WHITE PAPER

A human grounding layer for AI

'The Foundation of True Divergence'

Building deeper human foundations for chat systems and agentic workflows

Source-of-truth public edition | May 2026
Public website edition | Text only

simme.co.uk

Contents

Executive summary

AI is becoming part of daily work, planning, creativity, research, and automation. Yet many systems still rely on generic defaults, shallow personalisation, repeated prompting, and delayed behavioural alignment. The result is familiar: users spend too much time correcting tone, re-explaining preferences, steering behaviour, and rebuilding context across chats, workflows, and agents.

Simmes SymSoul is designed to address that gap. It is an offline soul-building application that helps users create structured, deeply personalised governance files for AI chats and agentic workflows through the ATQE — Adaptive Tacit Question Engine: a research-led question system designed to draw out the deeper human signal, preferences, behavioural patterns, working style, and tacit knowledge that ordinary onboarding usually misses.

Rather than relying on live prompting, weak memory accumulation, or machine-led attempts to gradually infer what matters, SymSoul gives the user a deliberate way to build a stronger human foundation for AI. The ideal time to add that foundation is before deployment, but SymSoul can also be introduced at any stage to strengthen alignment, improve behaviour, and reduce ongoing friction.

The system produces two main outputs. A Clean Soul is the lean, deployment-ready file designed for direct AI use. A Working Soul is the editable build-state version that preserves continuity, metadata, revision history, and development logic for later refinement. Together, these outputs are designed to improve alignment, reduce prompt friction, strengthen behavioural consistency, and make AI systems more useful in relation to the person behind them.

Simmes SymSoul is not the model, not the workflow engine, and not simply another prompt layer. It is better understood as a human governance layer for AI: a system that helps shape how intelligence is grounded, guided, and carried into live use. In practical terms, it sits between the human and the wider AI stack, giving both chat systems and agentic workflows a stronger starting state.

SymSoul also sits within a wider framework lineage. The product layer is supported by deeper thinking from the Euclid1 Paradigm, including Talents, Nodes, PFI-X, HyperSym, proportional structuring, and controlled divergence. This public white paper introduces those foundations at a product and principle level, while a separate technical foundation paper can explore the full Euclid1 architecture, Node history, HyperSym model, and project origins in greater depth.

This white paper is intended as a public source-of-truth document for SymSoul. It explains the problem the product addresses, the model it proposes, how the system works, what is already real today, what still belongs to future direction, and how to assess whether the product is a strong fit for a given workflow or organisation.

1. The problem now

Most AI systems still produce output that feels generic, overly standardised, or insufficiently human in relation to the person using them. Users often have to spend significant time trying to get the system to understand their preferences, communication style, priorities, and desired output quality. In chat-based AI, this creates re-prompting cycles, drift, and shallow alignment. In agentic workflows, the same problem becomes more serious because agents may begin acting before they have any meaningful understanding of their owner's behavioural direction.

In normal use, people often find themselves repeatedly explaining who they are and how they want the AI to behave, only to repeat the process again as context shifts or fades. They may rely on memory, customisation settings, or reverse prompting, but these methods usually capture only a small fraction of what is actually needed. The result is not only generic output, but wasted time, wasted tokens, weaker behavioural alignment, and reduced trust in the system over time.

The deeper problem is not that AI lacks capability. It is that capability alone does not create meaningful human grounding. Without a stronger foundation, AI systems risk becoming efficient at producing output while remaining inefficient at producing the right output for the right person in the right way. That weakness becomes even more important as AI moves from occasional use into regular workflow infrastructure and, increasingly, into agentic environments where weak early alignment can multiply into larger downstream inefficiencies.

SymSoul starts from the view that better AI use is not only about model quality. It is also about the quality of the human foundation the model receives. If that foundation is shallow, rushed, generic, or assembled through scattered prompts, the system may remain impressive yet insufficiently aligned. If that foundation is deliberate, structured, and deeply personal, more useful behaviour can emerge earlier and with less repeated correction.

2. Why SymSoul was built

SymSoul did not begin as a branding exercise or a thin wrapper around existing AI tools. It began as a response to a lived frustration: the feeling that powerful AI systems still too often defaulted to behaviour that was polished, generic, and disconnected from the complexity of real human thought.

That frustration became more serious over time. What began as curiosity about AI became a practical need for stronger support, clearer thinking, and more human-compatible behaviour from the systems increasingly present in everyday life. The longer the experimentation continued, the clearer the gap became. AI could be capable, even impressive, while still requiring too much correction, too much explanation, and too much effort to become meaningfully aligned.

The project also became personal in a deeper way. In periods where life became emotionally and cognitively heavy, AI was no longer just an interesting tool. It became part of the practical scaffolding needed to keep moving, keep thinking, and keep functioning. That made the weakness of generic AI behaviour impossible to ignore. When someone is already under pressure, the last thing they need is a system that demands further effort simply to understand the person it is meant to help.

