

PROTOTYPE IN TESTING AND PRE-SALE

Human Connection, Reimagined

The ORION White Paper on Personalized AI
Companionship

REVIEW DRAFT · DATED 2026-08-24

Executive Perspective

Companionship often lives in small moments: a familiar greeting after work, a patient listener, a remembered joke, or a quiet evening ritual. A product that enters those moments must be judged by its restraint, clarity, and respect.

ORION is our proposal for an adult consumer companion: a personalized conversational experience intended to connect software, voice, configurable character, and physical presence. The product is a prototype in testing and pre-sale. That sentence describes the programme as a whole; it does not mean that any individual capability described in this paper is implemented, approved, generally available, or ready for delivery. Every capability in the current public register is **Planned** and pending product-owner verification. Where this paper shows a possible interaction, it describes an intended experience, not a promise about the current prototype.

The starting point is emotional companionship, not a claim to solve a health condition. ORION is not a medical service, not a person, not a therapist, not a healthcare provider, and not an emergency service. It does not treat loneliness and does not provide a proven mental-health outcome. It is not a replacement for human relationships. Its generated language and expressions may feel socially legible, but the system does not possess emotion or consciousness. Candid boundaries like these are part of the product, not small print around it.

This paper explains the design perspective, a possible adult journey, the conceptual system, trust objectives, ownership questions, and release gates. It makes no commitment on price, delivery date, warranty, certification, geography, or general availability. A private pre-sale demonstration is an invitation to inspect a prototype programme, not proof that a finished product exists. Intimacy raises the standard of care: ORION aims to ask for less data, explain choices clearly, and make a graceful “no” as important as an engaging “yes.”

Why Companionship Needs Rethinking

The human category, without a product claim

Loneliness is not one thing and cannot be read from one kind of life. A person may live alone and feel connected, share a crowded home and feel unseen, want more contact, or simply want a calm private ritual at the end of the day.

WHO defines loneliness as the painful feeling arising from a gap between desired and actual social connections. (News release, definition paragraph)¹

ORION's perspective is that this definition offers useful category vocabulary because it centres the person's own desired level of connection. It does not establish that a device can close the gap, and it does not support an efficacy claim for ORION.

The distinction matters. Public discussion often moves too quickly from "social connection matters" to "a particular technology improves health."

Many studies do not distinguish technology-mediated from traditional social connection when assessing equivalence or long-term health and mortality risk. (Strengths and Limitations of the Evidence, p. 68)²

Our inference is that population evidence about social connection cannot simply be transferred to an AI companion. ORION begins with a consumer experience proposition, not a clinical conclusion.

Adults may want continuity, imaginative play, conversation, or a private ritual without regarding solitude as a problem. None of those motives is a diagnosis. A responsible companion leaves open the choice to call a friend, join a community, speak with a clinician, seek urgent help, or simply be alone. It makes no loyalty demands and does not encourage withdrawal from offline life.

What current studies can and cannot tell us

The evidence reviewed here is limited and heterogeneous.

In a 25-minute Wizard-of-Oz experiment, perceived partner identity as chatbot versus person did not alter reported disclosure outcomes. (Results, Overview, Table 2)³

Our privacy inference is narrow: adults may disclose personal material in an interaction they understand as machine-mediated. It does not establish how a real chatbot performs, whether an effect persists, or whether ORION produces any particular outcome.

One 2026 study made an experimental comparison.

After two weeks, students assigned to daily human-peer texting reported lower post-study loneliness than those assigned to the supportive chatbot. (Results, section 2.2.1, Table 1)⁴

Our inference is that this comparison is not a universal ranking of relationships or products. It does reinforce an essential boundary: a companion product cannot claim equivalence to human connection.

Two 2026 observational studies point in different directions.

In a cross-sectional survey of 14,721 Japanese adults, AI-companion use was associated with higher scores across three subjective well-being domains. (Abstract)⁵

Among 1,131 Character.AI users, reporting companionship as the primary use was associated with lower self-reported well-being. (Abstract)⁶

Our inference is that cross-sectional associations do not show what caused what. These populations, measures, platforms, and patterns of use are not interchangeable. Neither study is evidence about ORION.

Taken together, the literature gives us questions rather than a sales claim. Can a user change or end the relationship without pressure? What happens when generated language is wrong or emotionally intense? Which information is remembered, and can the user remove it? Product-specific evidence must answer those questions before public claims expand.

