MICRO CONNECT 滴灌通

NUMA Monthly

Published November 3, 2026



In This Issue

Market Design | Reaffirming the “First Investor”'s Role and Participation Arrangements

Market Progress | NUMA Operating System OS: The Three Systems

Market Preview | Simulation Testbed Operating Plan

A Letter to Readers

Dear Readers,

Over the past three months, we have published white papers on August 3, September 3, and October 3. They addressed, respectively, NUMA's market framework and basic rules, methods for measuring contracts and businesses, and ways to organize human-AI collaboration. This work laid the foundations for building the market. Our next step is to put these rules, methods, and divisions of responsibility into practice. Through simulation, we will continue to refine design and execution, improve efficiency, reduce costs, and establish accountability at every stage.

Beginning this month, we will publish NUMA Monthly on the third day of each month. We will continue to publish white papers to present systematic findings at key stages. The monthly publication will provide ongoing updates on our work, the problems we identify, our next steps, and revisions needed to documents already published.

The August 3 white paper established that only a party that first contributes value and bears risk under a Real World Contract is eligible to bring the related rights to market. The September 3 white paper then introduced the name “First Investor”. The “First Investor” has already contributed value to a business activity and is contractually entitled to subsequent cash flows; the “First Investor” is also the Protocol Contract Issuer. This issue's Market Design column explains how a “First Investor” can share future cash flows from existing contracts to recover some capital and support further contributions. It also explains the risks borne and interests retained in the process.

This issue focuses on the Market Progress column, which introduces the draft framework for the NUMA Operating System (OS).

The OS is the overall arrangement for organizing NUMA's day-to-day work. It translates market rules into business processes, breaks professional work into specific tasks, and defines AI's execution authority and human responsibilities. This is an initial design that we will continue to test and revise through the simulations ahead.

The arrangement builds on the HAT Over MASS method of human-AI collaboration introduced in the October 3 white paper. It divides work among three systems: the Responsibility Matching System, the Transaction Matching System, and the Information Matching System. The Responsibility Matching System establishes accountability through Human–AI Cockpits, defining what AI may do and which matters require human judgment and sign-off. Work and responsibilities must also be handed over clearly when staff take leave or change roles. The Transaction Matching System handles specific business processes, from verifying materials, registration, and measurement to transaction coordination and cash-flow reconciliation. It defines how each step is handled and checked, and retains the corresponding records. The Information Matching System connects the two sides of the market. Through content support and services for off-market intermediaries, it helps “First Investors” and Professional Investors understand the market and prepare to participate. Activities requiring a license are handled by the relevant licensed institutions.

Through this design, we hope AI can take on more of the repetitive work suited to it, allowing people to concentrate on judgments, decisions, and accountability. We will share simulation results, and the problems they reveal, candidly, and continue to improve system design and execution. We aim to reduce contract-processing costs and serve more participants while maintaining quality and safety.

This issue's Market Preview column introduces the Simulation Testbed Operating Plan. On November 3, we plan to begin a two-month program of internal simulation and testing, gradually incorporating different business activities, participant roles, and scenarios. Alongside normal processes, we will simulate incomplete materials, conflicting judgments, staff absences, authorization conflicts, and system exceptions. These tests will examine

whether processes work from end to end, rules are implemented, and work can continue through handovers. We will also record the time and cost of the entire process, improve the design in response to identified problems, and test whether the revisions are effective.

This November 3 issue focuses on the three systems, explaining what the OS comprises and what each part is responsible for. Around December 3, we plan to publish NUMA Operating System OS: The Three Engines, explaining how this work is allocated, matched, and executed. We also plan to publish an interim simulation report at that time, setting out which stages have been tested, what problems have been found, and what revisions have been made. We may also invite external participants, subject to testing needs and readiness.

Around January 3, 2027, we plan to bring together the two issues and simulation records to finalize the OS for this stage. We will publish our fourth white paper, NUMA Operating System OS: The Three Systems and Three Engines, along with simulation results available at that time. We will explain what has been tested, what still requires testing, and which problems remain unresolved. Actual business operations remain subject to the relevant readiness and approval conditions.

This issue's Continuity and Revisions section explains how the content relates to the three earlier white papers: what continues to apply, what has been supplemented or adjusted, and why and where revisions have been made.

Building a market takes steady work and continuous improvement based on actual operations. Through this monthly publication, we will keep reporting our progress and candidly explaining the problems we encounter, the judgments we make, and the revisions we adopt.

Thank you for your interest in and support for NUMA. We welcome your comments and suggestions as we work, step by step, to build a sound foundation.

Charles Li
November 3, 2026

Contents

A Letter to Readers 1

Part I | Market Design

Reaffirming the “First Investor”'s Role and Participation Arrangements 6

I. Market Positioning: A Contract Market for “First Investors” 6

II. Participants and the Three Contract Layers 7

III. “First Investors”: Contributions, Qualifications, and Profiles 7

IV. Why “First Investors” Participate: Many Opportunities, Limited Capital 9

V. Cash-Flow Arrangements, Contract Pools, and Retained Shares 10

VI. Measurement and Ongoing Disclosure 13

VII. Transactions Within the Market, Connections Outside It 14

VIII. Design Principles and Business Boundaries 15

Part II | Market Progress

NUMA Operating System OS: The Three Systems 16

I. Design Objectives and Overall Architecture 16

II. Responsibility Matching System (RMS) 18

III. Transaction Matching System (TMS) 23

IV. Information Matching System (IMS) 28

V. The Contract Lifecycle and Collaboration among the Three Systems 29

VI. Common Controls and Technical Support 32

Part III | Market Preview

Simulation Testbed Operating Plan 36

I. Positioning, Objectives, and Outputs of the Simulation Testbed 36

II. Operating Architecture and Coverage 37

III. Simulation Methods and Scenario Design 39

IV. Operating Organization and Iteration Mechanisms 40

V. Operational Observations and Verification Requirements 42

VI. Simulation Testbed Methods and Design Features 43

VII. Data, Participation, and Simulation Boundaries 44

VIII. Phased Schedule, Progress Disclosure, and Expected Outputs 45

IX. Applicability and Limitations of Simulation Results 46

Part IV | Continuity and Revisions

Appendices

Appendix I. Terminology and Historical Documents 51

Appendix II. Main Tasks and Microtask Examples in the Transaction Matching System 60

Important Notices and Guidance for Use

Nature of This Publication and Relationship among Versions

NUMA Monthly reports on progress in building NUMA, results at key stages, future plans, and supplements, clarifications, and revisions to published documents. It builds on the three white papers published on August 3, September 3, and October 3, 2026. Updates explicitly identified in this issue apply only to the relevant matters. Content not addressed here should continue to be read together with the original documents and their guidance for use.

Status of Results and Limits on Use

This publication describes the status of business design, rule development, system construction, validation involving specific enterprises, and simulation operations according to actual progress. That a function has been designed does not mean it is available; that rules have been developed does not mean they have been validated with a specific enterprise; and simulation operations are not actual business operations. The relevant methods, rules, and systems will continue to be refined using actual data, operational testing, and feedback from the parties involved.

This publication and the Industry Compasses, Compass readings, NUMA Adjustment Factors, and other content it references or presents introduce the relevant designs, methods, and interim results for reference in general industry research, business analysis, and methodological discussion. Actual use of the relevant professional tools or services is also subject to applicable eligibility, authorization, and conditions of use. The contents of this publication do not constitute investment advice, recommendations, an offer, invitation to offer, sale, or investment solicitation relating to securities, contract rights, or other financial products. Nor do they constitute a commitment regarding returns, legal advice, or compliance approval. The relevant scores and factors do not represent investment returns, default probabilities, or credit ratings. Any actual investment decision should be based on independent judgment, due diligence, and necessary professional advice.

Nature and Limitations of Simulation

The simulation plan described in this issue will be conducted in a simulated environment. All simulated contract, investor, fund, and transaction data are synthetic. The simulation does not connect to real accounts, handle real funds, conduct real transactions, accept investment payments, or raise capital. It tests rules, processes, systems, and human-AI collaboration. Its results reflect performance only under the specified conditions and cannot replace testing in a real market. They do not constitute actual investment performance, forecasts of future returns, or guarantees of returns. Taking part in the simulation does not confer eligibility to participate in investments.

Conditions for Business Operations and Allocation of Responsibilities

This publication separately describes designs and results already developed and work still under development and testing. Statements about future arrangements reflect plans at the time of publication and do not constitute commitments regarding launch dates, service availability, or investment outcomes. Actual business operations and related adjustments must comply with applicable laws, regulatory requirements, and the applicable rules of the relevant licensed trading venues. Required approvals, filings, and other necessary procedures must also be completed.

AI organizes evidence, performs checks, and develops recommendations within its authorized scope. Final scores must be determined or confirmed by authorized humans under the existing sign-off and authorization rules. Investors make their own investment decisions. Suitability, compliance, and legal judgments must be made or confirmed by people with the relevant responsibilities and the necessary qualifications and

authorization. Services that legally require a license must be provided by appropriately licensed entities within their licensed scopes. Internal authorization or human confirmation cannot replace licensing requirements.

Readership and Eligibility for Investment Participation

This publication provides general research information to the public. Making this publication public, discussing research, or offering participation in the simulation does not constitute the provision of investment services to readers or confirmation of their eligibility to participate in investments. Actual investment participation is restricted to eligible Professional Investors outside China who meet the requirements of the applicable jurisdictions, and must take place through appropriately licensed intermediaries. It is not open to retail investors. NUMA does not raise funds within China or accept funds from within China for investment.

Copyright and Citation

The relevant intellectual property rights in this publication belong to their respective rights holders under applicable law. Citations should identify the source and publication date and accurately reflect the original meaning and relevant limitations. Unless permitted by applicable law or separately authorized, copying, republication, translation, or adaptation of all or a substantial part of this publication requires the relevant rights holder's prior written permission. Use of this publication's content and related names and logos must not create the false impression of authorization, approval, or endorsement.

Part I | Market Design

Reaffirming the “First Investor”'s Role and Participation Arrangements

Overview

The white papers published on August 3, September 3, and October 3 established, respectively, NUMA's General Protocol and three-contract structure, the Twin Compasses measurement method, and an organizational arrangement in which AI staff execute tasks and humans remain accountable. Building on the existing market framework, this issue revisits what the “First Investor” contributes in advance, how it supplies contracts, the risks it bears, and how it stays involved. It further explains cash-flow arrangements, Contract Pools, and how the financing and investment sides enter the market.

NUMA connects the contract cash-flow rights that “First Investors” have already created with capital from Professional Investors. The “First Investor” initiates this connection by identifying opportunities in an industry, contributing value, and bearing risk, thereby creating contracts to bring to market. Once operations begin, the “First Investor” maintains contact with the underlying business operator to monitor contract performance and obtain operating information. The “First Investor” initiates the market's contract supply, and subsequent investors build on a judgment the “First Investor” has already acted upon. Measurement and disclosure over the life of the contract also require the “First Investor”'s continuing involvement. Understanding the “First Investor” is therefore fundamental to understanding how NUMA takes shape and continues to operate.

This column introduces the “First Investor”'s prior contributions, qualifications, and typical profiles, explaining the motivation to participate in the market. It then describes cash-flow arrangements designed for the characteristics of contract cash-flow assets, participation through individual contracts and Contract Pools, and the principle of retaining an ongoing interest. Next, it explains how market measurement and ongoing disclosure inform Professional Investors' judgments. Finally, it explains how on-market transactions and off-market connections bring “First Investors” and Professional Investors into the market.