SymSoul was built in response to that gap. Its purpose is not to romanticise human complexity, but to give it a clearer place inside the AI relationship. The project grew from the belief that if AI is going to become more useful, more trusted, and more embedded into daily work, then it needs stronger human grounding from the start rather than endless correction afterwards.

This is also why SymSoul has developed beyond a simple product concept. It has roots in a wider attempt to understand how human preference, personality, divergence, reasoning style, values, and tacit knowledge can be carried into AI systems in a structured way. That wider foundation includes the Euclid1 Paradigm and related framework work, but SymSoul is the practical product expression of that thinking: a way for people to build, review, save, and deploy a stronger human foundation for AI.

3. Why current approaches fall short

Most current AI customisation methods rely on repeated prompting, light custom instructions, slowly accumulated memory, or shallow onboarding questions. These tools can be useful, but they are not designed to perform a deep, structured, human-led soul-building process.
A prompt tells the AI what to do in the moment. Custom instructions provide a lightweight guidance layer. Memory mainly captures useful details over time. Reverse prompting still asks the machine to decide which questions matter. Each of these approaches can help, but none is designed to create a strong behavioural foundation before the system goes live.

This matters because these methods usually capture fragments rather than foundations. They can improve a conversation, but they rarely generate a complete, portable, inspectable, and reusable human governance layer. The result is that the user often ends up teaching the AI through fragmented interactions, correction cycles, and long periods of use, with no guarantee that the most important parts of their personality, preferences, working style, behavioural expectations, or deeper divergence will ever be captured properly.

Generic onboarding also tends to miss tacit knowledge. A user may know how they think, decide, judge quality, respond to pressure, handle ambiguity, or communicate nuance, but they may not know how to express those things unless the right questions draw them out. This is one of the reasons SymSoul uses the ATQE rather than a static questionnaire. The goal is not simply to collect answers. The goal is to surface the kind of human signal that ordinary prompting, memory, and form-based onboarding often fail to reach.

Simmes SymSoul is different because it creates a structured human governance layer first. Instead of depending on scattered prompts, vague memory patterns, or generic machine-led questioning, it uses a dedicated soul-building environment to translate the human user into a more complete, deliberate, and portable form. In agentic environments this distinction becomes especially important, because generic onboarding often gives the agent work before it gives the agent a meaningful understanding of its owner.

4. The new organising model: human-centred AI grounding

Simmes SymSoul should not be understood simply as a personality tool, a prompting shortcut, or a narrow customisation layer. It is better understood as a human grounding layer for AI: a system designed to build the behavioural, identity, and preference foundation that many current AI environments still lack.

At the heart of SymSoul is a simple belief: machines, like people, need a meaningful foundation if they are to develop in useful and human-compatible ways. Out-of-the-box AI may be powerful, but it does not arrive with lived context, inherited values, or the equivalent of human developmental grounding. SymSoul is designed to provide a structured form of human grounding so that more meaningful divergence can appear earlier, more clearly, and more efficiently.

This matters because human divergence is not noise. It is one of the main sources of creativity, invention, interpretation, judgement, tone, and progress. SymSoul is designed to preserve more of that difference inside AI-guided systems, rather than allowing increasingly generic defaults to flatten behaviour across users, contexts, and workflows.

The category claim here is therefore operational, not decorative. SymSoul sits between the human and the wider AI stack. It does not replace the model, the workflow engine, or the orchestration layer. It improves the starting state of those systems by giving them a stronger human foundation to work from.

5. What Simmes SymSoul is

Simmes SymSoul is an offline soul-building application that helps users create structured, deeply personalised governance files for AI chats and agentic workflows through the ATQE, the Adaptive Tacit Question Engine. The ATQE is the guided question system at the centre of SymSoul. Its role is not simply to ask more questions, but to help surface the tacit human knowledge, preferences, behavioural signals, working style, and deeper personal context that an AI system would otherwise have to infer slowly and imperfectly over time. Its purpose is to improve AI alignment, reduce prompt friction, and support more efficient behaviour across both live chat use and more complex multi-agent environments.

The application is available in a free version and a paid eXpanded version. The free version allows users to understand the onboarding logic and the core value of structured soul building. The paid version extends that into a fuller local product experience and deeper build environment. From the Simmes website, the user can access the free version, explore the wider thinking around the product, and move into the paid version when ready.

At the centre of the product is the Soul Builder, a guided human-centred environment in which the user answers a carefully designed sequence of adaptive questions. As the user responds, the system progressively constructs a structured soul file that can be reviewed live, refined over time, and exported in multiple forms depending on the intended AI deployment.

The Builder Chat is the interface through which this takes place, but it is not the whole product. The Builder Chat is the guided adaptive question interface inside the wider Soul Builder environment. The product also includes progress logic, quality signalling, continuity mechanisms, security and verification features, talent shaping, and export logic. In other words, SymSoul is not simply a clever questionnaire. It is a full soul-building system.