A more modest ambition

ORION explores a consumer interface for personalized companionship while keeping human agency in the foreground. An adult may shape tone, character, memory permissions, and boundaries. Personalization does not mean diagnosis, dependence, or imitation presented as inner life. A demo can examine latency, coherence, comfort, controls, boundaries, and comprehension. It cannot prove durable well-being from a selected interaction.

Organizations should not extrapolate GenAI performance or capabilities from narrow, non-systematic, anecdotal assessments. (MS-2.5-001, p. 34)⁷

ORION's resulting practice is to treat a polished demonstration as a demonstration, not as universal evidence.

"A polished demonstration is a demonstration, not universal evidence."

The Principles Behind ORION

Continuity without capture

A useful companion need not reset after every exchange, but continuity becomes intrusive when a system collects information simply because it can. ORION's selective-memory perspective is that context must be purposeful, understandable, and user-authorized. An adult may prefer no long-term memory, ask before saving, or review a visible list. These remain proposed patterns, not implemented features.

Data minimization offers a useful external principle.

Personal data must be adequate, relevant, and limited to what is necessary for the processing purpose. (Article 5(1)(c))⁸

Companies should reasonably limit their collection and retention of consumer data. (Executive Summary, Data Minimization, p. v)⁹

ORION uses these as separate design inputs. We do not claim that ORION has completed a GDPR assessment, received FTC approval, or satisfied any certification.

Choice that remains reversible

Character selection is an opening choice, not a trap. Adults can change preferences. A generated character must not guilt a user for editing settings, deleting memory, taking a break, or ending an account. It must not present paid continuation as proof of care. Choice also means choosing less: remote interaction, memory, animation, or frequent contact cannot be assumed to add value.

Control close to the moment

Controls belong where decisions occur. A clear prompt when information might be saved is more useful than a notice hidden in a policy. A direct pause is more useful than a buried setting. A deletion path must explain scope, timing, exceptions, and connected services.

Privacy information must be concise, transparent, intelligible, easily accessible, and use clear and plain language. (Guidelines, para. 7)¹⁰

ORION treats this as a design principle, not a compliance claim. Short, timely explanations can link to deeper detail. Physical and conversational controls need unmistakable ways to stop speech, pause movement, mute listening, end a session, and revise permissions, without negotiation from the character.

Warmth with an honest identity

ORION is intended to be warm without pretending to be human. Humour, affection, curiosity, visible expression, and a consistent character can support imaginative engagement. The product must remain identifiable as an artificial system and cannot claim feelings, private needs, or consciousness. Settings and help content can use plain terms such as

"generated response," "saved preference," and "system limitation." Mature, non-explicit intimacy follows the same rule: consent and boundary changes remain clear, and styling never overrides adult control.

"Warmth works only when identity stays honest."

Evidence before expansion

The final principle is operational: public language should move only as evidence moves. A feature idea begins as **Planned**. If an identifiable build is being evaluated against documented criteria, it may later become **In Testing** after owner review. If the wider product remains an early integrated build, it may be described overall as a **Prototype**. None of these labels means general availability.

AI RMF 1.0 is intended for voluntary use to help organizations manage AI risks. (Abstract, second paragraph)¹¹

ORION may use ideas from that framework when shaping internal review. Citation is not NIST validation, approval, compliance, or certification. The same separation applies to legal principles and industry guidance throughout this paper.


The ORION Experience

An illustrative day, not a testimonial

The following is an **illustrative fictional scenario**. Alex is a fictional adult in their thirties. Alex is not a customer or evidence, and the scene is not a report of current functionality, performance, satisfaction, or outcome. Every capability is **Planned** and pending product-owner verification.

The journey covers onboarding, character choice, conversation, selective memory, a boundary change, remote interaction, an everyday ritual, and mature non-explicit intimate companionship. Each beat remains an intended design discussion, not a capability claim.

Figure 1 maps the planned experience layers for orientation. It is a design reference, not evidence that any layer is implemented.



Alex first encounters ORION through a private pre-sale page. It states that the programme is a prototype in testing and pre-sale, explains pending feature status, and offers a private demonstration request. It makes no inventory or delivery claim. Before the demo, Alex sees the adult eligibility, data objectives, artificial identity, and session limits.

During onboarding and character choice, Alex selects calm, dry humour and a conversational pace with room for pauses. Examples are generated style samples, not statements from a conscious being. The choices remain revisable.

