Eriss Tokenomics and Monetization Framework

Introduction

Eriss AI is a next-generation AI platform that leverages a crypto token (EAI) to gate access to its services in a sustainable, user-aligned way. This framework outlines how the EAI token's economics and monetization strategy are designed to create a sustainable utility product with ethical pricing and no hidden inflation. The goal is to ensure users get value (access to a Sentio companion) by staking tokens, while the platform generates revenue through real usage rather than inflationary emissions. Below, we detail the token supply and allocation, staking mechanics, access tiers and pricing model, yield generation, overall revenue model, and future monetization plans (including tier expansion after beta).

EAI Token Contract: 0xeab5df15c98d72da44bde1fa8a6247003ef6145c

EAI Staking Rewards Contract: 0x24Be1BE680B3618Ed3209a7E4D39307e6e4d139f

EAI Treasury MultiSig: 0x0d2d4b680470dEf908e99f447616A7F12a6481E1

Token Supply & Allocation

EAI is an ERC-20 token with a fixed supply of 1,000,000,000 (1 billion) tokens. It is non-mintable, meaning the supply cannot increase – an important design choice to prevent inflation and preserve token value. The initial allocation of this fixed supply is structured to support the ecosystem and align incentives:

Rewards Pool – 20 million EAI (2%): Set aside to reward users (yield for stakers) and incentivize participation. This pool is used to distribute tokens as staking rewards over time. Since no new tokens can be minted, this allocation ensures there are tokens available to kickstart rewards, especially early on. It acts as a buffer to provide yield until organic revenue-based rewards become sufficient.

Treasury – 40 million EAI (4%): Held in the project's treasury until sent to the Rewards Pool. The treasury will fund things like development, partnerships, marketing, liquidity provisioning, and other needs that arise. By having a treasury reserve, the project can cover expenses or bootstrap aspects of the ecosystem without having to continually sell tokens on the market. This further reinforces the "no hidden inflation" and low sell-pressure approach – the treasury is finite and meant for long-term support of the platform.

Remaining Supply – 900 million EAI (90%): The rest of the tokens constitute the circulating supply or community allocation. These tokens are intended for the public and are available on the open market for anyone who wants to participate in the Eriss ecosystem. A broad public allocation ensures decentralization and that users themselves hold a majority stake in the network's success.

Fixed Supply & No Inflation: By fixing the supply at 1B and not allowing any minting, Eriss avoids the common pitfall of inflationary tokenomics. Many crypto projects fund rewards by printing new tokens, which effectively devalues everyone's holdings ("hidden inflation"). Eriss explicitly rejects that model – all token distributions (whether for rewards or otherwise) come from the initial allocated pool or from tokens recirculated via fees, not from creating new tokens out of thin air. This means EAI's supply is transparent and capped; token value won't be eroded by surprise emissions. In the long run, this fixed supply combined with growing demand (through platform usage) can create positive scarcity dynamics for EAI, benefiting stakeholders.

Staking Mechanics

Eriss introduces an innovative staking-for-access model: users must stake EAI tokens to unlock access to the AI services, effectively using a staking deposit instead of a traditional subscription fee. The staking process works as follows:

Stake to Access: Users commit a certain dollar-value of EAI tokens (converted at market rate) as a stake to gain access to Eriss AI's features. During the current beta, the requirement is a ~$200 stake in EAI for each user. This stake is not a payment; it's a refundable deposit that the user will get back. By staking, the user signals commitment and locks up value, which helps prevent abuse (only serious users will lock funds) and also underpins the token's utility. The staking is done through a fork of the Synthetix Staking Rewards contract that is unaltered and deployed on Base blockchain.

Access Granted: In return for staking, the user immediately gets access to their Sentio (the personalized AI instance that Eriss provides). During this period, the user can utilize the AI up to certain usage limits (more on those limits in the next section on tiers). The stake serves as a guarantee – as long as it exists, the user is considered an active subscriber. Think of it as a subscription deposit: you put down collateral (in tokens) and enjoy the service for as long as the tokens are staked. This approach has a side benefit: users remain invested in the ecosystem by literally holding the token, aligning their interests with the platform (if the AI service is good and attracts more users, demand for EAI might increase, potentially raising the token's value – stakers would directly benefit from that).