6. How the system works

The process starts when the user accesses Simmes SymSoul through the Simmes website. They can begin with the free version of the app in the browser or move to the paid eXpanded Soul version for deeper local use and a fuller build environment.

The user is first guided through the main dashboard, where they can review their current progress and move directly into the Soul Builder Chat. This gives them a clear entry point into the build process rather than leaving them to invent the structure for themselves. The dashboard also makes progression visible, surfacing what is complete, what is locked, what comes next, and where more depth can be added later. For users who want a simpler and more focused build experience, SymSoul also supports a Hyperfocus route. Hyperfocus is not a separate mode and does not change the underlying question engine. It presents the existing Soul Builder Chat in a quieter, distraction-reduced page so the user can concentrate only on answering the questions and saving their working soul.

As the user answers questions, they are guided through the ATQE, a carefully designed adaptive flow that helps shape the soul in a structured way. Each question is designed to draw out high-value human signal rather than simply collect more text. The always-visible answer-value signal helps the user understand the strength and usefulness of their input, encouraging answers with clearer specificity, behavioural relevance, contrast, context, and practical value for the resulting soul. The user can edit earlier answers at any point, skip questions to return later, and use X-Ray style visibility to understand where an answer feeds into the wider soul structure.

The Soul Builder then uses the ATQE and its internal scoring logic to progressively construct the soul file. This developing structure can be viewed live in the YAML review area, giving the user transparency into how their answers are being translated into machine-readable form rather than hiding the result until the end.

The build is staged rather than dumped into one long, exhausting session. The user moves through core shaping, deeper refinement, optional reflective layers, and agent-oriented guidance as appropriate. This staged approach is part of the product's philosophy: building a high-value soul should feel guided and insightful, not rushed or cognitively punishing.

7. Core outputs, Talents, and shaping

SymSoul produces two main outputs: the Clean Soul and the Working Soul.
A Clean Soul is the refined, deployment-ready file designed for direct use inside an AI system. It contains the structured personality breakdown, behavioural guidance, key user-alignment signals, and, where selected, a Talent layer containing relevant Nodes from the latest version of the Euclid1 Paradigm. It is designed to improve alignment, save time, reduce unnecessary token usage, and carry a consistent human-centred foundation across AI chats and agentic workflows.

A Working Soul is the full editable build file. It contains not only the core soul output, but also metadata, change logs, build memory, rollback points, and supporting information needed to preserve the development state of the soul. It is used for saving, backing up, refining, and continuing the build locally over time. A Working Soul makes the process continuous rather than disposable. It allows the user to reopen a build, revisit earlier answers, add or change a Talent, adjust their current thinking, or fork one build into several different soul outcomes for different projects and workflows.

SymSoul also includes a shaping layer built around Talents and Nodes. A Talent is a curated combination of interdependent Nodes drawn from the Euclid1 Paradigm. When added to a soul, a Talent gives that soul stronger grounding within a specific domain, helping it develop more effective skills in that area. This is especially useful in multi-agent workflows where different agents require different behavioural strengths. This distinction between Talent and Skill matters. In the SymSoul model, Talent is not just another feature toggle. Talent provides the stronger natural grounding from which a more specific functional skill can later be developed. Skills can be learned or brought in from trusted sources. Talent is the deeper orientation that makes those skills cohere more effectively.

Nodes are the building blocks that sit underneath this shaping logic. They are designed to work both independently and in combination, allowing them to be reused across multiple Talents while preserving auditability, structural clarity, and efficient file design. This makes SymSoul's shaping model both more modular and easier to refine over time.

8. The Euclid1 foundation

SymSoul is the public product layer, but it is not an isolated idea. It is built from a deeper body of work known as the Euclid1 Paradigm: a structured cognitive and behavioural framework designed to improve how AI systems organise behaviour, reasoning, output quality, proportional judgement, and human alignment.

Within SymSoul, Euclid1 is most visible through the Talent and Node system. Talents are curated shaping layers that can be added to a soul when a user wants stronger grounding in a particular domain. Nodes are the underlying behavioural and structural components that make those Talents coherent. They are not random feature blocks. They represent reusable pieces of reasoning, safety, calibration, creativity, ethical control, persona shaping, output optimisation, contradiction handling, and structural balance.

This matters because SymSoul is not only collecting user preferences. It is translating human input into a more structured foundation that can influence how an AI system behaves, prioritises, writes, reasons, and adapts. The Euclid1 Paradigm provides part of the deeper architecture behind that translation.

One of the deeper ideas inside the Euclid1 foundation is HyperSym, short for Hyperbolic Symmetry. HyperSym was developed as a way to think about richer, less generic AI reasoning without allowing that reasoning to become uncontrolled, chaotic, or detached from useful output. In simple terms, it combines expansion and contraction: the ability to move outward into wider associations, alternative framings, hidden patterns, creative possibilities, and less obvious connections, then bring that wider exploration back into proportion, structure, usefulness, and practical form.