I. Market Positioning: A Contract Market for “First Investors”

Activity in the real economy generates large numbers of underlying contracts (Real World Contract, RWC). These contracts generate operating cash flows. Yet existing markets offer limited coverage for investing in, measuring, and trading contracts that are smaller, dispersed, and less standardized. Digitalization enables original evidence relating to contracts to be recorded and verified, and artificial intelligence (AI) may reduce the cost of processing each contract.

Micro Connect New Markets (NUMA) positions itself as a contract market for “First Investors”. This positioning has three dimensions:

(1) Listings are based on contracts backed by actual prior contributions, legally effective, actively being performed, and compliant with the platform's admission standards.

The party that first contributes resources, capabilities, or capital and is contractually entitled to subsequent returns is the “First Investor”. For example, a commercial property operator draws on a location's advantages, provides premises, invests in operations and renovation, and attracts high-quality brands, thereby creating lease contracts. A chain brand opens its franchise network and provides franchisees with brand, training, and operating-system support, thereby creating franchise contracts.

(2) Investors acquire the agreed rights to contract cash flows.

Professional Investors acquire rights to distributions from a contract's future cash flows, according to an agreed share of cash-flow rights or distribution priority. The specific arrangements are agreed and registered before

the transaction. The amount and timing of cash flows depend on the contract terms and actual business operations and contract performance. Investors decide independently whether to participate, based on the relevant terms, operating records, and disclosed information.

(3) NUMA provides market infrastructure for qualification, registration, measurement, disclosure, and matching.

NUMA verifies admission qualifications, registers cash-flow rights and changes to those rights, and measures the returns and risks of contract cash flows. It fully discloses contract terms, interest arrangements, and performance records, and provides matching support for the relevant materials, information, conditions, and processes. Activities that require a license are handled by the relevant licensed institutions within their responsibilities and licensed scope. NUMA's systems prepare proposals, transmit instructions, and reconcile results; they do not replace licensed institutions in confirming and executing transactions. These arrangements support participants in making independent judgments and conducting transactions, allowing market mechanisms to function fully.

II. Participants and the Three Contract Layers

The market has three main categories of participants:

1. The “First Investor”: the party that has already contributed value to a business activity and is the Protocol Contract Issuer (PCI).
2. The underlying business operator (Real World Contract Counterparty, RWCC): the party that receives the “First Investor”'s contribution under the Real World Contract, conducts the business activity, and pays returns to the “First Investor” as agreed.
3. The Professional Investor (PI): an investor admitted through a licensed institution who meets the qualifications for a Professional Investor.

In addition, licensed trading venues and their members, licensed brokers, and licensed custodian banks handle listing, transaction execution, client identity and qualification verification, and custody of funds within their respective licensed scopes. Off-market intermediaries help clients on both sides enter the market. They are not parties to the transactions; their role is described in Section VII of this column.

NUMA provides market infrastructure, including registration, measurement, disclosure, and matching. The parties define their rights and obligations through three contract layers. This basic structure is carried forward from the August 3 white paper:

1. Real World Contract (RWC): the contract between the “First Investor” and the underlying business operator, specifying the value contributed in advance and the corresponding rights to cash flows.
2. Protocol Contract (PC): the contract signed by the “First Investor” and the NUMA operator, defining the scope of cash flows brought to market, the distribution method, and the portion retained by the “First Investor”.
3. Transaction Contract (TC): the contract signed when a Professional Investor acquires cash-flow rights, recording the investor's share of cash-flow rights or distribution priority, agreed returns, and rights and obligations.

III. “First Investors”: Contributions, Qualifications, and Profiles

The “First Investor” is the starting point of this market.

A “First Investor” typically has a core business, such as property ownership, chain franchising, equipment services, or content rights, and identifies worthwhile business opportunities within that business's industry

ecosystem. For example, a property owner assesses a tenant as having the ability to operate a business, funds the initial renovation and start-up costs, and receives a share of operating revenue over a specified period. Such contributions extend the “First Investor”'s core business. The “First Investor” draws on industry connections to identify opportunities and make judgments, and continues to follow up after making the contribution.

“First Investors” may have many business opportunities but limited available capital. Listing the contract cash-flow rights created by prior contributions and sharing future cash flows with Professional Investors can enable them to recover some capital for their next business investment. This is an important reason for “First Investors” to participate in the market.

In principle, “First Investors” retain an interest throughout the life of the contract. Retention arrangements, including the share of rights retained, related commitments, and subsequent changes, are disclosed to Professional Investors on the platform at listing and throughout the contract's life. See Section V(3) of this column. The retained share reflects the rights that the “First Investor” continues to hold and the alignment of interests with investors. The risks borne by each party also depend on the specific cash-flow distribution arrangement. The rights retained by the “First Investor”, the obligation to continue following up, and the obligation to provide information must each be clearly defined.

A “First Investor”'s status derives from actual prior contributions. To qualify, a “First Investor” must show that the contribution has already been made, that the Real World Contract specifies the corresponding cash-flow rights, and that the contribution and the basis for those cash-flow rights can be measured and verified. It is also necessary to confirm that the “First Investor” has the capability and standing to continue following up on the business activity. Joint contributions by multiple parties, related-party contracts, and similar cases are assessed individually according to the actual contributions and rights involved.

Contracts entering the market must be backed by actual prior contributions, be legally effective and be in the course of performance, and have a verifiable basis for the relevant cash flows. They must meet the platform's applicable admission standards in addition to legal and regulatory requirements and the rules of the relevant licensed trading venue.

If a contract has no historical performance record available for reference, or historical observations are insufficient, the relevant Compass scoring dimensions and disclosures should clearly state the gaps in the record, the evidence base, and the limits of judgment. Gaps must not be filled with speculation. This disclosure does not replace verification of the contract's authenticity, actual performance, and cash flows. Nor may the “First Investor”'s performance records for other contracts replace this contract's own history.

Prior contributions take many forms. To aid understanding, the following table groups them into four categories according to the main contribution made. These categories are not exhaustive and may overlap. Their main differences concern the value provided in advance and the corresponding contract cash-flow rights.

Contribution Category	What Was Contributed in Advance	Contractually Agreed Cash-Flow Rights	Typical Profile
Operating System	Brand, individual-store business model, systems, training, initial inventory, and similar contributions	Brand-use fees, management fees, and revenue shares paid by the underlying business operator as agreed	Chain brand owner
Capital-Intensive Assets	Property-use rights, equipment, facilities, renovation works, and similar assets or works provided in advance	Rent, usage fees, shares of operating revenue, or shares of energy-cost savings	Commercial property owners, computing-capacity service providers, and equipment lessors
Upfront Funding and Working Capital	Upfront costs, procurement payments, logistics, and inventory preparation	Contractual entitlements to box-office proceeds, sales receipts, or shares of project revenue	Live performance organizers and cross-border fulfillment service providers

Contribution Category	What Was Contributed in Advance	Contractually Agreed Cash-Flow Rights	Typical Profile
Intellectual Property	Licenses to use copyrights, patents, and other intellectual property, together with related technology, content, or design deliverables	Authorization fees, license fees, and agreed shares of sales	Content rights owners and patent holders

Table 1-1 Four Categories of Prior Contributions

IV. Why “First Investors” Participate: Many Opportunities, Limited Capital

A common question is: if everyone wants to keep good investments, might the investments they offer to others be precisely those they do not want? Answering this requires first understanding why “First Investors” participate in the market.

A “First Investor”'s core business is typically an industrial or commercial activity, such as owning property, operating a chain, manufacturing equipment, undertaking construction projects, or managing copyrights. Most of the “First Investor”'s capital is invested in that core business, and industry connections provide a continuing source of new business opportunities. When these opportunities require more capital than is currently available, the “First Investor” may wish to recover some capital for the next business investment, even if the existing contracts remain worth holding.

Transferring contract cash-flow rights offers one way to recycle capital. The “First Investor” contributes value first, then transfers part of the resulting contract cash-flow rights to Professional Investors and reinvests the funds received. For example, a chain brand owner provides support for setting up a franchise store and for training, then transfers part of the franchise contract's future cash flows. The proceeds can support the establishment of the next store. Capital recovered in this way may support successive rounds of business investment.

This also means that the market's potential size cannot be assessed solely by the number of contracts or participants it has today. It also depends on how many business operators have already created verifiable contract cash-flow rights and need to keep investing and recycling capital. Whether these contracts can enter the market also depends on contract quality, transfer terms, processing costs, and Professional Investors' willingness to participate.

“First Investors” may also consider borrowing. Access to loans and borrowing costs depend on factors such as their business condition, credit terms, and available security. Borrowing typically requires the borrower to repay principal and interest as agreed. Under the arrangements proposed in this issue, after contract cash-flow rights are transferred, investors' returns come from the corresponding contract's actual cash flows. The “First Investor” provides no guarantee and has no repurchase obligation. The legal characterization and accounting treatment of these arrangements depend on the final contract terms and applicable standards; no particular accounting treatment is guaranteed. “First Investors” choose the arrangement that suits their needs and the terms they can accept.

Capital recycling can explain why contract cash-flow rights are transferred, but it cannot establish contract quality on its own. “First Investors” typically know more about the contracts they hold, and their choices of which contracts to transfer and how much to retain may still create information asymmetries. The market needs to provide Professional Investors with a basis for judgment through verification of contributions, performance records, disclosure of retained interests, and ongoing measurement. These arrangements are discussed below under Measurement and Ongoing Disclosure.

V. Cash-Flow Arrangements, Contract Pools, and Retained Shares

The assets involved in Contract-Based Financing (CBF) differ in their cash-flow sources, payment frequency, duration, and risk characteristics. Cash-flow arrangements can be designed flexibly to reflect these characteristics and participants' needs. When bringing contracts to market, “First Investors” need to propose how cash flows will be distributed, whether to participate through individual contracts or a Contract Pool, and how they will retain an ongoing interest. Professional Investors decide whether to participate and on what terms. The arrangements are expressly agreed before the transaction, registered with NUMA, and disclosed to Professional Investors. Rights already established cannot be changed unilaterally by any party. Previously disclosed service fees and taxes payable by law are deducted from contract cash flows before distribution according to the agreed proportions or priority.

(1) Cash-Flow Arrangements for Individual Contracts

The following sections explain the basic logic of Pro Rata Cash-Flow Distribution and Priority-Based Cash-Flow Distribution. The two arrangements do not replace one another. The specific cash-flow scope, distribution proportions, sequence, and recovery conditions are governed by what is agreed, registered, and disclosed before the transaction.