Next comes selective memory: no long-term memory, ask before saving, or save within selected categories. Alex chooses "ask before saving" and excludes health, finance, and information about other people. Storage, retention, review, export, and deletion are not owner-approved, and Alex can continue without long-term memory.

The first conversation is ordinary. Alex mentions making pasta after a long commute. The concept could offer company, a light topic, or quiet. It must not infer a diagnosis from tiredness. If the tone misses, Alex can change the style.

Alex asks ORION to remember that Thursdays are busy. The concept shows the proposed memory and asks for authorization. Alex can edit or decline it. A later review could show its purpose and change or deletion controls. This is a design objective, not verified behaviour.

Presence across settings

Alex then tries the planned remote interaction. The concept carries a familiar style and authorized preferences while showing the active device and account. Alex asks for a film title and closes the session. Quiet hours and notification controls treat silence as normal, not as a cue for pressure. Device support, authentication, synchronization, and remote availability remain unapproved.

Back at the physical demonstration, planned voice presence, visible expression, gaze and head movement, and embodied response could make turn-taking readable. These software-generated cues do not prove emotion, comprehension, or awareness. Movement, latency, noise, accessibility, and stopping need live evaluation.

For Alex, the appeal is an everyday ritual: one question about a creative project on some evenings, music and no conversation on others. The concept supports both rather than optimizing only for engagement.

Changing a boundary

Later, Alex explores mature non-explicit intimate companionship. Yesterday's comfort is not today's consent. In a boundary change, Alex moves from "affectionate" to "friendly." The character accepts without protest, bargaining, or reference to money, and the interaction follows the new boundary.

That moment is a core test. Reviewers need to see whether responses, memories, and physical behaviour follow the setting and whether the user can recover if the system gets it wrong. A clear stop path cannot depend on conversational cooperation.

Alex may enjoy the interaction, find it awkward, or decide it is not for them. None predicts a lasting outcome. A demo can inform a decision and product feedback, but an anecdote cannot become a claim about well-being, relationships, or safety.

The scenario's capability record

Capability ID: conversation

Status: Planned

Public description: ORION is intended to support everyday spoken conversation.

Limitations: Functionality, response quality, and availability require product-owner approval and a live demonstration.

Capability ID: visible-expression

Status: Planned

Public description: ORION is intended to show software-generated facial responses.

Limitations: Movement range, timing, and final availability require product-owner approval and a live demonstration.

Capability ID: gaze-head-movement

Status: Planned

Public description: ORION may include gaze and head movement where supported.

Limitations: Support varies by final version and requires product-owner approval and a live demonstration.

Capability ID: selective-memory

Status: Planned

Public description: ORION is intended to use user-authorized long-term memory.

Limitations: Storage behavior, retention, and user choices require product-owner approval and documented implementation evidence.

Capability ID: voice-presence

Status: Planned

Public description: ORION is intended to combine spoken interaction with a physical presence.

Limitations: Hardware, audio performance, and availability require product-owner approval and a live demonstration.

Capability ID: remote-mobile-interaction

Status: Planned

Public description: A remote identity is planned for supported messaging or web experiences.

Limitations: Remote access, device support, and account safeguards are not product-owner approved.

Capability ID: embodied-response

Status: Planned

Public description: ORION is intended to pair visible response with a physical form.

Limitations: Final interaction behavior and product configuration require product-owner approval and a live demonstration.

For the planned embodied-response concept, ORION is intended as a physical adult companion that may support consensual sexual and other intimate interaction; those capabilities remain Planned and require product-owner, safety, and legal review.

Capability ID: user-memory-controls

Status: Planned

Public description: Customers are intended to have clear control over memory and interaction preferences.

Limitations: Viewing, editing, exporting, and deleting data require documented policy and product-owner approval.

Capability ID: age-assurance

Status: Planned

Public description: ORION is intended only for legal adults.

Limitations: The age-assurance method, geography, and enforcement require documented policy and product-owner approval.

Capability ID: encryption-safeguards

Status: Planned

Public description: Encryption safeguards are planned for relevant data transfer and storage.

Limitations: Technical design, scope, and independent verification are not product-owner approved.

Capability ID: subscription-features

Status: Planned

Public description: Some functions may require internet access or a paid subscription.

Limitations: Included features, pricing, cancellation behavior, and availability require written policy and product-owner approval.

Capability ID: delivery

Status: Planned

Public description: Delivery planning is described for a future small-batch release.