This controlled movement matters because many AI systems either stay too close to generic patterns or expand in ways that become unfocused. HyperSym is part of the attempt to create a more balanced reasoning architecture: one that allows difference, creativity, and divergence to appear, while still remaining governed by ethics, safety, calibration, contradiction handling, and output discipline.

The public version of this white paper does not expose the full internal framework. That is intentional. Its purpose is to explain the product and its principles clearly, not to disclose every implementation detail. A separate technical foundation paper can explore the full Euclid1 Paradigm, Node architecture, HyperSym model, PFI-X history, ATQE development, and wider project origins in greater depth.

9. Export, continuity, and trust

The build process does not stop at generating structure. SymSoul is designed to support review, continuity, and trust across the whole lifecycle of a soul.
When the build is complete, the user moves to review and export. There they can inspect readiness, check completeness, review the quick-fix queue for weaker or missing areas, inspect the YAML, and run a safety scan designed to detect issues such as prompt-injection patterns before export. They can then export either a Clean Soul for AI deployment or a Working Soul for later continuation. Advanced export options support more specialised handoff formats for agentic workflows and existing frameworks.

SymSoul also includes a Soul Guardian sidecar verification file. This uses a fingerprint check to help confirm the authenticity and integrity of an exported Clean Soul or Working Soul. It allows users to store verification data separately from the main file so they can later confirm that the file remains original, untampered, and safe to re-import or deploy.

Where local storage is used to support continuity inside the application, it is treated as a convenience layer rather than the permanent source of truth. The exported Clean Soul, Working Soul, and associated verification data remain the portable assets the user can keep, inspect, back up, re-import, or deploy elsewhere.

This trust model is important. The product does not assume that quality alone is enough. It treats continuity, authenticity, verifiability, and safer deployment as part of the user's real-world needs. In practice, this means the user is not only building a soul. They are building something they may want to reuse, share, modify, verify, and trust across time.

10. Real-world use cases and deployment scenarios

SymSoul creates its clearest immediate value wherever AI is already being used regularly, but still lacks a strong enough human foundation. The more often a user has to correct tone, re-explain preferences, steer behaviour, or rebuild alignment across chats and agents, the more valuable SymSoul becomes.

For founders, operators, and small business users, SymSoul gives AI a more complete understanding of how they think, communicate, prioritise, and want work carried out. This reduces repeated correction and helps create a more stable working style across repeated use. For marketing, content, and brand work, it helps preserve voice, tone, identity, and human texture inside AI-assisted output rather than allowing the system to default towards generic patterns.

For local and self-hosted AI users, SymSoul is a natural fit because it aligns with the same values that drive local AI adoption: privacy, control, portability, inspectable files, and reduced cloud dependence. For higher-trust environments such as advisory, compliance, or specialist professional work, the value lies in having a stronger behavioural base before important outputs are produced.

For multi-agent systems, the use case is even clearer. One weak human foundation can multiply into many downstream inefficiencies. A lead agent may misunderstand priorities, a specialist sub-agent may use the wrong tone, or multiple agents may drift in different directions. SymSoul helps by giving the workflow a stronger behavioural base before those agents begin acting.

SymSoul also supports multi-soul logic. A user can build one strong foundation and then create multiple Working Souls or Clean Souls for different contexts, talents, workflows, or agent roles without starting from scratch each time. This is especially valuable for advanced users working across several systems or projects.

11. Trust, proof, and what is already real

SymSoul is credible because it is already a functioning application with a full user journey, live build environment, export system, and clear deployment logic for immediate AI use. It is not only a theory or a future roadmap.

The clearest visible proof points include the guided Soul Builder dashboard, the adaptive Builder Chat, the ATQE answer-value signal, the live YAML review panel, the Talents and Nodes layer, the export environment, the ability to produce both Clean and Working Souls, the Soul Guardian verification system, Hyperfocus, Save My Soul continuity, and safer export options designed to support more trustworthy deployment.

These features matter because they make the product transparent. The user can see the structure being built, understand the role of their input, inspect the resulting output, review refinement opportunities, and verify the state of the exported file. This is not personalisation hidden behind vague promises. It is visible, inspectable, and operational.

The current product already supports the full core build process, review flow, export flow, Working Soul handling, and broader Version 1 user journey. In practical terms, SymSoul is already a real product with a real workflow, not only an emerging idea.

12. Commercial fit and market role

Commercially, SymSoul is best understood as a human governance layer for AI rather than a direct competitor to the large models, chatbot platforms, or workflow engines themselves. It is not trying to replace the model, replace the agent framework, or replace the orchestration layer. Its role is to improve the starting state, and therefore the future state, of those systems by giving them a more structured, more personal, and more deliberate human foundation.