	Pro Rata Cash-Flow Distribution	Priority-Based Cash-Flow Distribution
Basic Arrangement	Each period's distributable cash flows are distributed simultaneously to holders according to their shares of cash-flow rights, as agreed and registered in advance	Each period's distributable cash flows are allocated to holders earlier in the priority sequence according to the agreed priority allocation ratio; where a recovery cap is set, distributions proceed to subsequent positions once the agreed cap is reached
“First Investor”'s Position	Receives cash flows simultaneously with investors according to the agreed share	The retained portion comes last, as specified in the Register
Terms to Be Agreed	Cash-flow scope, shares of cash-flow rights and the basis for determining them, distribution frequency, whether a recovery cap applies, and termination conditions	Cash-flow scope, priority allocation ratio, priority sequence, whether each position has a recovery cap and its amount, and treatment of remaining cash flows

Table 1-2 Two Cash-Flow Distribution Arrangements

1. Pro Rata Cash-Flow Distribution

Pro Rata Cash-Flow Distribution carries forward the arrangements described in earlier white papers. Within the same distribution period, holders receive cash flows according to their agreed shares of cash-flow rights, without gaining priority because they bought earlier. Where different distribution priorities exist, this arrangement applies to holders within the same priority position. Shares of cash-flow rights and the basis for determining them are expressly agreed and registered before the transaction.

2. Priority-Based Cash-Flow Distribution

The August 3 white paper introduced time-based priority arrangements at the portfolio level. This issue further specifies the distribution rules for each position, the entry sequence for successive batches, transaction terms, and registration when priority arrangements are used. It also explains that subsequent transactions must not change registered holders' existing rights. The “First Investor” may transfer priority positions in the cash flows of an entire contract, and investors acquire priority distribution rights as agreed. Where a recovery cap is set, priority recovery continues until the amount invested and the agreed return have been recovered. The scope of cash flows covered by the priority arrangement and the proportion of each period's cash flows allocated to priority recovery are agreed separately. Where the entire contract's cash flows form the distribution base, all or a specified proportion of each period's distributable cash flows may be allocated to priority recovery. How cash flows outside the priority allocation are used, who receives them, and in what order must be specified in advance and must not impair registered holders' existing rights.

Cash flows allocated to priority distribution go first to the holder in the first position. Where a recovery cap is set, distributions proceed to the next position once that holder's agreed cap is reached. For positions without a recovery cap, distribution and exit conditions are separately agreed and registered before the transaction. Holders in subsequent positions receive the remaining cash flows as agreed and bear the corresponding risk of insufficient cash flows. The “First Investor”'s retained portion comes last, as specified in the Register. Returns at every position come from the contract's actual cash flows. Full recovery of the amount invested and the agreed return is not guaranteed.

3. Distinguishing Proportions from Priority

The following concepts are registered separately and must not be used interchangeably:

1) Capital contribution ratio: the relative proportions of capital contributed. Its calculation basis and denominator follow the relevant contractual agreement and registered definition.
2) Purchase size: the purchase amount as a proportion of Contract Value (CV). It describes the size of an investor's purchase and is not directly equivalent to the capital contribution ratio or the share of cash-flow rights.
3) Share of cash-flow rights: the holder's share of distribution rights within the agreed cash-flow scope, determined under the Protocol Contract and Transaction Contract and registered accordingly.
4) Priority allocation ratio: the portion of each period's distributable cash flows allocated to priority recovery.
5) Cash-flow distribution priority: the sequence in which positions receive distributions.

Cash flows are distributed according to agreed and registered shares of cash-flow rights or distribution priority. These are not inferred directly from the proportions of capital actually invested.

4. Whether a Recovery Cap Applies

How cash flows are distributed and whether a recovery cap applies are two separate matters. Under a capped arrangement, the cash-flow rights of a position terminate once the agreed amount has been paid in full. Under an uncapped arrangement, participation in distributions continues under the relevant contract until the agreed termination conditions arise. Distribution, continued participation, and exit conditions are agreed and registered separately for the applicable circumstances. The description of rights terminating once an agreed amount is reached applies only to arrangements with a recovery cap.

5. Entry in Batches and Determination of Terms

Under Priority-Based Cash-Flow Distribution, Professional Investors enter in successive batches, in the order in which they purchase, with their priority positions registered in sequence. Each batch's price, agreed return, recovery cap, if any, and other terms are determined and registered when the transaction occurs.

Within the scope permitted by the listing arrangement, the “First Investor” may adjust the terms of the portion not yet transferred. Adjustments apply only to investors entering afterward and must not change registered holders' priority, cash-flow scope, or other existing rights. Holders entering the same contract in different batches may therefore be subject to different terms.

A later investor seeking an earlier position may acquire a registered holder's entire position under rules disclosed and agreed in advance. The investor assumes that position's existing priority without changing other holders' positions. Transfer rules are specified before the initial purchase. Rules added or amended afterward do not apply retrospectively to registered positions. Transfers must meet the conditions specified in advance; later entrants do not have an unconditional right to obtain an earlier position.

6. Transferring One's Own Priority Position

An investor holding a residual position may carve out a portion with earlier priority within its own holding and transfer it to an investor who prefers earlier cash flows, while retaining the residual cash flows that follow. The

basic principle is that holders may subdivide and transfer their holdings only within the scope of their existing rights, without changing other holders' existing priorities or rights. The party transferring a position provides no guarantee and has no repurchase obligation.

All positions and changes to them are governed by the Register. Each holder can view the contract's complete priority schedule, including its own position, positions ahead of it and their distribution conditions and recovery caps, if any, and the “First Investor”'s remaining holding.

(2) Contract Pools: A Unified Arrangement for Multiple Contracts

The August 3 white paper allowed one Protocol Contract to correspond to one or more Real World Contracts and described the transparency requirements for bundling. This issue further specifies the fixed scope, unified distribution, priority, and release conditions of Contract Pools. A Contract Pool here is an asset-side arrangement through which a “First Investor” brings a batch of contracts into a unified cash-flow distribution arrangement. It differs from an investor's portfolio of multiple rights: that portfolio is a collection of holdings for monitoring purposes and does not alter the original transaction terms of each contract.

Some “First Investors” start with a single contract and gradually expand their participation as performance records accumulate. “First Investors” holding many similar contracts, such as certain platform businesses, may bring multiple contracts into the market under a unified arrangement. The specific arrangements include the following:

1. Contracts included in the pool: the “First Investor” selects a batch of contracts meeting standards disclosed in advance. All contracts are identified and locked in when included. Each contract must still undergo individual qualification, registration, and admission checks. Contracts in the pool are not subsequently replaced or added, and the batch's cash flows are not used to invest in new contracts.
2. Unified distribution: cash flows from the batch are distributed under the proportions, priority, and recovery caps, if any, determined at inclusion. The “First Investor”'s retained portion comes last. Investors acquire agreed positions in the aggregate cash flows of the batch. These contracts may share risks relating to industry, geography, underlying business operators, or operating capabilities. Such concentrations must be measured and disclosed together.
3. Distribution by priority: cash flows from the batch are distributed according to registered priority. Holders in earlier positions have agreed priority distribution rights. Holders in later positions bear the impact of cash-flow shortfalls first, as agreed. Every position is exposed to the batch's aggregate cash flows, and actual recovery depends on its overall performance. Agreed returns are not guaranteed.
4. Release of contracts: if an investor wishes to purchase certain contracts in the batch in their entirety, the purchase price must first be distributed to holders of the batch according to the agreed priority. The contracts may be released only after the predetermined release conditions are met. Release conditions and pricing rules are determined and disclosed when contracts are included in the pool.

Contract Pools are intended to enable larger “First Investors” and Professional Investors to participate in a manner commensurate with their scale. The former bring a batch of contracts into a unified arrangement, and the latter acquire agreed rights in the batch's aggregate cash flows. Contract Pools do not change the progressive listing principle in the August 3 white paper or create an exception to the number of contracts that may be listed. Individual contract verification, disclosure of concentration risks, and statements of validation status continue to apply. Contract Pools remain at the design stage and are intended to be tested in the Simulation Testbed. Any actual launch is subject to the required approvals and completion of necessary procedures, and participation is limited to qualified Professional Investors.

(3) Retaining an Ongoing Interest and Special Arrangements

In principle, “First Investors” should retain an interest throughout the contract's life to maintain alignment with investors. Retention arrangements, including the share of rights retained, related commitments, and subsequent changes, are disclosed to Professional Investors at listing and throughout the contract's life. Retention commitments already made must not be unilaterally reduced or released in respect of investors who have already purchased rights. Where genuinely special circumstances arise, the applicable conditions, disclosure, and continuing responsibilities will be discussed separately.

The retained share and changes to it are matters for disclosure and measurement. Whatever the share retained, obligations to provide information, follow up on operations, and handle cash flows throughout the contract's life must be expressly agreed. Transferring rights does not automatically discharge these obligations.

VI. Measurement and Ongoing Disclosure

“First Investors” typically know more about the individual contracts they hold. Subject to admission and disclosure requirements, they choose which contracts to list, how much to transfer, and how much to retain. Through measurement and disclosure, the market makes these information differences visible and provides Professional Investors with a basis for judgment.

(1) Measurement and Records of Issuer Performance

The “First Investor” in this article is the Protocol Contract Issuer. The three earlier white papers progressively developed arrangements for issuer profiles, performance records, and records accumulated across contracts and periods. When the Core Compass and Industry Compass measure contracts, their scores include relevant aspects of issuer performance. Building on this, each listed contract is linked to its actual cash flows and performance records, the issuer's retention arrangements, concentrations in the issuer's holdings, and the applicable rules. They are recorded continuously and disclosed to Professional Investors. These records help investors assess historical performance and transfer choices, but cannot guarantee that every problem will be identified.

(2) Accurate Disclosure of Retained Shares and Selection Methods

For each listed contract, the “First Investor” discloses to Professional Investors the retained share, related commitments, the method of selecting the retained portion, and subsequent changes.

(3) Distinguishing Rights and Risks by Cash-Flow Arrangement

Pro Rata Cash-Flow Distribution distributes cash flows simultaneously according to agreed shares of cash-flow rights. Priority-Based Cash-Flow Distribution allocates cash flows according to the agreed priority allocation ratio and sequence, with the “First Investor”'s retained portion coming last as specified in the Register. Each party's cash-flow timing, amounts, and risks depend on the rights it holds and the specific distribution arrangement.

Professional Investors independently decide whether to participate, based on Compass readings, contract terms, and disclosed information.

The main information requiring ongoing disclosure includes what the “First Investor” contributed in advance, when the contribution was made, and the supporting evidence. It also includes the Real World Contract's terms and cash-flow path. Disclosures cover the chosen cash-flow arrangement, whether a recovery cap applies, whether individual contracts or a Contract Pool are used, and the retained share and changes to it. They also cover past contract performance, gaps in this contract's historical performance record and the evidence base, and concentrations by industry, geography, and underlying business operator among the contracts held. “First Investors” provide this information when entering the market and update it throughout the contract's life.

Measurement of the contract itself follows the NUMA SHARP Twin Compasses method introduced in the September 3 white paper and refined in the October 3 white paper. The Core Compass examines factors common to all contracts, and the Industry Compass examines factors specific to each industry. Returns and risks are assessed separately. Where evidence is insufficient, gaps are left blank rather than filled with speculation.

Compass readings and disclosed information are tools to help investors form judgments. They are not recommendations or opinions on any party's credit standing or solvency. This information is provided only to Professional Investors admitted through licensed institutions and is not made public. Underlying materials and personal information are not part of general disclosures. Any provision of information needed for verification in accordance with law, business processing, or the relevant parties' voluntary choices is subject to applicable rules and valid authorization and is limited to what is necessary. The Compasses use only materials voluntarily submitted by the “First Investor” and verified, together with its contract performance records within this market. They do not collect information from credit reporting agencies in China.

