the operational truth of the organisation, including the undocumented reality that no process map captures and no system records. This makes them indispensable to honest diagnosis: you cannot understand the current state without them. It also makes them indispensable to design: the people who know how the work actually happens are the ones who can shape what it should become.

Agency. People are not only sources of information about the present. They are critical to shaping what comes next. In a human-centric operating model, people are participants in the design, not subjects of it. They are not moved through a change that is done to them. They help create the change, which is also what makes the change hold.

Humanity. People have feelings, fears, and needs. In a human-centric model these are not friction to be managed around or risks to be mitigated. They are real, they are legitimate, and they must be built into the plan and addressed continuously, not handled in a one-off change-management phase that ends when the programme does. An operating model that ignores how change feels to the people living through it is not a complete design.

Running through all four is a single organising question, the question that connects human-centric to everything else in the framework.

How can AI enable this person, this role, this team to do better, to do more, and to contribute to their fullest potential?

That question reveals the framework's deepest commitment. The point of the AI is human flourishing. The efficiency is real, but it is in service of the people, not the other way around. The human is the end. The AI is the means. An operating model designed around this question produces something that maximal automation cannot: people doing better work, more of it, growing into their fullest contribution, in an organisation that is stronger because of it rather than hollowed out.

Human-centricity is not something that begins once design starts. It begins in the diagnostic phase, in how the organisation chooses to understand itself. Diagnosing with people, treating them from the first conversation as the holders of operational truth and as participants rather than subjects, is already the first act of the human-centric approach. It is also where the culture that a continuously adaptive model depends on starts to form. The agency and trust that allow an organisation to keep changing are built, or undermined, in how it chooses to diagnose, long before the design work begins.

Principle Two: AI-Enabled

AI-enabled means that AI is applied deliberately, in service of the human contribution and the health of the operating model. It rests on three ideas.

Deliberate, not maximal. The governing discipline is thoughtfulness, not technical capability. The fact that AI can do something is not a reason to use it for that purpose. AI can technically write the content, develop the code, and handle the customer service. The question is never whether it can, but whether doing so serves the human contribution and creates real value. Applied this way, AI delivers its greatest value precisely because it is directed where it matters rather than spread everywhere it is possible.

Two purposes, equally weighted. AI-enablement does two things, and they matter equally. It enables people, removing the low-value, repetitive, transactional, and data-heavy work that consumes human capacity, and supplementing human activity so people can do more and do it better. And it solves structural operating model problems that were previously too complex, too expensive, or too burdensome to fix. The same deliberate application can do both at once: addressing operational debt that traditional redesign could never reach, while freeing and amplifying the people who work within the model. This is the two-birds logic. Fixing the structure and enabling the people are not competing uses of AI. They are the same use, done well.

The cost of getting it wrong. Applying AI maximally rather than deliberately does not merely fail to add value. It destroys it. It removes the human contribution that was the source of organisational value, delivers short-term cost savings that are mistaken for real value, and generates the reputational, regulatory, customer, institutional-knowledge, operational, and future-looking risks that those savings do not offset. The deliberate application is not the cautious option. It is the one that actually creates value and protects against the asymmetric downside of getting it wrong.

Principle Three: Continuously Adaptive

Continuously adaptive means the operating model is designed to keep changing as conditions change, without requiring a restructure or a transformation programme each time. There is no stable business-as-usual to return to, because the pace of change in technology and markets no longer allows for one. The model must be built to absorb change as a permanent condition rather than to survive it as an occasional event. This principle has two halves, and they depend on each other.

The structural half: an operating model built to change. The model must be able to sense, to decide, and to act. It monitors for signals that something needs to change, whether in technology, market conditions, or its own performance, rather than waiting for a crisis to reveal that it no longer fits. It has the decision architecture to act on what it senses, because sensing without the ability to decide and act is merely awareness of a problem. And it is able to change its

own details in response, adjusting how work flows and how value is delivered without requiring the whole structure to be torn down and rebuilt.

The cultural and human half: an organisation where adaptation is possible. People do not naturally thrive in uncertainty, and uncertainty is now the permanent condition. This is the human reality the model must be designed around, not designed in spite of. It means the human-facing disciplines have to change. Change management, people management, and communications were built for a world where organisations changed, then stabilised, then changed again. That world is gone. These disciplines now have to engage in a fundamentally different way, supporting people through continuous change rather than through discrete transitions with a beginning and an end.