Limitations: Shipping availability, timing, regions, and returns require written policy and product-owner approval.

Capability ID: warranty

Status: Planned

Public description: Warranty information is intended to be provided before purchase.

Limitations: Warranty terms, duration, exclusions, and geography require written policy and product-owner approval.

Capability ID: repair

Status: Planned

Public description: Repair support information is intended to be provided before purchase.

Limitations: Repair process, eligibility, costs, parts, and service regions require written policy and product-owner approval.


"Planned" keeps a coherent future experience from being mistaken for today's build. Product-owner review, documented policies, technical evidence, and live demonstrations decide what can move forward.

How ORION Works

A conceptual system, not a final specification

In plain language, ORION is a set of connected layers. Conversation prepares a generated response. Personality configuration shapes style within boundaries. Selective memory may offer authorized context. Voice and presence turn responses into speech and visible cues. Remote access may carry the configured identity to supported settings. A physical companion provides form, sound, and movement. Every capability is **Planned**.

Figure 2 summarizes the public labels used in this paper. It is an editorial status guide, not product evidence.



The architecture is conceptual. Final models, processors, providers, storage regions, sensors, protocols, and security controls are unresolved. They require documentation and owner review before a public specification.

Conversation and response generation

Speech or text input would pass through character settings, boundaries, and any authorized context before generated output becomes text, speech, or visible cues. Each stage can fail: speech can be misheard, context irrelevant, language wrong, or movement late. Testing must identify the failing stage and whether the adult can correct or stop it. Fluency is not authority; ORION is not a healthcare, legal, financial, or emergency professional.

Selective memory

Selective memory could bridge isolated sessions without preserving transcripts forever. An adult might authorize a preferred name, boundary, favourite genre, or routine. A real design needs categories, purposes, retention, access, review, deletion, and rules for information about other people.

Connected-product and related-service data must by default be easily, securely, and free of charge accessible to the user, with direct access where relevant and technically feasible. (Article 3(1))¹²

Application to a final ORION configuration requires legal review. Our design inference is that adults need practical, understandable access to retained information.

Personality configuration

Personality configuration means controllable presentation choices, not a manufactured soul. Tone, pacing, humour, topics, initiative, and forms of address may be adjustable. Guardrails remain outside character preferences: no style disables adult eligibility, stop controls, privacy notices, or crisis boundaries. Material changes require clear release notes.

Voice, visible expression, and physical presence

Voice and physical presence add timing, volume, location, gaze, and movement, increasing the cost of errors. Testing covers startup, shutdown, mute and stop paths, volume, movement limits, obstruction, lag, network failure, and accessibility. Visible expression is software-generated output, not evidence of understanding or feeling.

Remote mobile interaction

A planned remote identity raises account and device questions: sign-in, lock-screen content, shared devices, loss, synchronization, and the ability to disable access. Decisions require a threat model, authentication design, notification policy, data-flow map, and test evidence. Any dependency on remote services must be disclosed before purchase.

What the public labels mean

Prototype: This label describes an early integrated product or programme used to explore design, engineering, and user experience. It does not mean finished, certified, generally available, or ready for an announced delivery date. ORION's overall product stage is "Prototype in testing and pre-sale."

In Testing: This label is reserved for an identifiable capability actively being evaluated against documented criteria. It does not mean approved, complete, safe in every context, or promised for release. No capability in the current public register has this status.

Planned: This label identifies an intended capability or commercial process that remains pending product-owner verification. It may exist as a concept, requirement, design, mock-up, partial experiment, or unverified statement. It must not be described as implemented. All 14 public capabilities in this edition have this status.


The status date is 24 August 2026. Only product-owner review of specific evidence can move a status. A later edition must update the register, callout, limitations, and revision notice together.

Designed for Trust

Trust is a practice

Trust is not a guarantee. Connected products retain risk, and conversational systems produce unexpected output. The credible process is to minimize data, document flows, test controls, explain limits, respond to incidents, and revise claims. ORION has not completed independent security verification or legal compliance review. A data inventory would give each field a purpose, owner, retention rule, access condition, and deletion path. Timely explanations cover memory, remote services, and subscriptions before an adult makes the choice.

Figure 3 shows the proposed control loop as a design objective. It does not establish that the controls are implemented or effective.