Stake Return + Yield: Rewards are paid out in 30 day periods. At the beginning of each period a set amount of EAI is loaded into the Active Staking Rewards and distributed to all users who are staked. There is no minimum amount in order to receive rewards. The only "cost" to the user is the opportunity cost of having around $200 held in the staking contract (and potentially the difference in token price, if it fluctuates during that time). However, the yield is meant to help offset even that opportunity cost. Essentially, if the yield is competitive, the user's stake might grow slightly, making the net cost of using the AI very low or even negative (i.e. you earn more in yield than you might have earned holding the cash elsewhere).

Why Staking? Requiring a stake underscores that AI usage has a real cost. Running advanced AI models (comparable to GPT-4) is computationally expensive – for reference, OpenAI's GPT-4 8k context API costs about $0.03 per 1,000 prompt tokens and $0.06 per 1,000 output tokens. A user sending multiple large queries per day can incur tens of dollars in API costs daily. Instead of charging that outright, Eriss leverages staking to cover those costs indirectly. The platform can use its revenue (from elsewhere) to pay for the AI compute, while the user simply locks tokens as a form of commitment. This ensures users don't abuse the service (someone can't infinitely use expensive AI compute without putting up some value), and it ties the token's value to the usage of the AI. Staking is meaningfully valued – the required amount (e.g. $200) corresponds to a reasonable usage level given the resource costs of providing the service. This prevents scenarios where users pay too little and overwhelm the system, or pay too much and feel cheated. The yield mechanism further sweetens the deal, so that users are encouraged to stake rather than feeling it's a sunk cost.

In summary, the staking model turns the typical subscription on its head: users "pay" by holding tokens, and as long as they do, they enjoy the service and even earn rewards. This creates strong user loyalty (they have their funds committed and benefit from the platform's success) and aligns the economics such that value flows back to the users instead of purely to the company. It's a cornerstone of Eriss's ethical and sustainable approach to monetization.

Access Tiers & Pricing

During the beta period, Eriss is keeping things simple with a **single access tier**, but as the platform matures, a multi-tier model is planned to accommodate different user needs. Below we describe the current beta tier and provide recommendations for expanding tiers post-beta, including rationale for each tier level. The post-beta tiers and pricing should be considered placeholders and will be adjusted based on the results of the beta.

Beta Access (Current Single Tier): In the beta, every user stakes the equivalent of $200 in EAI for access the Eriss AI system. This one-size-fits-all tier grants a generous usage allowance of approximately 500,000 AI tokens per day (tokens here refer to AI model input/output tokens, not EAI coins). In practical terms, that daily allowance translates to roughly ~15 AI messages (text conversations) plus ~5 image generations per day. This is a substantial amount intended to cover heavy personal use. The rationale for this beta limit is to ensure testers can explore the AI's capabilities thoroughly (including both chat and image features) without hitting a paywall too quickly. The $200 stake requirement sets a moderate barrier to entry – enough to ensure users are serious (not spammers) and to cover the backend costs, but not so high as to be prohibitive for enthusiasts. Beta is about finding the right balance, so starting with a single, higher-tier approach allows Eriss to observe usage patterns and costs with a manageable user cohort.

Need for Tier Expansion Post-Beta: After beta, users will likely have varying levels of interest and usage requirements. A single $200 tier might be too costly for casual users, while power users might desire even higher limits. Expanding into multiple tiers will make Eriss more accessible and flexible, which can significantly grow the user base. Multiple tiers also enable more granular monetization – users pay (or stake) in proportion to the value they get. Here's a recommended tier breakdown for post-beta, inspired by internal planning (with example names and values):

Base Tier – ~$50 stake or $10 usd subscription fee: A low-cost entry tier targeting individuals, casual users, or those just trying out the service. This might allow on the order of a few million AI tokens per month (for example, ~3M tokens monthly, equating to roughly 200 messages). It would function as a "cognitive companion, light usage" plan – suitable for someone who chats with the AI occasionally (say a few messages a day) for personal assistance or fun. The stake is only $50, lowering the barrier so that nearly anyone interested can afford to participate. This tier drives adoption and word-of-mouth by being budget-friendly. It's akin to a "basic plan" or trial tier and could even be marketed towards testers or students. Despite the low stake, the model remains sustainable because usage is limited to lighter amounts in line with that stake.