Its market role is therefore not better chatbot, but better grounding and growth layer. It acts as the bridge between the human and the wider AI stack. Where many current systems ask the user to shape behaviour through live prompting, scattered correction, or slowly accumulated memory, SymSoul is positioned as a pre-deployment system that creates a stronger behavioural base first.

That gives SymSoul a specific strategic advantage. It can sit across multiple models, multiple chat environments, and multiple agentic workflows without needing to become the model provider itself. In the near term, the commercial path is a freemium-to-paid journey built around the free and eXpanded versions. Over time, that can develop into a broader ecosystem of higher-value outputs, talent-driven refinement, advanced exports, and potentially team or enterprise-facing versions.

The strongest near-term users are not everyone interested in AI. They are serious AI users who already feel the cost of generic output, repeated steering, weak alignment, and behavioural drift. That includes agent builders, workflow operators, local AI users, creator-operators, and founder-led or service-led businesses that already use AI heavily enough for weak grounding to become expensive.

13. What exists now and what comes next

What already exists today is the SymSoul application in both its free version and its paid eXpanded version. The current product already allows the user to build, review, edit, and export advanced soul files for immediate use with AI. The full core build process, review flow, export flow, Working Soul handling, and broader Version 1 user journey are already in place.

What is still evolving is the ongoing auditing, refinement, and expansion of the app as AI systems, user needs, and deployment environments continue to change. The near-term direction is continued refinement through real customer use, structured audits, and iteration. The longer-term vision is broader: accessible tools that help steer consumer and operator AI in better directions, including more advanced human-guidance systems for future autonomous and agentic environments.

The separate technical foundation paper will support that longer-term work by recording more of the Euclid1 Paradigm, Node architecture, HyperSym development, PFI-X origins, and ATQE design history in one deeper reference document. This allows the public paper to remain clear and commercially readable while the technical paper carries the fuller framework depth.

That distinction matters because it keeps the paper honest. It separates present reality from future ambition and lets the reader see clearly what users can benefit from now, what is still being refined, and what belongs to the wider long-term direction of the Simmes project.

14. Philosophy and ethical position

At the heart of SymSoul is the belief that machines, like people, need a meaningful foundation if they are to develop in useful and human-compatible ways. Out-of-the-box AI does not arrive with lived context, inherited values, or the equivalent of human developmental grounding. It may be highly capable, but it does not begin with the unique personal foundations that shape how a real human thinks, feels, prioritises, and understands the world.

One of the core problems with current AI is that it takes too long to understand the user's preferences, personality, priorities, and intended way of working. Even with a brain-dump approach or a reverse-prompt cycle at the start, the process is often shallow, rushed, and incomplete. If the goal is to create an AI with a truly distinctive and useful behavioural outlook from the beginning, there needs to be a more deliberate and carefully developed foundation. That is the role SymSoul is designed to play.

Simmes takes the view that if AI is going to become safer, more capable, more inventive, and more useful, it must begin from a stronger human grounding. The more powerful these systems become, the more important it is that they are shaped by meaningful human direction rather than left to generic defaults and shallow personalisation.

In ethical terms, SymSoul aims to make AI more human-centred, more grounded, and safer to use by encouraging deeper personal alignment, clearer behavioural direction, and a more deliberate relationship between the user and the machine. The project does not start from the idea that human difference should be normalised away. It starts from the opposite belief: that human divergence is part of what makes better guidance, better judgement, and richer invention possible.

15. How to assess fit

SymSoul is likely to be a strong fit if one or more of the following are true.

1. AI is already part of your regular workflow rather than something you only use occasionally.
2. You or your team spend a noticeable amount of time correcting generic output, re-explaining preferences, or steering the system back towards the desired style or behaviour.
3. You care about portability, inspectability, and a cleaner separation between the human foundation and the live model environment.
4. You are building or managing multi-agent systems where weak early alignment could multiply into larger downstream inefficiencies.
5. You want stronger continuity between one project, one chat, one workflow, and the next.

SymSoul is probably less urgent if AI is only used rarely, if highly generic output is already acceptable for the work in question, or if the user does not yet feel any practical cost from weak alignment. In those cases the product may still be interesting, but its full value may not yet be felt.

A simple evaluation lens is this: does your current AI stack already have a reliable way to build, preserve, verify, and redeploy a structured human foundation before live use? If the answer is no, then SymSoul is addressing a real gap.

16. Why this matters now

SymSoul matters now because AI is moving from being an occasional tool into a daily working layer for millions of people, yet many systems are still shaped by generic defaults, shallow personalisation, repeated prompting, and delayed behavioural alignment. As AI becomes more agentic, more autonomous, and more embedded into real workflows, the cost of weak human grounding becomes higher.

The significance of SymSoul lies in the fact that it addresses a problem that is likely to grow rather than shrink: the gap between machine capability and meaningful human grounding. Models are becoming more powerful, faster, and more accessible, but that does not automatically make them more personally aligned, more behaviourally stable, or more deeply attuned to the human being using them.