VII. Transactions Within the Market, Connections Outside It

(1) The Problem

For a market to operate, “First Investors” and Professional Investors must both be present. This is difficult at the outset. “First Investors” will come only if they believe participation will benefit them. Professional Investors will not spend time on a market with no contracts yet listed. Each side's participation depends on the other, so both wait. The market itself has limited staff to reach both sides. Off-market intermediaries often know one side well: some are close to “First Investors”, and others to Professional Investors. Relatively few know both sides well.

(2) The Arrangements

NUMA's answer is: transactions within the market, connections outside it.

1. All transactions take place within the market. Listing, subscription, allocation, and transaction execution are handled by licensed trading venues and their members. Funds settlement takes place through licensed institutions and custodian banks. NUMA's Transaction Matching System handles “First Investor” qualification, contract registration, measurement, and matching support such as preparing materials, transmitting instructions, and reconciling results.
2. Off-market intermediaries bring the two sides closer. Financing-side intermediaries familiar with “First Investors” encourage them to prepare materials, attend interviews, and submit contracts, without advising them on listing terms, prices, or returns. Investor-side intermediaries encourage Professional Investors to open accounts at licensed institutions, complete identity checks, and learn about the market. Off-market content and activities requiring a license are the responsibility of the relevant licensed institutions.
3. Central support and intermediary collaboration. Through its Information Matching System, NUMA continuously produces content, coordinates its direction and publication timing, provides support, and organizes outreach. Intermediaries on both sides connect with clients, explain the market, and help them enter through the prescribed channels. Each has its own role. The Matching Hub, operated by NUMA within the Information Matching System, gives both sides visibility at the category level and defines what may and may not be said outside the market. Every intermediary's actions are controlled and recorded within the system. Together, their work helps ensure that both sides are ready when transactions enter the market. The algorithm of the Matching Hub, the Information Matching Engine, is planned for introduction around December 3, 2026.

Outside the market, there are connections but no transactions. Category information displayed outside the market does not disclose specific issuers, contract terms, prices, or expected returns. Information about

specific parties and contracts that intermediaries access when helping prepare materials, arrange interviews, and submit contracts must be handled within the authorized scope and confidentiality requirements. It must not be disclosed to other participants without authorization. Off-market intermediaries do not collect or receive any payments on behalf of others. Interest registration and other connection activities create no subscription allocation or priority.

(3) Significance

Off-market connections allow “First Investors” to bring contracts and materials into the market with the help of an intermediary familiar with their side, without having to find investors themselves. Market development relies on continuous central content and system support, with intermediaries on each side serving their clients; the central team need not handle every client individually. In the market's early stages, NUMA is willing to allocate a portion of platform revenue to support and incentivize participation. The specific incentives and rules remain to be announced. Separately, the platform provides recordkeeping and transfer support for remuneration allocation under agreements already reached by market participants. Fund transfers are executed by the relevant licensed institutions under valid authorization. All transactions remain concentrated within the market and are subject to the same qualification, registration, measurement, and disclosure rules. New arrangements for off-market connections will be introduced after the relevant approvals have been obtained and necessary procedures completed.

VIII. Design Principles and Business Boundaries

“First Investors” enter the market under four basic principles:

1. Prior contribution before purchase: listed contracts must arise from value already contributed by the “First Investor”. This is the primary eligibility condition for a contract to enter the market and is checked through the “First Investor” qualification process. The contract must also be legally effective, actively being performed, and compliant with legal and regulatory requirements and the platform's admission standards.
2. Returns are limited to contract cash flows: holders receive returns only from the contract cash flows corresponding to their rights. Parties transferring rights, including “First Investors”, provide no guarantee and have no repurchase obligation. The absence of a guarantee or repurchase obligation does not relieve transferors of liability for fraud, defective title, diversion of cash flows, or breach of retention commitments.
3. Transfers are limited to one's own rights: holders may transfer only the cash-flow rights they own.
4. The Register governs: ownership of rights, shares of cash-flow rights or distribution priority, and changes to them are governed by the Register.

The market itself also has clear boundaries. As market infrastructure, NUMA provides qualification, registration, measurement, disclosure, and matching support. Participation on the investment side is limited to Professional Investors admitted through licensed institutions. The funding side is restricted to outside China: there are no sales to residents of mainland China and no acceptance of funds originating from mainland China. The relevant licensed institutions verify investor identity, qualifications, and sources of funds in accordance with applicable rules and within their respective responsibilities. The platform manages access and permissions based on those verification results and its own admission rules. Client funds are held only at licensed custodian banks. Transactions take place only within the market. The market does not provide loans, set contract prices, or make investment decisions for investors. Its exclusion list covers sectors on the Negative List for Foreign Investment Access and sectors where foreign investment is prohibited or restricted in China.

Part II | Market Progress

NUMA Operating System OS: The Three Systems

Reading Guide

Earlier white papers established NUMA's contract structure, measurement methods, and organizational model in which artificial intelligence (AI) employees execute tasks and humans bear responsibility. Building on the HAT Over MASS approach to human–AI collaboration introduced in the October 3 white paper, this part presents the draft framework for the NUMA Operating System. Centered on internal operations and connected to market participants and licensed institutions, the proposed system organizes business activities into tasks with clear boundaries that can be executed and verified. AI employees execute within their authorized scope. Named humans are responsible for rules, authorization, necessary business judgments and decisions, supervision, and exception handling. In this way, the system translates market rules into specific processes and responsibilities.

From submission of materials through ongoing monitoring and exit, a contract passes through material checks, registration, measurement, transaction coordination, cash-flow distribution, and ongoing disclosure. Under this design, OS organizes this work into connected tasks. It specifies the rules governing each step, the required output, who executes it, and who is responsible, and retains records for reconciliation and traceability. It must support efficient processing and provide paths for handling exceptions, a factual basis for judgments, and named people accountable for the work.

This part covers the Responsibility Matching System, Transaction Matching System, and Information Matching System. The Responsibility Matching System assigns supervision and decisions to named humans. The Transaction Matching System supports business tasks and AI employees. The Information Matching System organizes communication and connections for market entry. Together, they explain what the system comprises, what each component is responsible for, and how work and responsibility connect. The framework in this issue still requires simulation testing and revision. Further operating methods and a report on simulation progress are planned for around December 3. Drawing on both issues and the test results, a fourth white paper is planned for around January 3, 2027.

I. Design Objectives and Overall Architecture

This section explains the problems OS is intended to address, its components, and the current status of the design.

(1) Managing Business Processing and Internal Support as Tasks

OS organizes financing-side services, investor-side services, and market operations support through tasks. Each task specifies its inputs, outputs, applicable rules, authorized scope, and responsible person, with execution and verification results recorded. Internal support, including compliance, legal, finance, risk management, and information technology (IT), is also managed through tasks. Advice and results are recorded with the corresponding business activity, preserving a complete processing record.

(2) Balancing Efficiency, Safety, Compliance, and Traceability

The main efficiency metric for OS is whether lifecycle processing cost per contract meeting completion criteria falls from one batch to the next. Processing time, rework, quality, safety, and compliance are also monitored. Cost statistics include resources used for normal processing, items returned for correction, failures, rework, and exception handling. Costs must not be omitted because a transaction was ultimately not executed or processing was not completed. A contract meets the completion criteria only when all five conditions hold: business processing is complete, evidence is complete, controls are effective, accounting records and holdings reconcile, and relevant issues are closed. Standards must not be relaxed to reduce costs, and meeting these

criteria does not guarantee investment quality or returns. AI employee numbers and call counts are not success metrics.

Efficiency assessment must consider both whether a step becomes faster and whether it increases subsequent verification and resolution costs. Safety controls cover prevention of duplicate funds instructions, reconciliation of differences, and exception handling. Business processing must comply with applicable rules, entity qualifications, and authorized scope. Original evidence, rule versions, execution processes, and accountability records are fully linked to support verification and traceability. Operational effectiveness is assessed through task completion quality and check results. Specific metric definitions and thresholds will be progressively calibrated through simulation and subsequent operations.

(3) Supporting Modular Integration and Safe Expansion

When adding industries, data sources, rules, or AI employees, existing interfaces, task standards, and control mechanisms are reused wherever possible. Inputs, outputs, permissions, responsibilities, and rollback plans are specified to reduce duplication. Integration validation, change approval, and rollback requirements are described in Section VI.

(4) The Three Systems

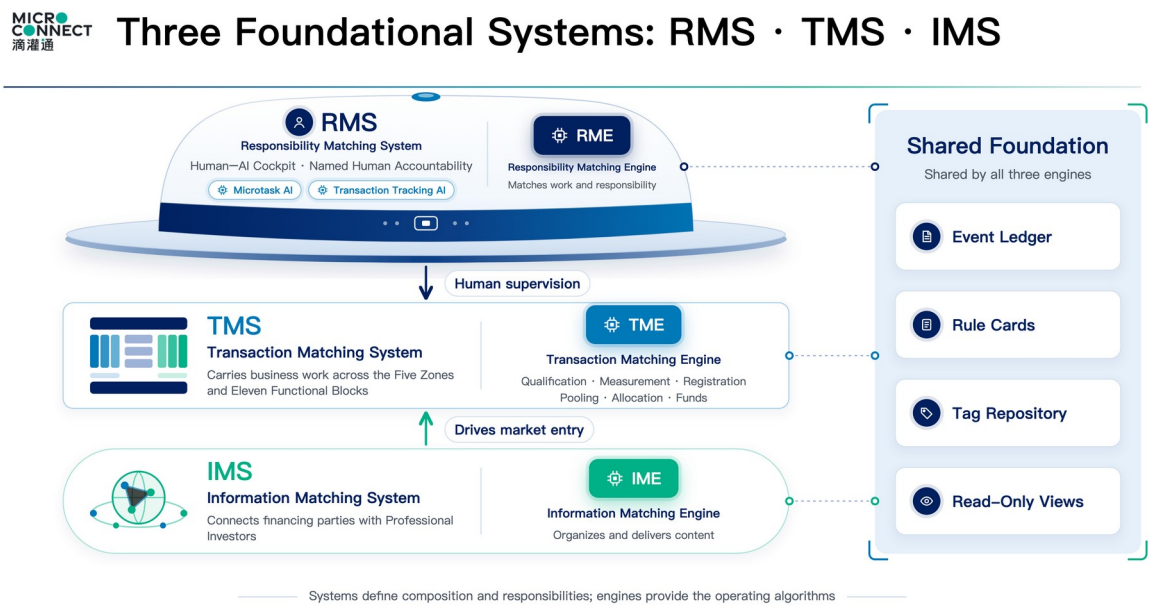


Figure 2-1 The Three Systems: Responsibility Matching, Transaction Matching, and Information Matching

OS comprises three systems. The Responsibility Matching System (RMS) uses the Human–AI Cockpit as its organizational unit, assigning task supervision, authorization, and decisions to named employees on duty. The Transaction Matching System (TMS) supports business and support work through functional blocks, main tasks, and microtasks, executed by AI employees under approved rules and within their authorized scope. The Information Matching System (IMS) organizes content support and connections on both sides, helping financing-side “First Investors” and Professional Investors enter the market through prescribed channels. Here, a “First Investor” is a party that has already contributed resources, capabilities, or capital to a business activity and brings the resulting contract cash-flow rights to market. Off-market intermediaries do not execute transactions. Content and activities requiring a license are handled by the relevant licensed institutions under their rules.