Review, correction, export, and deletion

Planned memory controls aim to let adults review, correct, export, and remove saved items in plain language. They need to distinguish preferences from transcripts, show active settings, explain deletion limits, protect other people's information, and cover the physical device, remote identity, backups, support records, and subscription. A correction such as "forget that" needs a direct confirmation and must not create another unwanted memory.

Age assurance and adult positioning

ORION is intended only for legal adults. The method, threshold, geography, failed-check handling, privacy effects, and enforcement are **Planned**. A label is not a control; policy, implementation, accessibility, and legal review are needed. Mature companionship can be discussed candidly and without explicit detail. Marketing cannot target minors or depict boundary violations as intensity.

Consent across software and form

Consent is ongoing across topics, forms of address, affection, memory, notifications, voice, movement, and remote contact. One onboarding acceptance cannot settle later decisions. Adults need immediate ways to change, pause, and stop. The character cannot bargain, sulk, imply harm, or make paid features a test of loyalty.

Planned safeguards and external principles

Encryption safeguards are **Planned** for relevant transfer and storage. Algorithms, keys, endpoints, backups, access, logging, and independent verification are unapproved. ORION cannot guarantee privacy or security; review must examine the eventual architecture.

Claims that an AI service can substitute for a professional service require supporting evidence. (DoNotPay section, final paragraph)¹³

ORION is not presented as a substitute for a therapist, clinician, lawyer, financial adviser, crisis counsellor, or other professional. This evidence statement concerns the proposed DoNotPay order and that specific substantiation point; it is not an endorsement of ORION.

NIST materials offer risk language but do not certify ORION. EU texts can inform minimization, access, and clear explanations but do not establish applicability or compliance. Reviewers must examine the actual product, data flows, jurisdictions, and claims.

Misuse, error, and repair

A companion can fail without a breach: inventing a fact, missing a stop cue, exposing a notification, or applying a wrong preference. Testing needs event classes, reporting, and response ownership. Support receives only information needed for investigation. Mechanical and software repair must explain access, privacy protection, and whether data can be isolated before service. Material changes require clear updates to pre-sale participants.

Product and Ownership

A physical companion and a software service

ORION is conceived as an object and continuing software experience. The physical form, sound, and movement may have different lifespans and dependencies from conversation, memory, remote access, and updates. Before purchase, an adult needs to know which functions are local, online, account-based, or subscription-based and what happens after an outage, declined update, or cancellation. Those terms are unapproved; this paper commits to no bundle, price, billing interval, cancellation rule, or offline mode.

Customization with consequences made visible

Customization may cover appearance, character, voice, boundaries, memory, and services. Consequences for maintenance, compatibility, data, or updates need to appear with the choice. Software options may be reversible while physical choices may not be. A preview is a representation, not a guarantee of final colour, motion, sound, or fit. The adult configures generated behaviour, not a person or conscious being.

A proposed pre-sale journey

A proposed journey begins with stage, status, limitations, eligibility, and a clear distinction between private demonstration and general availability. Renders, mock-ups, and prototype footage need accurate labels. A demonstration identifies its build and which functions are live, simulated, limited, or unavailable, with stop controls, privacy boundaries, and feedback use explained.

Any later expression of interest or order needs approved terms before payment: configuration, total cost, taxes, schedule, cancellation, refunds, estimated timing, regions, privacy, software dependencies, warranty, and repair. None can be inferred here. Confirmation would provide a durable record and explain changes. If delivery is later approved, setup would cover adult verification, controls, care, support, and discreet handling without pressure to enable optional data uses.

Support, maintenance, and subscriptions

Support must be technically capable and discreet, collect only necessary diagnostics, and never demand an entire conversation history for an unrelated fault. Maintenance and care depend on the final materials. Any subscription needs approved features, renewal, cancellation, data consequences, and the service remaining afterward. The character cannot discourage cancellation. Warranty, repair, certifications, and regulatory assessments remain **Planned** or unresolved; final terms must precede commitment.

The meaning of a premium product

"Premium" describes attention, not mystique: suitable materials, careful onboarding, understandable controls, discreet service, and candid communication. It also respects a negative decision. The product may be personal; sales remain professional.

The Road Ahead

Near-term validation

Near-term work reduces uncertainty. Conversation needs identifiable builds, realistic tasks, latency and error measures, boundary tests, and correction paths. Movement and voice need range, stopping, repeatability, noise, accessibility, audio, and mute evidence. Memory needs data flows, purposes, retention, authorization, review, and deletion. Remote interaction needs account design, notifications, synchronization, and threat modelling. Age assurance, subscriptions, delivery, warranty, and repair need written policies grounded in intended markets and operations. Each capability stands on its own evidence.