There is also an efficiency dimension. The better grounded an AI system is from the start, the less wasted interaction is needed to reach useful alignment. That means less prompt churn, less corrective steering, less friction, and potentially less avoidable compute waste across repeated use. Although this issue may lessen over time as local model deployment becomes more practical and supporting hardware improves, today's frontier models still consume meaningful energy at scale. If approaches like this were implemented widely, they could contribute to reducing avoidable AI waste rather than merely shifting it downstream.

17. Closing

Simmes SymSoul exists because the current direction of AI, while extraordinarily capable, is still too often built on weak human grounding, shallow personalisation, repeated prompting, and generic behavioural defaults. As AI becomes more powerful, more autonomous, and more deeply integrated into everyday work, the cost of that weakness increases.

SymSoul is designed as a direct answer to that problem. It gives the user a structured way to build a deeper human foundation for AI before deployment, rather than relying on live correction, vague memory accumulation, or machine-led attempts at understanding. In doing so, it creates a bridge between human individuality and machine capability, allowing AI systems to begin from a more grounded, more deliberate, and more useful starting point.

If that direction succeeds, the long-term impact of SymSoul may reach well beyond one application. It may help establish the principle that AI should not only become more intelligent over time, but more meaningfully grounded in the people it is meant to support. That is the real significance of the project, and the reason Simmes SymSoul matters both as a product in the present and as an idea for the future.

Appendix A. Terminology reference

A1. Core product terms

Simmes

The wider human-led brand and framework ecosystem within which SymSoul sits. Simmes is the broader context for the product philosophy, framework development, and future tooling around stronger human grounding for AI.

Simmes SymSoul

The offline soul-building application designed to create structured, deeply personalised governance files for AI chats and agentic workflows. SymSoul is the product itself: the environment in which the user builds, reviews, refines, and exports a structured human foundation for AI.

Soul

A structured human governance file created through SymSoul. In practice, a soul contains the user's behavioural, personality, preference, and guidance architecture in a form that can be carried into AI systems or retained for later refinement. This is the core object the whole product is designed to build.

Soul Builder

The guided human-centred build environment inside SymSoul. It is the wider environment within which the user develops the soul, rather than the chat interface alone. The Soul Builder includes the question flow, progress logic, structure visibility, continuity mechanisms, and export logic.

Builder Chat

The adaptive question interface through which the user develops the soul. It is the conversational layer inside the Soul Builder that guides the user through a carefully designed question sequence, but it is not the whole product.

ATQE / Adaptive Tacit Question Engine

The central question system inside SymSoul. The ATQE guides the user through adaptive, structured questioning designed to surface tacit knowledge, behavioural preferences, communication style, working patterns, values, context, and deeper human signal. Its purpose is not to collect generic answers, but to help build a more useful, personalised, and deployment-ready human foundation for AI.

Dashboard
The main operating view of the Soul Builder. The dashboard gives the user a clear starting point into the build, surfaces progress, shows what is complete or still locked, and makes the overall build process legible rather than leaving the user to invent the structure for themselves.

Hyperfocus
A focused Builder Chat route designed to remove unnecessary visual and navigational distraction during the soul-building process. Hyperfocus is not a separate question mode. It uses the same Builder Chat and underlying ATQE, but presents them in a simpler environment so the user can concentrate on answering questions and saving the working soul.

ATQE answer-value signal
The always-visible quality indicator attached to each question during the ATQE build process. It uses internal rubrics, schema, and scoring logic to help show the strength and usefulness of the user's input. Its purpose is to encourage clearer, more useful answers without turning the process into opaque scoring.

X-Ray
The visibility layer that helps the user understand where an answer feeds into the wider soul structure. Rather than hiding the consequences of an answer, X-Ray makes the mapping between the guided input and the resulting build structure more transparent.

Quick Core
The early-stage build layer used to establish the strongest essential foundation of the soul quickly. It gives the user a guided way to build the most important grounding before moving into deeper or more optional shaping layers.

Deep Refinement
The later build layer used to add missing detail and strengthen underdeveloped areas once the core foundation is in place. It extends and sharpens the soul beyond the first-pass structure.

Dream Mode
An optional reflective layer used to add more depth, especially around reflective, imaginative, or less immediately task-bound aspects of the soul. It sits later in the process because it is additive, not foundational.

Agent Mode add-on
The build-stage extension that hardens execution and integrity for more agentic workflows once the earlier refinement layers are complete. It is effectively the applied workflow-facing part of the later-stage soul shaping sequence.

Free version
The browser-based entry version of SymSoul, designed to let users understand the onboarding logic and the core value of structured soul-building before moving into the fuller build environment.

eXpanded version / eXpanded Soul
The paid, fuller local product experience that extends the build depth, continuity, and overall working environment beyond the entry version. It is intended for users who want the deeper SymSoul workflow rather than only the initial browser experience.