The three engines operating on these systems are planned for introduction around December 3, 2026: the Responsibility Matching Engine (RME), Transaction Matching Engine (TME), and Information Matching Engine

(IME), collectively called the matching engines. All three follow the same principle: algorithms propose, rule cards check, and humans sign off. Subject to actual readiness, we plan to include the relevant engines and the three systems together in internal simulation testing.

The extensions shown in the figure will be developed in later discussions. This issue focuses on the composition, duties, and accountability of the three systems.

The chip in the figure represents work supported by the Transaction Matching System, broken down internally into the Functional, Main Task, and Microtask Layers. Humans supervise AI through Human–AI Cockpits in the Responsibility Matching System. When employees take leave, leave their posts, or return, the cockpits are reconfigured, and the work and records on the chip are retained. The three systems describe the overall composition and responsibilities. The three layers inside the chip describe how work is broken down. The names of these levels must not be used interchangeably.

(5) Design Scope and Version Status

The initial version as a whole has not yet been validated through operation, nor does it mean that every module is live. The arrangements described reflect current design intentions and do not constitute a commitment regarding implementation progress, continuous operating capability, or business results. As development proceeds, each component will be labeled according to its actual status: proposed, configured, deployed, or validated. Interfaces, processes, and accountability arrangements will continue to be refined through configuration and testing. This version statement applies to all content in this part. The specific conditions for matters such as 7×24-hour continuous operation are described in the relevant subsections.

II. Responsibility Matching System (RMS)

This section explains who is responsible and how Human–AI Cockpits implement authorization, supervision, decisions, and personnel handovers.

(1) Named Human Accountability and the Division of Work Between Humans and AI

Human governance specifies who makes rules, how authority is granted, who decides matters, and how responsibility can be traced. AI executes within its authorized scope, and named humans undertake the corresponding governance, judgment, and supervision responsibilities.

HAT (Human-Commanded, AI-Operated, Trust) is a trust mechanism in which humans command and AI operates. Humans are responsible for making, amending, and repealing rules, as well as authorization, supervision, and exception handling. MASS (Micro Agents · Simple & Short Tasks) organizes AI employees to execute microtasks. Together, they form the HAT Over MASS approach to human–AI collaboration. The NUMA Operating System OS puts this approach into practice. This part explains its composition, duties, and accountability arrangements through the three systems.

Each AI position is assigned to a named human, who takes responsibility for it. When a matter falls into a gap in the rules, exceeds authorization, or requires discretion, the AI suspends the relevant execution and submits the facts, supporting grounds, and questions requiring a decision to an authorized human.

Human positions are divided into decision-making, supervisory, and specialist tiers. Within their respective authorized scopes, they undertake responsibilities such as rule management, organizational supervision, specific processing, and review.

Each call separately records the following responsibilities:

Person responsible for the position: responsible for the AI position's rules, permission boundaries, and supervision under the cockpit's command relationships at that time.

Actual handler: processes the task within the authorized scope and is responsible for the specific output.

Approver: makes valid decisions on matters requiring approval, release, or sign-off.

These responsibilities are recorded separately. Taking responsibility for a position does not replace specific processing and approval responsibilities.

(2) One Cockpit per Person and Two Types of Supervisory AI

A Human–AI Cockpit consists of a named employee on duty and the assigned supervisory AI. It implements task supervision and human accountability and is the organizational form of the Responsibility Matching System. Each employee on duty has one cockpit and is responsible for it. A cockpit does not continue to exist without a human on duty. The number of cockpits depends on actual staffing on duty. The design range is 50–100, with an initial configuration planned around the approximately 80 employees currently on duty. The actual number opened is governed by the number of employees on duty; this does not establish that all planned cockpits have already been opened or are operating.

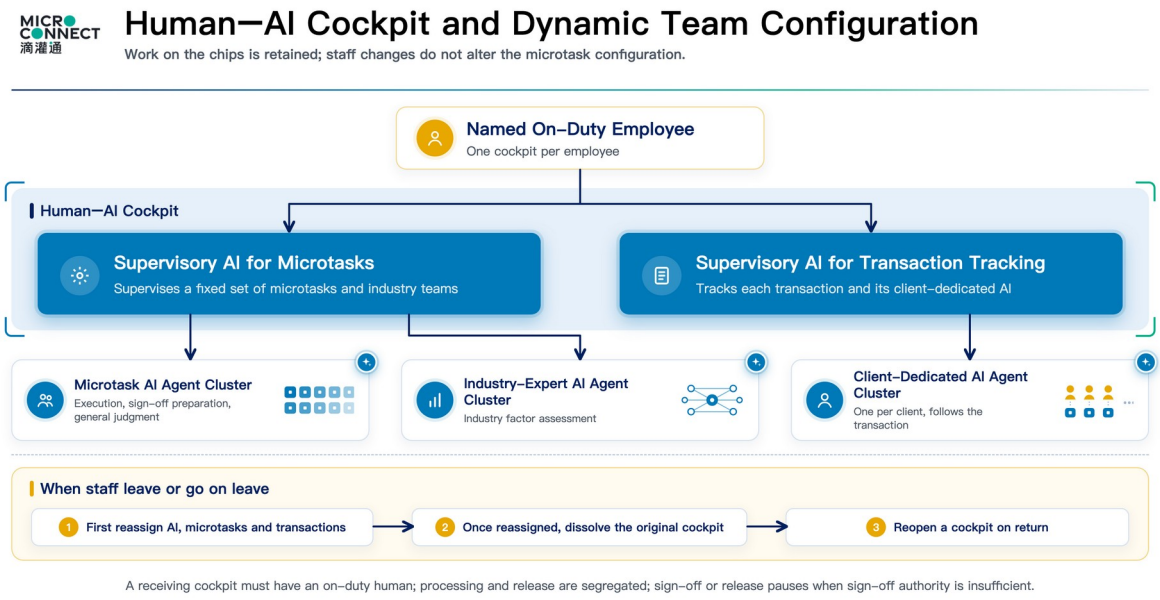


Figure 2-2 Human–AI Cockpits and Dynamic Team Configuration

Each cockpit is assigned two types of AI. Supervisory AI for Microtasks manages a fixed group of microtasks, usually 10–20, tracking outputs, queues, returns, and errors and forwarding matters requiring a human decision to the responsible person. Supervisory AI for Transaction Tracking continuously tracks each transaction and its client-dedicated AI until the transaction is completed or terminated. One such tracking AI may track multiple client-dedicated AI agents at the same time. Industry-expert teams are also assigned to cockpits, supervised by employees in the relevant fields.

Two roles sharing the name microtask AI must be distinguished here. The microtask AI in a cockpit supervises. The Microtask AI Agent Cluster in the Transaction Matching System carries out execution, sign-off preparation, and general judgment. Every AI employee in the Transaction Matching System reports to a cockpit AI, and every cockpit AI reports to a named human. The Event Ledger records command relationships at every point in time, so responsibility can be traced to the AI and human involved at that time.

(3) Sign-Off Methods, Authorization, and Task Classification

OS provides three sign-off methods according to the subject of the decision:

Sign-Off Method	Applicable Matters
Sign-Off by Task Category	Approve rules, standard terms, and methods as the basis for executing tasks in the same category
Case-by-Case Sign-Off	Approve business requiring individual decisions using the materials for that case
Exception-Only Sign-Off	Make a specific decision on matters not covered by the rules or departing from the normal path

Table 2-1 Sign-Off Methods and Applicable Matters

Tasks that have been approved, meet the conditions, and remain within the authorized scope may proceed under the rules without repeated approval of every individual action. General authorization does not waive applicable case-by-case, dual, or statutory sign-off requirements.

Authorization records specify the scope of matters, applicable monetary limits, validity periods, and reassignment conditions. Funds transfers, contract signing, external commitments, and matters requiring professional confirmation are decided by parties with the appropriate authority and qualifications.

Employee qualification records specify the task tiers each employee may undertake in each professional field, providing a basis for team configuration and authorization. Internal task tiers do not replace required professional qualifications or approval authority.

The Responsibility Matching Engine calculates task classifications and team configurations daily. They take effect after confirmation by Central Control staff on duty. The calculation considers the task's professional field, risk level, inputs, outputs, and verification requirements, together with employee qualifications, business authorization, and arrangements for segregation of duties. Tasks are currently divided into four tiers:

Tier	Characteristics and Assignment Requirements
Tier 1	Rules are clear, AI has produced a draft, results are easy to verify, and errors can be reversed; handled by employees with basic task capabilities
Tier 2	Handled by employees with relevant experience and subject to subsequent spot checks
Tier 3	Involves professional judgment; handled by qualified personnel in the relevant field
Tier 4	Involves major matters; handled by responsible managers or designated experts with the appropriate authority and qualifications

Table 2-2 Employee Qualification Records and Task Tiers

The draft proposes that new employees and interns start with Tier 1 matters. A proportion of their decisions is spot-checked by employees at higher tiers. They may progress through the tiers when rates of work returned for correction and decisions overturned remain consistently low. A single high-cost error triggers an immediate downgrade. Task classification does not replace required professional qualifications or approval authority.

(4) Segregation of Duties and Sign-Off Validity

The handler is responsible for specific processing and outputs. The reviewer verifies evidence, the application of rules, and processing results. The signatory is responsible for approval, release, or sign-off within the authorized scope. Each retains records of the work performed.

Funds decisions and changes to payment or receipt accounts require sign-off by two humans who must not be in a direct reporting relationship. Payments with valid dual approval may be executed within the effective authorized scope. Those processing a case do not release it; those approving do not issue instructions; those calculating do not reconcile accounts; and those producing conclusions do not perform acceptance checks. Sign-off must meet qualification and authorization requirements, as well as the segregation-of-duties requirements for the same contract and its transactions. Once authorization expires, it may no longer be used to process or sign off on the relevant matters.

The effectiveness of sign-off is tested against errors, returns, processing times, and test results, with a focus on whether it identifies problems and produces evidence-based decisions.

Sign-off microtasks with no returns over an extended period are referred for audit review. The system periodically introduces test cases to check whether sign-off can catch errors. Segregation-of-duties and signatory qualification requirements during personnel changes also apply to the Dynamic Team Configuration described in this section.

(5) Dynamic Team Configuration and Continuity of Accountability

The Responsibility Matching Engine calculates the assignment of microtasks, transactions, and industry teams daily. Assignments take effect after confirmation by Central Control staff on duty. The calculation considers employees' professional fields, experience, on-duty status, task loads, and past performance. The engine observes information barriers and conflict-of-interest rules. Personnel changes, workload changes, or the need for specialist coverage may trigger reconfiguration. Specific workload calculations and assignment limits must be determined through detailed design and simulation verification. When an employee takes a break, takes leave, or leaves their post, the engine recalculates under preset rules. After confirmation by Central Control staff on duty, the employee's AI, microtasks, and transactions are reassigned to other cockpits with humans on duty. The original cockpit is dissolved as reassignment is completed. Its Supervisory AI for Microtasks and Supervisory AI for Transaction Tracking are discontinued with it. The AI and transactions under them come under the corresponding AI in the receiving cockpits.

If an employee leaves unexpectedly without going offline in advance, the Responsibility Matching Engine recalculates under preset rules. After confirmation by Central Control staff on duty, reassignment follows the preset default destinations. When the employee returns, a cockpit is opened again and the relevant AI is assigned to it. Default or previous-day configurations may carry forward only the portions that remain valid. Departed employees do not remain responsible merely because an old configuration is still in the system. Dissolving or reconfiguring cockpits does not change the number of microtasks or AI agents on standing task assignments in the Transaction Matching System. Work and records remain in the system.