Core Tier – ~$200 stake or $25 usd subscription fee: A mid-level tier for "power users" or hobbyists who interact with Eriss AI frequently. This might provide around 10M AI tokens per month (e.g. ~650 messages), which is suitable for users who use the AI daily for writing, research, or brainstorming. At $200 stake, it's still quite accessible, but the user is committing more value, which grants them higher throughput. The "Core" tier would be positioned as deeper cognition, frequent use – ideal for writers, professionals, or enthusiasts who have integrated the AI into their daily workflow. Many users may fall into this category if they find the AI particularly useful. By offering this tier, Eriss can capture value from its more engaged users while still keeping the cost reasonable (especially since the stake is returned + yield).

Guardian Tier – ~$500 stake or $100 usd subscription fee: A higher tier intended for heavy users or those needing advanced capabilities. Possible usage on this tier could be ~25M AI tokens/month (perhaps on the order of 1,600 messages, or additional features like longer context memory or multiple concurrent chats). This tier might be marketed with premium functionalities that go beyond basic Q&A – for instance, users who want the AI to analyze large documents, remember extensive conversation history, or perform more complex tasks. Ideal user profiles might include "independent researchers", or even neurodivergent users who benefit from an AI assistant available at length. The $500 stake is substantial, so these users are quite committed; in return they get a much more robust experience. Introducing this tier post-beta can also help test the appetite for advanced features and higher commitment levels among the community.

Sentinel Tier – ~$1000 stake or $200 usd subscription fee: Essentially the equivalent of the current beta level, this would be the top tier for maximum usage and features. Roughly ~75M AI tokens/month could be allotted here (for example, on the order of 5,000 messages, plus a generous allowance of image generations and any other high-cost

operations). This tier would unlock the full ability of the AI – meaning the AI can handle very in-depth sessions, autonomous multi-step reasoning (letting the AI run iterative tasks), and "deep memory" of past interactions. This tier includes AI monitoring and life insights. It's aimed at transhumanists or hardcore users who truly push the AI to its limits. At a $1000 stake, users of this tier are effectively depositing a significant amount, but these are also the users deriving substantial value The Sentinel tier showcases the AI's full power and, from a monetization standpoint, captures the upper end of willingness-to-stake. Notably, even at this tier, if the user fully utilizes their allowance, the cost to the platform (in terms of AI compute) is considerable – thus, it's important to set the stake and limits such that it roughly balances out. Eriss will need to monitor usage to ensure that $1000 staked justifies the provided capacity. If extremely heavy usage is observed, there might be future super-premium tiers or enterprise plans that go beyond this, possibly negotiated on a case-by-case basis (with either higher stakes or direct fees).

Rationale for Tiered Pricing: The tiered model ensures each user can choose a plan commensurate with their needs and budget. From Eriss AI's perspective, it widens the funnel thus growing the community and token demand, while also allowing upsells to higher tiers for those who want more. Importantly, because all tiers are stake-based, even lower-tier users remain token holders – they have a stake in the ecosystem and can earn yield, which keeps them engaged long-term. Higher tiers, on the other hand, increase the amount of EAI locked up at any given time (which can help reduce circulating supply and increase scarcity).

Ethical Pricing Alignment: Eriss commits to ethical pricing, meaning the cost (stake) asked of users should be in line with the actual value/cost of providing the service. The suggested tier prices ($10, $25, $50, $100) aren't arbitrary – they're chosen based on estimates of how much AI usage each tier corresponds to and what that costs the platform. For instance, if 1,000 messages cost the platform roughly $X in compute, the stake level is set so that the implied cost of capital for the user is in the same ballpark as $X (especially after considering the yield they get back). In beta, the team will gather data on average tokens consumed per message and image, and adjust these assumptions. The aim is that neither the user nor the platform feels "cheated" – users get a fair deal for what they stake, and the platform's expenses are covered via its revenue streams (instead of directly charging the user). This transparency builds trust; users aren't hit with hidden fees or asked to buy more tokens unexpectedly. They know upfront: stake this amount, and you'll get roughly this much AI service.