A2. Shaping and framework terms

Clean Soul
The refined, deployment-ready output designed for direct AI use. A Clean Soul is leaner than a Working Soul and is intended to provide immediate personality clarity, behavioural grounding, and structured guidance without carrying the full build-state metadata with it.

Working Soul
The editable build-state output designed for refinement, backup, re-import, and multi-soul development. A Working Soul preserves continuity, metadata, revision history, and development logic so the user can reopen, revise, fork, or continue the build later.

Multi-soul logic
The ability to build one strong foundation and then create multiple soul variants for different contexts, roles, projects, or workflows without starting from scratch each time. In practice, this often means using one core human foundation and adapting it through different Talents, agent settings, or refinement choices.

Talent
A curated combination of interdependent Nodes that gives the soul stronger natural grounding within a specific domain. A Talent does not act like a short-lived add-on. It shapes the soul more deeply, especially in workflows where different roles or agents require different behavioural strengths.

Skill
A more specific capability developed within a Talent's domain. In the SymSoul model, Skill is downstream of Talent: Talent gives the soul stronger orientation in a domain, and Skill is then built more effectively within that grounded area.

Node
An individual behavioural or structural component within the Euclid1 Paradigm. Nodes are designed to work both independently and in combination, allowing them to be reused across multiple Talents while preserving auditability, structural clarity, and efficient file design.

Node depth
The selectable level of Node intensity or context weight applied during shaping. In practical terms, Node depth controls how deep and token-heavy the shaping system becomes, allowing the user to choose a lighter or more expansive structural layer depending on their intended use.

Euclid1 Paradigm
The structured framework of Nodes used to shape behaviour, reasoning, and output orientation. It is built around principles of balance, proportion, controlled divergence, ethical control, calibration, persona shaping, and useful reasoning discipline.

PFI-X
Persona Fusion Interface-eXtended, the foundational personality architecture behind the Soul Builder and ATQE. It began as a top-level persona layer inside the wider framework and evolved into the core personality-building logic that now sits underneath the guided soul-building process.

HyperSym
Hyperbolic Symmetry, a foundational concept within the framework designed to support wider, more divergent conceptual expansion while remaining under structured control. It is part of the attempt to introduce richer and less generic modes of reasoning into AI behaviour.

Chain of Concept
The framework's controlled pathway for conceptual expansion before later contraction into useful structure. It allows the system to move beyond the immediate surface of a prompt and explore wider meaning, deeper connections, alternative framings, and less obvious patterns before settling into a final answer.

Chain of Focus
The complementary contraction pathway that brings expanded conceptual material back into clearer structure, proportion, priority, and usable output. It helps prevent wider reasoning from becoming detached from the user's practical need.

Controlled divergence
The principle that human difference, creative expansion, and non-generic reasoning should be allowed to appear, but under enough ethical, structural, and practical control to remain useful. This is one of the reasons HyperSym matters inside the wider Euclid1 foundation.

Agent Mode
A later-stage mode for shaping boundaries, preferences, and behavioural guidance for agentic workflows. It appears after the core soul foundation has been established so that agent-level guidance is built on top of a stronger and more valuable personal base.

Human grounding layer
The organising model proposed by SymSoul. It describes the product's role as the structured human foundation that sits between the user and the wider AI stack, improving the starting state of chat systems and agentic workflows without replacing the models or orchestration tools themselves.

Human governance layer
A closely related framing to human grounding layer. It emphasises that the soul is not merely descriptive or cosmetic, but provides structured behavioural guidance that helps shape how the AI interprets, prioritises, and responds in use.

Behavioural alignment
The degree to which an AI system's outputs, tone, priorities, and decision patterns fit the actual user's needs, preferences, and intended way of working. In SymSoul, stronger behavioural alignment is one of the primary outcomes of a better-built human foundation.

Divergence / human divergence
The richer variety of real human thought, preference, creativity, interpretation, and judgement that SymSoul is designed to preserve inside AI-guided systems. In the product philosophy, divergence is treated as a strength to be carried into AI rather than normalised away by generic defaults.

A3. Workflow, export, and trust terms

Export
The stage in which the user reviews readiness, inspects structure, chooses the soul type, and packages the result for use outside the Builder. In SymSoul, export is not a trivial file-save action; it is part of the product's continuity and trust model.

Quick-fix queue
The guided review layer used before export to surface weaker, missing, or incomplete areas of the build. Its purpose is to help the user strengthen the soul before deployment rather than exporting a structure that is technically complete but still underdeveloped.

YAML review area
The live inspection area in which the user can see the developing soul structure in machine-readable form. This transparency is important to SymSoul's philosophy because it allows the user to inspect what is being built rather than waiting for a hidden result.

YAML safety scan
The export-stage scan designed to check the soul content for issues such as prompt-injection patterns before export. It functions as part of the product's trust and safer deployment model.

Advanced export dashboard
The deeper export environment for users who need more than the basic handoff flow. It supports broader export context and more specialised packaging options for users working with fixed stacks, frameworks, or more advanced workflows.