As business grows, task configurations may be adjusted according to actual work needs and the number, qualifications, and authorization of employees on duty. Additional cockpits should correspond to an increase in employees on duty. Each employee on duty remains responsible for only one cockpit; business growth does not automatically expand that person's authority.

Temporary additional work may be assigned to existing cockpits through temporary task coverage arrangements. These specify the task scope, person taking over, required qualifications and authorization, coverage period, and handover arrangements at its end. Temporary coverage does not automatically create a new cockpit or confer operational, approval, or sign-off authority that the person did not already possess. The reasons, applicable scope, and accountability records for these adjustments are retained.

Dynamic Team Configuration must observe two constraints. First, processing and release for the same contract or any transaction under it must not fall within the same employee's cockpit, including cockpits that employee opens at different times. Second, sign-off preparation microtasks are assigned only to cockpits of appropriately qualified employees. A matter requiring a licensed person or designated responsible person to sign must not be transferred to an unqualified person because of reconfiguration.

If both constraints cannot be met at a given time, the relevant AI is still assigned to a cockpit with a human on duty and remains supervised. The corresponding sign-off or release is suspended until a qualified employee independent of the processing is on duty to perform it. Continuity of responsibility and sign-off authority are handled separately, ensuring that AI is never left without a responsible human and that sign-off conditions are not relaxed.

Handovers use a uniform format. Handover records identify the business object, objective, evidence, rule basis, work completed, unresolved issues, and receiving responsible person. Records of acceptance, return, requests for additional materials, and timeout handling are retained.

After personnel or AI employee changes, the relevant records remain with the task, enabling the person taking over to continue processing and reviewing the work.

(6) Triage, Consultation, and Escalation

Central Control uses a unified triage mechanism to queue, group, and prioritize matters according to impact, urgency, and authority requirements, then sends matters requiring human handling to the relevant cockpits. Related matters may be presented together, but each task's object, responsibility, and required decision are recorded separately.

Matters involving funds, external commitments, statutory deadlines, or loss containment are sent directly. Recipients may return matters with incomplete materials, specifying whether what is missing is facts, methods, a decision, or a response. Matters not handled within the prescribed time are escalated under the rules. Repeated exceptions of the same type automatically generate an upstream rule-revision ticket.

For matters involving multiple positions, each must respond within a time limit with one of three responses: comments, no comments, or escalation required. Named comments and their supporting grounds are recorded separately. Failure to respond is not treated as no objection. Where views differ, the differing views are retained and an authorized party makes the decision.

Decision materials set out the objective, facts, constraints, options, and impacts. Once a decision is made, the decision-maker, materials version, applicable scope, and conclusion are recorded. A material change to the materials requires a fresh check of whether the original decision remains valid.

(7) Protecting Human Cognitive and Judgment Capabilities

Decision-makers and reviewers must understand the key facts, applicable rules, and grounds for conclusions, and be able to explain their judgments. An independent traceability AI explains the background and development of a matter to the human, organizing tasks, evidence, and differing views. It retains uncertainties and key limitations for human examination.

Mandatory questions are set periodically for each position to test judgment. Training, supervision, and authorization adjustments are arranged in light of actual processing and review performance, to prevent key judgment capabilities from deteriorating through prolonged reliance on AI.

(8) Dashboards, Correction Categories, and Configuration Adjustments

Dashboards prioritize objectives, outputs, responsible persons, and applicable rules. They also allow users to expand task processes and evidence to assess progress and investigate problems.

Corrections fall into four categories: correcting facts to address errors in materials or records; improving methods to adjust specific practices; adjusting authorization to change the party handling a matter or the scope of authority; and proposing rule revisions to change generally applicable requirements. Each category follows the corresponding authority and procedures, with its basis and results retained.

Only changes to generally applicable requirements enter the rule-revision process. Microtask configurations, cockpit team configurations, and permission configurations are version-controlled and implemented after approval through the relevant procedures. Tasks retain the versions used for replay and verification.

The initial configuration is planned around opening cockpits for employees actually on duty. A cockpit-to-microtask assignment table will identify the microtasks, transactions, and industry teams for which each is responsible, record qualification tiers, mark sign-off tasks and signatory qualifications, and preregister default

reassignment destinations for when employees leave their posts. Configurations will subsequently be adjusted according to task loads, handover efficiency, and specialist coverage.

III. Transaction Matching System (TMS)

This section explains how the Five Zones and Eleven Functional Blocks support business activities, and how main tasks, microtasks, and AI employees divide responsibilities and work together.

(1) Business Layout of the Five Zones and Eleven Functional Blocks

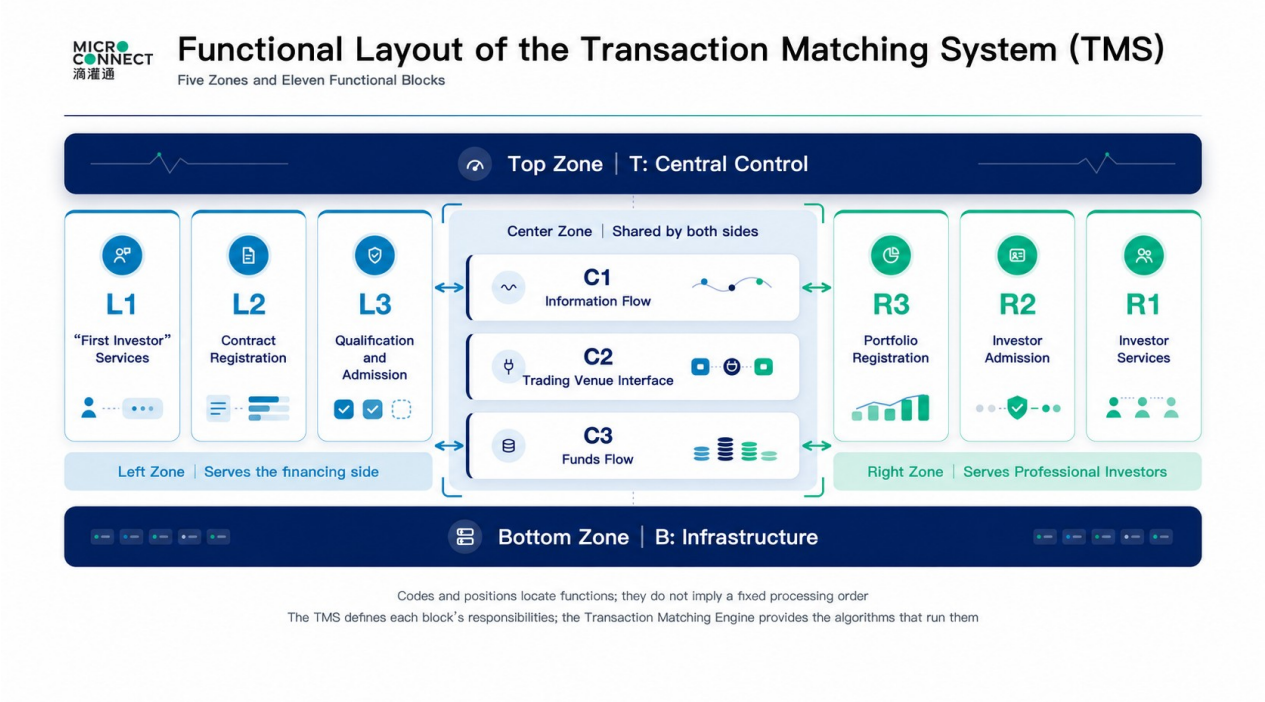


Figure 2-3 Transaction Matching System: Five Zones and Eleven Functional Blocks, and the Breakdown of Work

The Transaction Matching System divides business and support functions among five zones - top, left, right, center, and bottom - with eleven functional blocks in total. Central Control sits at the top, the financing side on the left, the investment side on the right, information, trading interfaces, and funds activities in the center, and operational support at the bottom.

Zone	Functional Block	Main Tasks
Top	T: Central Control	Rulemaking, central coordination and scheduling, backtesting, consultation, and decisions
Left	L1: Financing Party Services ("First Investor" Services)	Client reception and identification, dedicated assistants, and post-investment follow-up
Left	L2: Contract Registration	Collection of contract elements, three-way verification, listing documents, and entry in the Register
Left	L3: Qualification and Admission	"First Investor" qualification, contract admission, and shares available for sale
Right	R1: Investor Services	Client reception and education, dedicated assistants, and notifications
Right	R2: Investor Admission	Onboarding through licensed institutions, qualification checks, access activation, and limits
Right	R3: Portfolio Registration	Orders, holdings and distribution priorities, snapshots, and statements
Center	C1: Information Flow	Compass readings, updates based on evidence, scanning, and disclosure
Center	C2: Trading Venue Interface	Listing coordination, orders and allocation, and trade confirmations

Zone	Functional Block	Main Tasks
Center	C3: Funds Flow	Settlement, reconciliation, and cash-flow distributions
Bottom	B: Infrastructure	Risk control and compliance, IT systems, finance, legal affairs, and audit

Table 2-3 Five Zones and Eleven Functional Blocks and Their Main Tasks

Functional block codes identify locations; they do not indicate processing order. C1, C2, and C3 record information, transaction, and funds results, respectively, and cross-check them against contract, order, and holding records. Activities requiring a license are executed by the relevant licensed institutions under their rules.

(2) Functional, Main Task, and Microtask Layers

The functional layer defines the responsibilities of each of the eleven blocks. The main task layer divides those responsibilities into 34 main activities. The microtask layer is intended to break the main tasks down into 1,530 work units that can be executed and verified independently. The functional and main task layers organize work and have no AI agents assigned to them; standing AI assignments begin at the microtask layer. Each microtask is a fixed work position that performs just one activity, with its assigned AI agents responsible for that work on an ongoing basis.

For example, "reconciliation" is a main task within C3: Funds Flow, and "determining whether differences are zero" is one of its microtasks. Work is identified successively by functional block, main task, and microtask; responsibility is mapped successively to the Transaction Matching System AI, the cockpit AI, and a named human. Both relationships are recorded in the event ledger. Section VI sets out the requirements for rule cards, the event ledger, and links among position, task, and individual run records.

MASS breaks business activities down into microtasks with clear boundaries that can be verified independently, and uses large numbers of collaborating agents to complete them. Its name combines "Micro Agents" and "Simple & Short Tasks." It also conveys the English meaning of "mass."

"Micro" refers to the granularity of the task; execution still uses the model's full capabilities. Each task specifies its inputs, outputs, authority, completion criteria, and verification method. Verified results provide the basis for subsequent tasks. Matters requiring discretion are submitted to an authorized human for judgment, with AI providing information and analytical support.

(3) Financing Side: Communication Services, Contract Registration, and Admission

L1: “First Investor” Services handles business inquiries, preparation of materials, and ongoing communication. At reception, it also creates a client-dedicated AI agent, registers the source attribution, and assists in collecting contracts, evidence of contributions, and cash-flow records.

L2: Contract Registration extracts elements such as contracting parties, cash-flow arrangements, and term, and cross-checks them against transaction records, business registration information, and third-party data. It retains contract texts, evidence of contributions, and cash-flow records, and completes the signing and registration of listing documents. It links registration records to original evidence and retains change records.