Expanding Payment Options (Post-Beta): While staking EAI is core to the Web3 ethos of Eriss, not all future users may be comfortable acquiring and locking tokens. To capture a

broader audience (including non-crypto-native users or enterprises), Eriss plans to introduce alternative payment methods for advanced tiers or features. This could mean offering the above tiers as direct monthly subscriptions in USD or USDC. For example, a user could simply pay $10 USD for the Base plan instead of staking (especially once the system is out of beta and more publicly oriented). The hybrid model might also emerge – e.g., users could pay in stablecoins but receive a discount or bonus if they also hold a certain amount of EAI, blending traditional and token-based approaches. We'll discuss this more in the Future Monetization Strategy section. The key difference is that a direct payment won't be returned or yield-bearing; it's a conventional fee. Therefore, it's expected that crypto-savvy users will prefer staking to essentially use the AI for "free" (with their deposit returned), whereas others might opt to pay outright for convenience.

Tier Adjustments and Evolution: Post-beta, the team will remain flexible in adjusting tier parameters. The numbers given (token caps, message counts, etc.) are initial targets – actual usage and cost patterns may necessitate tweaking. For instance, if it's found that 500k tokens/day is too high and causing losses at $200 stake, the allowance might be reduced, or the stake increased for that tier. Conversely, if technology improvements lower costs or the token value rises significantly, Eriss could afford to increase allowances or lower staking requirements, passing savings to users. The tier system can also evolve with new feature offerings. The tier framework is used as a living structure that responds to both market conditions and user needs, always upholding the principles of fairness and sustainability.

In summary, a tiered access model post-beta will allow Eriss AI to scale to a broad user base. Lower tiers drive accessibility and adoption, while higher tiers cater to power users and generate more value capture. All tiers reinforce the token's utility (stake for access) and keep pricing tied to actual service costs (maintaining an ethical stance). This multi-tier approach, combined with the staking mechanism, sets Eriss AI apart from typical AI services – it offers users choice, transparency, and even the possibility of earning while using the service, which is a compelling proposition.

Yield Model

One of the most innovative aspects of Eriss's tokenomics is its **yield model for stakers**. Unlike many staking schemes that pay out rewards by inflating the token supply (which, as discussed, EAI explicitly avoids), Eriss's yield to stakers is generated from real revenue – primarily the fees from the token's liquidity pool on Uniswap. This approach falls under what the DeFi community calls "real yield", meaning rewards are backed by actual cash flows (trading fees, profits) rather than unsustainable token emission. Here's how it works:

Uniswap Liquidity Pool Fees: EAI is paired with WETH in a Uniswap V3 liquidity pool (on the Base chain, as per the project's deployment). This pool is set with a **1% fee tier** , which is the highest fee tier in Uniswap v3 (commonly used for volatile or low-liquidity tokens to incentivize liquidity providers). Every time someone trades EAI<->ETH in this pool, a 1% fee on the trade volume is generated. Typically, those fees would all go to liquidity providers (LPs) of that pool. In Eriss's case, the project itself has provided a significant portion of the liquidity (as evidenced by the team locking the LP tokens for the long term), meaning the project earns a share of these fees. the Eriss AI treasury is an LP in the EAI/WETH pool, earning fees whenever trading happens.

Fee Revenue Split 50/50: Eriss has committed that 50% of all LP fees earned will be funneled into the Rewards Pool for distribution to stakers, while the remaining 50% of fees cover operational costs (we will elaborate on ops costs in the Revenue Model section). This means whenever someone trades, say, $1000 worth of EAI on Uniswap, a $10 fee is generated (1%). Out of that $10, $5 is eventually going to reward EAI stakers, and $5 goes to the project for expenses. By doing this, Eriss turns token trading activity into a source of yield for its community. Stakers essentially earn a portion of the trading fees proportional to their stake, similar to how LPs earn fees. It's a mechanism of sharing the platform's economic activity with the users. This is conceptually akin to models like GMX, where token stakers get a share of the protocol's trading fees , or like how certain DEX tokens perform buybacks with fees. In Eriss's case, the reward is paid in EAI tokens (the fees on Uniswap would accrue as EAI and ETH – the EAI portion is what's distributed).