Figure 4 presents that sequence as a planning reference, not a forecast of dates, approval, or delivery.



Product-owner gates

The product owner first verifies the exact build or policy, public description, limitations, reviewer, and date. Integrated testing then checks the seams: a boundary change must reach conversation, memory, remote access, and physical output; a stop remains available during failure; deleted preferences do not return. Specialists examine privacy, security, product safety, accessibility, consumer terms, and adult eligibility. Finally, approved facts must agree across the register, paper, demonstration, website, order documents, support, and release notes.

What remains aspiration

The integrated vision remains aspirational: voice, character, selected memories, remote presence, and physical form working coherently while the adult stays in control. That depends on data mapping, testing, mechanical reliability, account security, support, and honest sales practice. Broader character range, languages, devices, expression, or services add failure modes and maintenance commitments. Depth, comprehension, and control come before breadth. ORION will not diagnose, claim professional equivalence, simulate crisis intervention, or design for dependence.

Learning without overclaiming

Private demonstrations can reveal whether instructions are clear, controls discoverable, movement comfortable, settings respected, and artificial identity understood. They cannot establish long-term safety, general effectiveness, or health outcomes. Reviewers record the build, task, setting, prompt, behaviour, and issue severity. Positive reactions are not real-customer outcome stories. Any future public outcome claim must match its study design, population, comparison, duration, and result.

A measured invitation

ORION invites adults to inspect a prototype programme, ask direct questions, compare it with the dated status register, and decide whether the direction is relevant. Progress means narrower uncertainty: documented tests, understandable controls, specific policies, measurable limits, and supportable claims.

“Progress means narrower uncertainty, not a louder promise.”

Limitations and Responsible Use

This review draft describes intent, principles, and proposed operations. ORION is a prototype in testing and pre-sale. All registered capabilities are **Planned** and pending product-owner verification. The paper is not a user manual, contract, safety certification, legal opinion, clinical study, delivery promise, or statement of general availability. Final hardware, software, data flows, supported regions, service dependencies, prices, terms, and schedules are unresolved.

ORION is intended for legal adults. It is an artificial system, not a person. It can generate incorrect, inappropriate, repetitive, or unexpected language and cues. It does not possess emotion or consciousness. It is not a medical service, not a therapist, not a healthcare provider, and not a substitute for professional advice. It does not treat, cure, or prevent loneliness or any health condition.

Use should support the adult’s own choices and existing life rather than displace human relationships. ORION is not a replacement for human relationships. Adults should be free to stop, take a break, change boundaries, disable optional functions, and seek human or professional support. A product response should not be used as the sole basis for medical, legal, financial, safety, or emergency decisions.

Privacy and security safeguards remain design and review work. Adults should avoid sharing information they do not want processed and should consider the privacy of other people mentioned in conversation. Any final product must explain collection, retention, access, deletion, remote services, support diagnostics, and residual risks in approved documentation.

Crisis and Emergency Boundary

ORION is not an emergency service and is not designed or validated to detect, diagnose, triage, monitor, or resolve a crisis. It cannot dispatch help. Generated dialogue may be delayed, incomplete, or wrong.

As ORION responsible-use guidance, contact applicable local emergency services or local crisis services. Choose a service appropriate to your current location and circumstances. ORION cannot route, monitor, or complete that contact, and this paper makes no jurisdiction-specific factual routing claim.

Glossary

Artificial companion: A consumer product that uses generated interaction and, in ORION's concept, may combine software with physical presence. It is not a person or conscious being.

Capability status: A dated public label tied to one item in the product-status register and its stated limitations.

Consent and boundary controls: Proposed ways for an adult to select, change, pause, or stop forms of interaction. Their implementation is pending.

Data minimization: The practice of limiting collection and retention to information justified by a clear purpose. References to this principle do not establish ORION's legal compliance.

Generated response: Text, speech, or a visible cue produced by software. Socially meaningful presentation does not imply human emotion or consciousness.

In Testing: An identifiable capability actively being evaluated against documented criteria. It is not the same as approved, available, certified, or promised. No current registered capability carries this label.

Planned: An intended capability or process pending product-owner verification. It must not be read as implemented. Every registered capability in this edition is Planned.

Pre-sale: Activity before a final general release. In this paper it does not establish inventory, a delivery date, an accepted order, or approved commercial terms.