Soul Guardian
The verification sidecar used to help confirm file authenticity and integrity. Soul Guardian allows the user to keep verification data separate from the main soul file so they can later confirm whether a Clean Soul or Working Soul still matches its original exported state.

Fingerprint
The SHA-256 hash value generated for an exported soul file. It acts as the integrity signature used by Soul Guardian to confirm that the structure of the file has not been altered.

Threat Aware
The export-layer option designed to support safer deployment thinking by helping the user account for covered threat conditions at the point of handoff. Rather than treating trust as an afterthought, Threat Aware makes deployment risk part of the export decision itself.

Save My Soul
The continuity and save feature used inside the Builder Chat and Hyperfocus flow. Save My Soul allows the user to pause without losing progress by packaging the current Working Soul state for later continuation. In practice, this supports the product's view that soul-building should be guided, thoughtful, and recoverable rather than forced into one exhausting session.

Soul Timeline
The continuity layer that allows the user to return to earlier build states as the soul evolves. It supports safer editing, rollback confidence, and a sense that the build process is continuous rather than fragile or disposable.

Prompt friction
The repeated time, effort, and correction burden created when the user must continually retry prompts, restate preferences, and steer behaviour because the AI lacks a strong enough foundation. Prompt friction is one of the main operational costs SymSoul is designed to reduce.

Soul Signal
The useful human signal captured through the ATQE. A strong Soul Signal is not simply a long answer; it is an answer that gives the soul clearer behavioural direction, preference structure, contrast, context, and practical usefulness inside AI deployment.

Question Intelligence
The internal design approach behind the ATQE. Question Intelligence focuses on asking fewer, sharper, more useful questions that extract meaningful human grounding, rather than relying on long generic questionnaires or shallow onboarding prompts.

Technical foundation paper
The deeper companion document intended to explore the Euclid1 Paradigm, Node architecture, HyperSym, PFI-X, ATQE origins, and longer project history in greater detail than is appropriate for the public commercial paper.

Appendix B. Public Euclid1 Node reference

Purpose of this public Node reference

This appendix gives a controlled public overview of key Euclid1 Nodes that help explain the framework behind SymSoul. It is not intended to disclose the full internal Node architecture, dependency rules, scoring methods, prompt structures, or implementation details.

Ethics Node

Provides the highest-level ethical orientation for the framework. It helps ensure that outputs remain human-centred, responsible, and bounded by appropriate moral constraints.

Output Optimisation Node

Supports clearer, more useful, and better-structured outputs. It helps prevent weak formatting, unnecessary verbosity, poor prioritisation, generic response patterns, and avoidable friction in how information is presented.

Safety and Security Node

Provides safety-aware behavioural control, including caution around misuse, manipulation, unsafe instructions, hostile prompt patterns, and deployment risks that could undermine trust.

Calibration Node

Helps maintain behavioural consistency over time. It supports steadier tone, better expectation matching, improved context awareness, and reduced drift between the intended soul behaviour and actual output.

Golden Ratio Structuring Node

Uses proportional thinking as a structuring principle. Its purpose is not decorative mathematics, but better balance between emphasis, detail, clarity, depth, hierarchy, and usability.

MemoryKey and continuity Node

Supports the idea that important learning moments, exceptional responses, and meaningful session insights can be captured in a structured way for later improvement. In the wider framework, this contributes to continuity rather than treating each interaction as disposable.

HyperSym Node
Supports controlled expansion and contraction in reasoning. It allows wider conceptual exploration while ensuring that the final result returns to useful, structured, and human-relevant output.

Cognitive Control and Heuristics Node
Guides the system away from shallow shortcuts, brittle assumptions, or uncontrolled reasoning habits. It supports better judgement about when to expand, when to narrow, and when to question the path being taken.

Logical Integrity and Contradiction Handling Node
Helps detect tension, inconsistency, or conflict between instructions, values, assumptions, or outputs. It supports clearer reasoning and safer resolution of competing signals.

Failsafe Node
Provides a final protective layer when normal reasoning paths become unreliable, unsafe, contradictory, or outside the intended behavioural boundary.

Ethical Boundaries Node
Reinforces the framework's boundary system by helping distinguish useful assistance from behaviour that would be unsafe, manipulative, invasive, or misaligned with the intended human-centred purpose.

Integration Layer Node
Supports the way different Nodes, Talents, persona elements, and workflow requirements are brought together without losing coherence. It helps the system act as a structured whole rather than a loose collection of features.

Persona / PFI-X Node
Supports the construction and preservation of a coherent behavioural identity. In SymSoul, this contributes to the deeper personality, tone, preference, and behavioural architecture behind the soul.

Why the full Node system is handled separately

The public white paper explains Nodes at a product and principle level. The comprehensive technical foundation paper can go further, covering the full Node history, origins, dependencies, deeper mechanics, HyperSym development, and how the Euclid1 Paradigm evolved into SymSoul.