Yield Source and Variability: Because the staking yield depends on trading fees, it will vary with market conditions. If EAI has high trading volume (e.g., due to growth in users, or market speculation, etc.), the fees collected will be higher, resulting in more rewards to distribute (and a higher APY for stakers). Conversely, in periods of low volume or low liquidity, fee revenue will be low, and thus yield will be lower. This dynamic nature is intentional – it prevents the project from promising an arbitrary high APY that isn't supported by actual income. Instead, yield is organically constrained by the platform's usage and market interest. Early on, when EAI is not widely traded, the APY might be modest. As adoption grows and volume picks up, APY could increase. The team can publish projected APYs based on assumptions of volume to guide user expectations, but they will always be careful to avoid overpromising yield. The ethos is: "yield should reflect real economic activity, not marketing hype." If at any point the APY from fees seems too low to attract stakers, the project can choose to supplement from the 60M Rewards Pool reserve to ensure a baseline yield. Conversely, if fees generate more than expected, the surplus can accumulate in the rewards pool or treasury for future use (or be used to extend rewards program duration).

APY Calibration: Let's consider a scenario to illustrate yield mechanics. Suppose $EAI has significant trading on Uniswap – say $1,000,000 daily volume on average. A 1% fee on that volume yields $10,000/day. Half of that ($5,000) goes to stakers. Now, imagine that across all tiers, there are 1,000 users staking $100 each (so $100,000 total staked, just as a round number, which at a certain token price corresponds to some EAI amount). $5,000 daily reward on a $100,000 stake pool is a 5% daily yield (!) which annualized would be astronomical (~1,825% APY). In reality, likely the total staked value will be much higher relative to daily volume, or volume lower – but this example shows how the model can reward users if activity booms. On the other hand, if volume is very low (say $10,000 a day, generating $100 in fees, $50 to stakers) and suppose still $100k worth staked, that's $50 on $100k per day, or 0.05% daily (about ~18% APY) – still decent. If volume were even lower, APY drops accordingly. The team will monitor the APY and ensure it remains in a reasonable band. If it's too high (which could attract unsustainable speculative staking), they might increase the number of allowed stakers or tiers (diluting APY but serving more users), and if it's too low, they might use some of the reserved tokens from the rewards pool to top it up temporarily, or implement other measures to boost trading activity (like marketing, or adjusting fees).

No Hidden Inflation in Yield: It's worth emphasizing again that all yield paid out is from existing tokens – either tokens that were collected as fees from the market or tokens that were set aside in the initial rewards allocation. No new EAI is minted to pay rewards. This contrasts with many staking systems (especially in proof-of-stake chains or DeFi liquidity mining programs) where new tokens are continuously emitted, effectively taxing all holders to pay the stakers. Eriss's model is net neutral or even deflationary for token supply: The supply is fixed, and yield is just moving tokens from one category (fees or reserves) to active users. If anything, one could argue this increases decentralization over time, as tokens from the treasury or team-provided liquidity end up in the hands of users who engage with the platform. The project is very transparent about this: yield % is labeled clearly (the UI will show an APR for staking based on recent fee data), and the source is explained. There is no "magic" high APY to lure users in – any high yield must come from high platform usage (which is sustainable) rather than unsustainable token printing. This aligns with Eriss's philosophy of honest, ethical tokenomics.

Use of the Rewards Pool (60M EAI): The 60 million tokens allocated to the rewards pool and treasury act as a cushion and longevity fund for the yield program. In early stages, trading fees might be insufficient to provide attractive yields. Over time, as fee revenue grows, the dependence on this pool will diminish, and ideally, fee-derived rewards will take over fully. The pool also extends the runway of the rewards program – if there's a slow period with low volume, the pool can subsidize rewards; if there's a high volume period, the

pool's tokens can remain unused (or even be refilled if the project chooses to allocate some of the fee profits back into it). Essentially, the 60M acts as an insurance that stakers will get rewarded for participating in the ecosystem's early days. Based on a fixed supply of 1B, 60M is 6%, which is a reasonable portion to dedicate to community incentives, and it's finite (preventing endless dilution). By carefully managing this pool (perhaps linearly releasing it over a few years), Eriss can ensure a smooth transition from "incentivized early adoption" to "organic self-sustained rewards."