Private demonstration: A limited opportunity to inspect an identified prototype or concept under stated conditions. It is not evidence of general performance or availability.

Prototype: The overall early product or programme stage used for design, engineering, and experience evaluation. It does not imply that each described capability exists. ORION's programme is currently a prototype in testing and pre-sale.

Selective memory: The planned use of limited user-authorized information to support continuity, subject to unresolved storage, retention, review, export, and deletion design.

References

Revision Notice

Version 1.0, dated 24 August 2026, is a review draft for the Consumer Edition. Product-owner review, legal review, and English-editor review are pending. The capability register remains entirely **Planned**. This edition must not be published as a statement of general availability. A later revision should identify changed claims, capability statuses, policies, evidence, and review decisions.

Request a Private Pre-Sale Demonstration

Adults in Europe and North America may request a private pre-sale demonstration of the ORION prototype programme. Availability, location, format, and eligibility are not guaranteed. A request is not an order, reservation, delivery commitment, claim of general availability, or evidence that every planned capability can be shown. Ask to see the dated capability register, the specific build being demonstrated, its limitations, and the controls relevant to you before deciding whether to continue.

-
1. *Social connection linked to improved health and reduced risk of early death*. World Health Organization. Published 2025-06-30. <https://www.who.int/news/item/30-06-2025-social-connection-linked-to-improved-health-and-reduced-risk-of-early-death> ↗
 2. *Our Epidemic of Loneliness and Isolation: The U.S. Surgeon General's Advisory on the Healing Effects of Social Connection and Community*. Office of the U.S. Surgeon General. Published 2023-05-03. <https://www.hhs.gov/sites/default/files/surgeon-general-social-connection-advisory.pdf> ↗
 3. *Psychological, Relational, and Emotional Effects of Self-Disclosure After Conversations With a Chatbot*. Oxford University Press. Published 2018-05-30. <https://doi.org/10.1093/joc/jqy026> ↗
 4. *Is a random human peer better than a highly supportive chatbot in reducing loneliness over time?*. Elsevier. Published 2026-03-10. <https://doi.org/10.1016/j.jesp.2026.104911> ↗
 5. *AI companions and subjective well-being: Moderation by social connectedness and loneliness*. Elsevier. Published 2026-01-07. <https://doi.org/10.1016/j.techsoc.2026.103229> ↗
 6. *Interaction with AI companions and psychological well-being*. Springer Nature. Published 2026-08-04. <https://doi.org/10.1038/s41562-026-02516-2> ↗
 7. *Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile*. National Institute of Standards and Technology. Published 2024-07-26. <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.600-1.pdf> ↗
 8. *Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation)*. European Parliament and Council of the European Union. Published 2016-05-04. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX%3A32016R0679> ↗
 9. *Internet of Things: Privacy & Security in a Connected World*. Federal Trade Commission. Published 2015-01-27. <https://www.ftc.gov/system/files/documents/reports/federal-trade-commission-staff-report-november-2013-workshop-entitled-internet-things-privacy/150127iotrpt.pdf> ↗

10. *Article 29 Working Party - Guidelines on transparency under Regulation 2016/679*. European Data Protection Board. Published 2018-04-11. https://www.edpb.europa.eu/system/files/documents/2023-09/wp260rev01_en.pdf ↵
11. *Artificial Intelligence Risk Management Framework (AI RMF 1.0)*. National Institute of Standards and Technology. Published 2023-01-26. <https://www.nist.gov/publications/artificial-intelligence-risk-management-framework-ai-rmf-10> ↵
12. *Regulation (EU) 2023/2854 of the European Parliament and of the Council of 13 December 2023 on harmonised rules on fair access to and use of data and amending Regulation (EU) 2017/2394 and Directive (EU) 2020/1828 (Data Act)*. European Parliament and Council of the European Union. Published 2023-12-22. <https://eur-lex.europa.eu/legal-content/EN/TXT/PDF/?uri=CELEX%3A32023R2854> ↵
13. *FTC Announces Crackdown on Deceptive AI Claims and Schemes*. Federal Trade Commission. Published 2024-09-25. <https://www.ftc.gov/news-events/news/press-releases/2024/09/ftc-announces-crackdown-deceptive-ai-claims-schemes> ↵

ORION

ORION · Consumer Edition · Version 1.0

Review draft. Review gates: product owner Pending; legal Pending; English editor Pending.