The two halves are inseparable. A structurally adaptive model staffed by people who cannot cope with continuous change will not adapt. It will break under the strain. A resilient culture sitting on top of a rigid structure will exhaust itself trying to force change through a model that resists it. Genuine adaptive capacity requires both: a structure built to change and a culture and workforce supported to change with it.

WHY IT IS A SYSTEM

The three principles are not independent. Each depends on the others, and each protects the others. This is the most important thing to understand about the framework, because it is what separates Living Arch from a checklist of modern-sounding design values.

Human-centric depends on AI-enabled. Without AI applied deliberately to remove low-value work and amplify human capability, the commitment to human flourishing is an aspiration without a mechanism. AI-enablement is what gives the organisation the capacity to let people focus on the work only they can do.

AI-enabled depends on human-centric. Without the human-centric premise, AI-enablement collapses back into AI-first. The discipline of deliberate rather than maximal application only makes sense if people are the point. Remove the human-centric commitment and there is no principled reason not to automate everything that can be automated.

Continuously adaptive depends on both. It depends on human-centric because the people are the ones who sense that change is needed, who decide and act, and who carry the organisation through change, and because the culture that makes continuous change bearable is itself a human-centric achievement. That culture does not begin when the new model goes live. It begins in the diagnostic phase, in whether the organisation chooses to understand itself with its people or about them. It depends on AI-enabled because AI provides the sensing mechanisms and frees the human capacity that continuous adaptation requires.

And both human-centric and AI-enabled depend on continuously adaptive, because it is the structural answer to the future-looking risk. A continuously adaptive operating model does not make irreversible all-in bets. It preserves optionality. It keeps the human capability that allows the organisation to respond when the future turns out differently than expected, which it will. The principle that keeps the options open is what protects the organisation from the asymmetric risk that AI-first thinking ignores.

This is why the framework holds together. Adopt one principle without the others and it fails in a predictable way. Adopt all three and they reinforce each other into something more durable than any of them alone.

DIAGNOSIS BEFORE DESIGN

The Living Arch framework is a design framework, and design follows diagnosis. You cannot redesign a foundation you have never seen.

The diagnosis is structured around the Seven Signs™, the framework that identifies the most common patterns operating models carry. For organisations specifically assessing their readiness for AI adoption, that same framework is applied through a dedicated self-assessment tool. Current tools are available at sheffronadvisory.com and will evolve as the underlying technology does; the framework beneath them does not. Whichever form the diagnosis takes, it establishes where the organisation actually is: where its operating model foundations are sound, where they are not, and where AI deliberately applied could address weaknesses the organisation may not even know are costing it.

That assessment is what makes Living Arch design possible. You cannot design a human-centric operating model without understanding how people currently work and where the infrastructure supports or constrains them. You cannot design AI-enablement without understanding what the data and process foundations actually are. You cannot design a continuously adaptive operating model without understanding what is currently preventing the organisation from adapting.

The diagnosis is not a preliminary step to be rushed through before the real work begins, and it is not a neutral information-gathering exercise. It is the foundation the design is built on, and it is also the first expression of the framework's principles. How an organisation chooses to diagnose already reveals whether people are the centre or the periphery. A diagnosis done to people, extracting information while treating them as subjects, undermines the human-centric culture before design even starts. A diagnosis done with people, treating them as the holders of operational truth and as participants in understanding the organisation, begins building the agency and trust that a continuously adaptive model will depend on. The principles do not wait for the design phase. They begin in how the diagnosis is conducted.

This is where the stakes are highest, and where the standard case for diagnosis stops short. The danger is not simply that an organisation builds on a foundation it has not seen. It is that AI applied without that understanding cannot distinguish a load-bearing element from a removable cost. The relationships, the judgement, and the undocumented knowledge that hold an organisation together rarely declare themselves on an org chart or a process map. From a distance they look like overhead, which is to say they look like precisely the work AI was brought in to remove. Applied to a model that has never identified them, AI will automate them away. The diagnosis is how an organisation finds its keystone before it designs around it. And the keystone is not a stone you can remove to see what happens, because by the time the structure moves, it is already coming down.

The Living Arch framework helps you build for the future once you know what you have today.

You cannot redesign a foundation you have never seen.

To discuss how the Living Arch framework applies to your organisation, or to explore an engagement, get in touch.

Sarah Heffron

Sarah Heffron Advisory & Consulting

sarah@sheffronadvisory.com | sheffronadvisory.com

Living Arch™ is the proprietary operating model design framework of Sarah Heffron Advisory & Consulting.