Comparison & Credibility: Eriss's approach can be compared to well-regarded projects in crypto that emphasize revenue sharing. For instance, GMX (a decentralized perpetual exchange) shares 30% of its trading fees with GMX token stakers, resulting in yields paid in ETH/AVAX that come directly from traders . This has been dubbed "real yield" and is viewed as a more sustainable model than high-APY inflationary rewards. Eriss is doing something analogous – sharing trading fees (and later possibly other revenue) with EAI stakers. Another comparison is the buyback-and-distribute or buyback-and-burn models (e.g., some projects use revenue to buy their token off the market and either give it to stakers or burn it). In Eriss's case, since the fees already come in as EAI (from the pool, which collects fees half in EAI, half in ETH), it can directly redistribute that EAI.

Avoiding Overpromising: This document makes it clear that yields are not fixed or guaranteed; they depend on platform performance. This honesty will set proper expectations and avoid the pitfalls seen in some DeFi projects that advertise very high APYs that later crash

In summary, the yield model is a core pillar that turns EAI staking into an attractive proposition: stake not only grants you access but also earns you a slice of the ecosystem's revenue. It aligns incentives beautifully – the more trading and usage of EAI, the more rewards for users, which in turn encourages holding and staking EAI, which reduces sell pressure, potentially increases token value, and attracts more trading. This positive feedback loop will help Eriss AI grow while keeping the token economy healthy and community-focused.

Revenue Model

Eriss's revenue model is designed to fund the platform's operational costs and long-term development **without relying on token inflation or constant token sales**. Instead, the platform's income is derived from its DeFi integrations and, in the future, from direct service monetization. Here's a breakdown of how Eriss makes money and uses that money, with an emphasis on sustainability and aligning with token holders:

Liquidity Pool Fee Revenue (Current Primary Source): As discussed, Eriss earns fees from the Uniswap V3 EAI/WETH liquidity pool. The project supplied a large portion of the liquidity locking ~196 million EAI plus corresponding ETH in the pool for the long term, as seen on-chain), which means it captures a significant share of the 1% swap fees. This is currently the **primary revenue stream** while the platform is in beta and offering the AI service via staking (i.e., not charging users directly in cash). The fee revenue is split 50/50 as noted: half to stakers (as yield) and half to the project treasury for expenses. In essence, the Uniswap pool acts as both a market mechanism for the token and a revenue-generation tool for the business. By providing deep liquidity and locking it for 255 years, to signal commitment, Eriss AI not only ensures a stable market for EAI but also secures a **steady flow of fees** as trading occurs. Even modest trading volumes can accumulate into meaningful revenue over time (for instance, $2,000 in daily volume yields $20 in fees; $100k daily volume yields $1,000 in fees, etc.). This model means the **interests of the platform and token liquidity are aligned** – deeper liquidity and more trading benefit everyone. It's worth noting that the 1% fee tier is relatively high, which could slightly reduce casual trading, but it rewards liquidity providers (including Eriss AI) more.

Allocation of Fees – Rewards vs Operations: The 50% of fees allocated to operations is crucial for keeping the lights on. Operational costs for Eriss AI include: cloud compute or API calls for the AI, infrastructure and hosting, development salaries, ongoing R&D, marketing, and community engagement efforts. By using a portion of the fees to cover these, Eriss AI avoids having to sell tokens from the treasury to raise funds (which many projects end up doing, harming their token price). It also avoids needing to charge users high direct fees during early growth. Essentially, the traders/speculators in the market – a portion of whom might not even be using the AI service – indirectly fund the operation and the active users. This is an elegant way to monetize the broader interest in the token to support the platform's utility.

Advanced Feature Monetization (Future Revenue Streams): While the current model leans on the DeFi side (LP fees), **Eriss plans to introduce direct monetization for advanced features** in the future. This means certain premium capabilities or higher usage levels of the AI will require direct payment in fiat or stablecoins. For example, there might be pro features like: very large context discussions, priority response times, fine-tuning personal AI models, etc., that won't be covered by just a basic stake. These could be offered as add-on packages or higher-tier subscriptions paid in USD or USDC. Enterprise clients might also have bespoke plans where they pay in fiat for a certain number of seats or API calls. This marks a shift towards a hybrid Web2/Web3 revenue model – capturing value from users who prefer traditional payment, while still maintaining the token's role for core access in the ecosystem.