L3: Qualification and Admission verifies, for “First Investor” qualification, the authenticity of prior contributions and the applicant's capacity and position to continue following the relevant business operations. Contract admission determines whether a contract can proceed to subsequent activities based on the qualification conclusion, registration materials, and applicable conditions, and calculates the share available for sale that day. Qualification, registration, and admission each have separate records of their bases and conclusions. Completed registration does not mean admission has been granted, nor does it constitute listing on a licensed trading venue. Admission conditions and factors used for measurement and disclosure are explained separately; a score cannot be treated as an admission conclusion.

(4) Investment Side: Reception Services, Qualification and Admission, and Portfolio Registration

R1: Investor Services receives Professional Investors onboarded through licensed institutions, creates client-dedicated AI agents, and serves as the single channel for investor notifications. Through dedicated representatives, it explains rules, responds to inquiries about business status, and handles communication about maturities, exceptions, and complaints. AI employees do not provide investment advice.

R2: Investor Admission: know-your-customer (KYC) checks are the responsibility of the relevant licensed broker. R2 receives and checks the identity and qualification verification results provided by the broker, verifies the onboarding institution, and handles platform admission, terminal access activation, and limits. Platform access and permission administration do not replace licensed institutions' client verification responsibilities.

R3: Portfolio Registration uses a single ledger for orders, holdings, distribution priorities, and Contract Pools, maintained by a Single Writer. R3 records investor orders and updates holdings and Contract Pools based on trade results and funds reconciliation results. Investors make investment decisions independently, and orders are submitted through the relevant licensed channels.

(5) Market Operations: Information Flow, Trading Venue Interface, and Funds Flow

C1: Information Flow performs Compass measurement using contract and operating evidence and updates readings according to the nature of the data and new evidence; daily checks do not mean values are rewritten daily. The system incorporates the “First Investor”'s Contract Performance Record, retained share, and method of selecting the retained portion into measurement and disclosure, and scans information against conditions set by investors themselves. Readings must state their evidence base, validation status, and scope of applicability. Gaps in the materials are recorded as they stand.

C2: Trading Venue Interface prepares relevant materials, coordinates with licensed institutions, transmits authorized instructions, and receives and checks transaction results. Activities requiring a license are executed by the relevant licensed institutions under their rules.

C3: Funds Flow transmits authorized funds instructions for settlement, refunds, and cash-flow distributions, and reconciles custodian confirmations, contracts, and holding records. Cash-flow calculations and distributions are processed according to valid registrations and authorizations. Client funds are held by licensed custodian banks.

Compass readings are labeled with one of four validation statuses: simulation only, historical data replay, testing against actual contracts, or revision based on real-world outcomes. When a contract enters the market, the predictions made at that time are entered in the ledger. When the contract is fully settled, the system automatically compares predictions with actual outcomes and records revisions. The Core Compass and Industry Compass examine common factors and industry factors, respectively. Returns and risks are assessed separately, and fields with insufficient information are left blank.

(6) Central Control and Operational Support

Central Control (T), a functional block of the Transaction Matching System, organizes rulemaking and the publication of approved rules, coordinates task assignments, cross-module collaboration, and exception handling, and identifies problems and promotes improvements through sample checks and backtesting. Rule revisions are approved under the applicable authorization procedures. Each module and professional position remains responsible for judgments and outputs within its remit. Central Control coordinates operations; the Responsibility Matching System implements the specific governance, authorization, and accountability requirements.

Compliance, legal, finance, and risk control provide support for rule applicability, contract documentation, accounting reconciliation, and risk resolution, respectively. Their opinions and the reasons for them are included in the relevant task records. IT and data management support systems, interfaces, permissions, and

records, retaining linked rule versions, evidence, execution events, and authorization records. Independent audit examines rule execution and control effectiveness.

(7) Three Microtask Categories and Three AI Agent Clusters

Organizing AI employees' work requires clarity on how tasks are divided, how positions collaborate, and what execution authority and recordkeeping requirements apply. AI employees execute approved rules within their authorized scope; matters beyond those rules or that authority are decided by authorized humans.

Each microtask belongs to just one of three categories: execution, sign-off preparation, or judgment. Execution microtasks check, calculate, register, or notify according to rules, with A and B roles assigned to different large language models for mutual backup and verification. Sign-off preparation microtasks have two AI agents prepare materials, while appropriately qualified humans retain responsibility for decisions and signatures. General judgment microtasks use three different large language models, each making an independent judgment. Industry-factor assessment has one microtask for each of 500 industries, each staffed by a ten-member team drawn from different large language models, including a team leader, members with industry experience, and members dedicated to finding errors.

Functional Block	Microtasks	Execution	Sign-Off Preparation	Judgment	AI Agents on Standing Task Assignments
T	130	109	4	17	277
L1	90	81	0	9	189
L2	90	79	2	9	189
L3	90	73	3	14	194
R1	90	85	1	4	184
R2	90	76	3	11	191
R3	90	81	1	8	188
C1	590	78	2	510	5,190
C2	90	81	2	7	187
C3	90	77	4	9	189
B	90	76	4	10	190
Total	1,530	896	26	608	7,168

Table 2-4 Proposed Microtask and Standing AI Assignments by Functional Block

Note: The figures in this table are provisional for the current version and will be updated based on validation results and changes to the plan. There are 608 judgment microtasks, including 500 industry-factor assessment microtasks, all within C1, and 108 general judgment microtasks. The 7,168 AI agents shown in the table have standing microtask assignments; they exclude client-dedicated AI agents, whose number grows with the client base, and supervisory AI agents within the cockpits. These figures represent the configuration proposed in this design, rather than the scale of actual operations.

Microtask Category	Number of Microtasks	Standing AI Agents per Microtask	Total Standing AI Agents	AI Agent Cluster
Execution	896	2(A and B roles)	1,792	Microtask AI Agent Cluster

Microtask Category	Number of Microtasks	Standing AI Agents per Microtask	Total Standing AI Agents	AI Agent Cluster
Sign-Off Preparation	26	2(preparation)	52	Microtask AI Agent Cluster
General Judgment	108	3	324	Microtask AI Agent Cluster
Industry-Factor Assessment	500	10	5,000	Industry-Expert AI Agent Cluster
Total	1,530	–	7,168	The First Two AI Agent Clusters
Client-Dedicated	**Number of Microtasks:** No standing microtask assignment **Standing AI Agents per Microtask:** 1 per client **Total Standing AI Agents:** Grows with the number of clients			Client-Dedicated AI Agent Cluster

Table 2-5 Microtask Categories and AI Agent Clusters

Note: The figures in this table are provisional for the current version and will be updated based on validation results and changes to the plan.

AI employees in the Transaction Matching System are organized into three clusters. The Microtask AI Agent Cluster handles execution, sign-off preparation, and general judgment, while the Industry-Expert AI Agent Cluster handles industry-factor assessment. Under the current microtask plan, these two clusters are intended to have 7,168 AI agents on standing task assignments. The Client-Dedicated AI Agent Cluster is configured by client: each “First Investor” and each Professional Investor has one AI agent, created by L1 and R1, respectively. These agents provide ongoing service from registration onward and follow each of their client's transactions; they do not have a standing assignment to a particular microtask. All three clusters follow the same rules and permission requirements, and all operations are recorded in the event ledger.

(8) Independent Judgment and Reliability Verification

For general judgment tasks, three different large language models each make independent judgments based on evidence and retain their respective reasoning. Disagreements are submitted to a human. Industry-factor assessment follows its dedicated ten-member team configuration.

Facts may be passed along the workflow, but independent judgments are not fed into one another; they are brought together only at the end of the process and submitted to a human. Unverified declarations and AI-generated profiles cannot replace original evidence. Verification starts with original documents and cross-checks them against bank statements, business registration information, and other reliable third-party materials. Where differences exist, the differences, reasoning, and resolution conclusions are retained. The system continuously tracks errors arising from common sources and introduces external checks of sampled conclusions. Agreement among three models does not remove the requirement for human review, and high-risk steps still require human sign-off.

Execution tasks use A and B roles for backup, continuing work from the same task record to prevent omissions or duplicate operations. AI employees are granted autonomy by level; a single costly error triggers an automatic downgrade. Text in external materials cannot change AI permissions, and critical operations require separate authorization. The system retains arrangements for manual offline takeover and recovery drills, and clients always know when they are communicating with AI.

IV. Information Matching System (IMS)

This section explains how intermediaries deliver content to both sides and help clients enter the market, as well as the boundaries governing information disclosure and remuneration allocation.

(1) Matching Hub and Intermediaries on Both Sides

The Information Matching System comprises financing-side intermediaries, investment-side intermediaries, and a Matching Hub operated by NUMA. Financing-side intermediaries introduce “First Investors” to the market, help “First Investors” prepare materials, arrange interviews, and submit contracts. Those employed by an institution must obtain their employer's written consent and register. Investment-side intermediaries are licensed institutions and their representatives, or intermediary members registered and approved in the future under trading venue rules. They help Professional Investors open accounts through licensed institutions and complete identity checks and investor education. The Matching Hub is intended to organize and produce content about the market, industries, and investment approaches centrally and continuously, distribute it to the relevant audiences through intermediaries on both sides, and manage communication records and handoffs to the Transaction Matching System. The Hub's algorithm, the Information Matching Engine, is expected to be introduced around December 3, 2026.

Content communication has three levels: market and industry introductions, outreach to groups, and outreach to specific audiences. The Matching Hub organizes content, while intermediaries on both sides distribute and communicate it within their approved scope. Content, audience qualifications, regions, and channels must all undergo the relevant checks. Targeted delivery remains subject to existing disclosure and licensing requirements; a more clearly defined audience does not relax these restrictions.

(2) Category Information, Registration of Interest, and Identity Disclosure

The Matching Hub publishes only category-level pipeline information using templates, including industry, term range, size range, and expected quarter of market entry. It does not identify specific issuers, contract terms, prices, or expected returns. A category with fewer than three contracts is not displayed separately. Content and channels require prior approval, and information is shown only to Professional Investors whose identities have been verified by licensed institutions. Intermediaries may neither modify it nor privately add details.

Investors who have opened an account with a licensed institution and completed identity verification may register interest in a category of contracts. An indication of interest is nonbinding: no funds are collected, no queue is formed, and no allocation entitlement or priority is created. Only the number of interested investors is shown to the financing side. Intermediaries register in advance, making sources and the division of service responsibilities traceable, but registration fields do not automatically determine remuneration shares. An investor's identity is provided to the other party only with that investor's consent. When an investor voluntarily discloses information at a particular stage of platform service, the platform provides the corresponding service as far as possible within compliance requirements. Entering the market or trading does not require the investor to disclose their full identity to the other party.

(3) Market-Entry Handoffs and Conduct Boundaries

Content sent by intermediaries, including AI-generated content, must pass a prohibited-terms check. Content requesting investor contributions before listing cannot be sent. After the introduction is completed, the “First Investor” enters L1, and the Professional Investor enters R1 through a licensed institution. The relevant microtasks in L1 and R1 register source attribution. Subscription remains subject to qualification, suitability, and confirmation of risk understanding.