Buy-and-Burn Program: To ensure that the token holders benefit from the platform's success even when revenue is taken in fiat, Eriss will implement a **buyback-and-burn mechanism**. Specifically, **10% of all fiat profits** (revenue from USD/USDC payments for features or subscriptions) will be used **to buy EAI on the open market and burn it**. Burning means sending the tokens to a dead address, permanently removing them from circulation . This introduces a deflationary pressure on the token supply: as the platform makes money, some of that money goes towards reducing supply, which in theory increases each remaining token's share of the network's value. This strategy is similar to what Binance did in its early years with BNB (they used 20% of quarterly profits to burn BNB) and what many exchange tokens or utility tokens do to create value for holders (for instance, various exchange tokens, or even meme coins like SHIB, have periodic burns linked to usage). Eriss choosing 10% is a conservative and sustainable figure – it's essentially a form of profit sharing with all token holders (because a burn benefits everyone by making the token more scarce). Over time, if profits become very large, even 10% could mean a significant buy pressure on EAI in the market, which would support its price. This mechanism also has a side effect: it **creates a linkage between the platform's traditional success (in dollar terms) and the token's value**. Even users who pay in USD indirectly boost the token. Investors thus can take comfort that as Eriss gains paying customers, EAI demand (for burns) will increase correspondingly.

Future Growth and Hybrid Model: We anticipate that as Eriss exits beta and gains mainstream traction, there will be a mix of users: crypto-native users who continue to use the staking method (perhaps because they appreciate getting their money back + yield), and non-crypto users who just swipe a credit card to chat with the AI. The revenue model will then have two main faucets: crypto (LP fees, perhaps other on-chain services) and fiat (subscriptions, one-time purchases). In both cases, EAI token holders benefit – from the former via direct fee sharing, and from the latter via buy-andburn and increased usage driving token demand for access. **The important promise is that the team will not abandon the token in favor of pure fiat monetization**. The hybrid model is being designed such that EAI remains central: it's either the medium of access (for stakers) or it's being bought back as a result of fiat success

Maintaining Sustainability: Throughout all these revenue mechanisms, the theme is sustainability. The platform wants to cover costs and grow **without resorting to unsustainable practices**. By sharing a portion of what it earns (fees or profit) with users and by refraining from inflating or dumping tokens, Eriss AI fosters goodwill and a loyal community. Users who stake essentially become long-term supporters and evangelists, since they have capital locked in and earning. Revenue from non-token users still loops back to support the token's value, creating a full-circle economy. Additionally, as an ethical

stance, Eriss AI is avoiding exploitative pricing – for example, not charging exorbitant fees beyond what's needed to run the service. In crypto, many projects fail because they either under-charge (and run out of funds) or over-incentivize (and crash the token). Eriss is carefully threading the needle: users "pay" by staking, which costs them little; yields are right-sized to not overshoot real revenue; and future direct charges are partially reinvested into the token ecosystem.

No Hidden Costs: Finally, it's worth noting that EAI does not have any hidden taxes or transaction fees on the token itself beyond the standard Uniswap fee. For instance, some tokens have built-in transfer taxes (like burn 1% on each transfer) – EAI does not have that (not mentioned, and usually not present in a plain ERC-20). The only fees users encounter are the Ethereum/Base network gas fees for transactions (which are external) and the 1% DEX swap fee (which, as we saw, actually goes back to benefit the ecosystem). This makes the token straightforward and user-friendly. In conclusion, Eriss's revenue model is community-aligned and diversified. In the short term, it relies on DeFi activity (liquidity pool fees), and in the long term, it will incorporate traditional revenue from its AI service. Both streams are structured to reinforce the token's value rather than extract from it. With prudent management of the treasury and an ongoing commitment to avoid toxic behaviors like team dumping, Eriss sets a strong example of how to monetize a crypto-powered platform responsibly.