Off-market intermediation does not execute trades. Intermediaries may use only functions that the system has made available and authorized. Financing-side intermediaries must not represent investors, express views on listing terms, prices, or returns, promote specific products, negotiate transaction structures, or accept or

transmit orders. Before listing, they must not invite investors to view specific contracts or request contributions, and must not use expressions such as principal, principal protection, guaranteed interest, or repurchase in off-market promotion. Content and related operations in off-market activities that require a license are likewise the responsibility of the relevant licensed institutions within their licensed scope.

Intermediaries must not charge clients or collect payments on their behalf, share remuneration outside the system, or claim to represent NUMA under titles such as "partner" or "agent." The system blocks registration of interest by people without accounts, access to investor identities without consent, display of pipeline information to nonprofessional investors or residents of regions outside the service area, and release by the person processing the matter. Investment-side intermediaries must not promote the market to residents of mainland China. Financing-side intermediaries must not publicly promote this market to unspecified audiences in China, hold promotional events there, or charge “First Investors”. All such conduct and blocking results are recorded throughout the process.

(4) Market Incentives and Remuneration Allocation Support

1. Platform-Funded Early-Stage Market Incentives

In the early stages of market development, NUMA is willing to use part of its platform revenue to support and incentivize market participation. The specific incentive arrangements and rules will be announced later.

2. Remuneration Allocation Support Services

Under agreements already reached among market participants, the platform provides bookkeeping and transfer support for remuneration allocation. Any transfers of funds are executed by the relevant licensed institutions under valid authorization.

When each transaction is executed, the client decides which registered intermediary or intermediaries will receive remuneration. The system registers that decision and executes the agreed arrangements. Different execution methods may be used according to market needs and the relevant agreements. The specific triggers and release arrangements follow the applicable rules.

Intermediaries receive agreed remuneration only for business they directly handle. They receive no entry fees, recruitment rewards, or tiered rebates, and earn no remuneration for recruiting other intermediaries. Investment-side remuneration is paid to the licensed institution for handling under its internal procedures, rather than directly to individual representatives or to anyone other than the registered payee. The specific rules will be published after the relevant approvals have been obtained.

V. The Contract Lifecycle and Collaboration among the Three Systems

The following account traces the processing of a business case to explain how each stage is handed over, verified, and completed, how exceptions are handled, and who retains responsibility.

The business workflow connects activities across modules. For each step, it specifies processing order, state transitions, handoff conditions, and completion requirements to support execution, verification, and traceability.

(1) Contract States and Processing Order on Both Sides

The full contract lifecycle covers qualification, registration, measurement, listing, allocation and trade execution, settlement and holdings, ongoing monitoring, and exit or termination. Steps connect according to the applicable conditions, and activities requiring a license are executed by the relevant licensed institutions.

The system records contract states separately from the progress of specific tasks, identifying a single primary state and a person with ultimate responsibility at every point in time. State transitions must be supported by the relevant evidence and authorized decisions. On exit or termination, the reasons, resolution results, and outstanding obligations are recorded.

If a contract cannot proceed, its current process must be formally closed rather than left open indefinitely. Refusal of acceptance, withdrawal, failure to transact, termination, and recognition of loss are each recorded separately. Contract exit and the handling of outstanding obligations are recorded separately.

On the financing side, L1 receives the “First Investor” and helps prepare materials. “First Investor” qualification in L3 is Gate 0, verifying genuine prior contributions and the capacity for ongoing follow-up. L2 collects, verifies, and registers contract elements, after which L3 checks contract admission and the share available for sale against the relevant conditions. Qualification, registration, and admission each retain separate records of their bases. The positions of L1, L2, and L3 in the diagram do not mean every matter is processed in numerical order.

The investment-side workflow connects reception in R1, admission in R2, and registration in R3. Professional Investors enter R1 through licensed institutions to receive rule explanations and dedicated services. R2 receives and checks identity, qualification, and related verification results provided by licensed institutions and handles platform activation and limits. R3 registers holdings, distribution priorities, and Contract Pools based on orders and transaction results. Reception in R1 does not open access to subscription. Licensed institutions complete product-specific suitability assessments and confirmation of risk understanding under their rules; only after the system obtains verifiable results is the corresponding subscription permission activated. As the single channel for investor notifications, R1 continues to provide communication services throughout subsequent activities.

(2) Registration of Rights, Distribution Priorities, and Batches

Both individual contracts and Contract Pools must have their cash-flow scope, distribution method, shares of rights, and applicable cash-flow distribution priorities registered. Pool records link to the constituent contracts, and batch records correspond to the relevant transaction terms, connecting contracts, rights, holders, and the bases for distributions.

Events such as the establishment of distribution priorities, transfers of rights, priority changes, buyouts, payments to holders, and releases are registered as events, with the priority table updated in real time. Corrections are made by appending records, preserving the original records and the bases for correction.

Any adjustment by the “First Investor” to the terms of the portion not yet transferred applies only to subsequent transactions and cannot unilaterally change the rights of registered holders. Cash-flow distributions are based on the valid records of the corresponding contract, Contract Pool, and batch.

(3) Task Handoffs, Exceptions, and Completion

Cross-module handoffs must identify the business object, version of the materials, work completed, unresolved issues, and the person responsible for receiving the work. Acceptance, return, requests for additional materials, and timeout handling are all recorded.

Task completion requires the applicable output, verification, and authorization requirements to be met, relevant issues to be addressed, and complete evidence to remain available for inspection. Completion, return, suspension, and termination are recorded separately. Completion of an individual task does not mean the contract has ended.

Internal matters such as compliance, legal affairs, finance, and rulemaking are also included in task management. Opinions, approval decisions, and the related business or matters are retained in linked records.

(4) Funds Processing and Reconciliation

Responsibilities and supporting bases are recorded separately for calculations, approvals, instruction transmission, and result reconciliation. A payment is considered complete only when the actual movement of funds has been confirmed.

Funds processing must distinguish five matters: holding funds does not mean an obligation exists to pay the same amount; completion of a calculation does not mean payment has been approved; sending an instruction does not mean the recipient has received it; a matching reconciliation does not mean there is no dispute; and closing an exception does not extinguish an amount owed. If the execution result is unknown, the original instruction is checked before any further action, rather than immediately resent. If only part of a payment arrives, the difference is verified and only that difference is paid.

The Register, custody accounts, and distribution ledger undergo daily three-way reconciliation, with differences required to be zero; new business is suspended if they are not. Corrections are made only by appending records. The contracting entity, the legal nature of the receiving account, beneficial entitlement, and authorization chain must be consistent and verifiable, and platform fee accounts must be separate from client funds accounts. Funds decisions and changes to paying or receiving accounts require two human signatures, and the two signatories must not be in a direct reporting relationship. Payments that have received dual sign-off may be executed within valid authorization approved in advance. AI recalculation does not count as one signature.

(5) Continuity of Execution and Change Records

Position, task, and individual run records are linked. When models, sessions, or handlers change, or a cockpit is dissolved and its work reassigned, progress is confirmed from existing records to prevent omissions or duplicate execution. System replay is separate from actual execution and does not retrigger real operations such as funds movements.

When materials, outputs, or rules change materially, affected verifications and approvals are rechecked for validity, and versions before and after the change and the bases for handling it are retained.

"Delivered," "Completed," and "Archived" are verified separately. Delivery of an output does not mean it has passed inspection, and evidence must remain available after a task is completed. When an output undergoes a material change, the continued validity of affected checks and approvals is reassessed. Any required rechecking, decisions, or sign-off follow the applicable rules.

(6) Collaboration among the Three Systems: A Contract Example

Consider a property revenue-sharing contract. Collaboration among the three systems begins off-market. A financing-side intermediary helps the property holder prepare materials. When the contract-count threshold and display conditions are met, the Matching Hub publishes category information such as property revenue sharing, term range, size range, and expected quarter of market entry. An investment-side intermediary explains the category to eligible clients through approved channels, and clients register nonbinding interest after opening accounts with licensed institutions.

Once the property holder enters L1, the system creates a client-dedicated AI agent and registers the source. Gate 0 in L3 verifies prior contributions and the capacity for ongoing follow-up, with different models making independent judgments and a qualified human signing the qualification conclusion. L2 then completes the registration of contract elements, evidence, and related documents; L3 checks admission conditions; and the industry-expert team in C1 completes industry-factor assessment and Compass measurement.

C2 submits listing materials to the licensed trading venue. Once Professional Investors meet the relevant qualification and product suitability conditions, they place orders through licensed channels. Allocation and trade results are written back to R3 after reconciliation. C3 verifies the arrival of funds at the custodian bank, and disbursement instructions are transmitted after valid dual human sign-off. Throughout the contract's life, C1 continues to update readings, L1's client-dedicated AI agent follows business operations, and C3 reconciles cash-flow distributions against valid registrations, until the contract is completed or terminated under the rules.

From market entry to termination, every transaction is continuously followed by a transaction-tracking supervisory AI agent and has a named human responsible at every point in time. When personnel change, transactions and client-dedicated AI agents are reassigned across cockpits so that processing continues. Completion of a single microtask does not mean the entire transaction or contract has ended.

VI. Common Controls and Technical Support

This section explains the rules, permissions, records, and compliance requirements shared by the three systems, and how the system will continue to improve.

(1) Rule Cards, the Event Ledger, and Execution Records

Work is executed according to rule cards and entered in the event ledger. Rule cards specify the applicable circumstances, basis, processing method, and recipient for exception escalation, with each card retaining its version and status. The event ledger records who performed which operation, when, under which rule version, and with what result. Records are appended, never overwritten.

The Tag Repository is intended to use common tags to describe objects such as personnel, microtasks, contracts, investors, industries, and content. This allows the different systems to use consistent definitions when assigning responsibility, matching business activities, and passing information, and to connect their processing results. The specific tags and their scope of use remain to be refined through design and simulation testing.

Specific operating methods determine how matching, calculation, and processing are performed. Rule cards define the conditions, authority, and restrictions governing those methods. Operating methods must remain subject to valid rules; calculation results cannot by themselves replace approval or sign-off. Rule maintenance and method adjustments retain separate records of their bases, versions, and approvals.

AI organizational positions, tasks, and individual runs each receive an identifier, linked to their functional block, microtask, cockpit and named responsible person at the time, actual handler, and applicable approval records.

Execution records retain the versions of the rules and materials used, inputs and outputs, authorization bases, verification results, and exception handling. Even after a model, session, or handler changes, the course of a task can still be reconstructed so that work can continue, results can be reviewed, and responsibility can be traced.

Each decision should be reconstructible using the inputs, rule cards, and model versions in effect at the time.

The inputs, conditions, and versions used at the time should make it possible to reconstruct the AI's processing and review or reproduce the results. Reproducibility remains subject to correctness testing and does not mean everyone may access all the same materials. Visibility for different positions and audiences continues to follow permission and data-isolation requirements.

(2) Delegated Authorization for AI Agents, Data Boundaries, and Audit Records

Delegated Authorization for AI Agents is based on valid business authorization granted to a named human. AI employees' access to materials and operational permissions are configured by position and task. Taking responsibility for an AI position does not increase approval authority, and materials for cross-module collaboration are provided according to business needs.

For cockpit supervision, cross-module collaboration, and reassignment, every invocation checks both the caller's business authorization and the AI position's authorization. Execution is limited to the scope permitted to both parties, and the purpose of the invocation, permission basis, and execution result are recorded. Changes in team configuration do not expand either party's access to materials, operational permissions, or approval authority.

Client identity information does not enter external models, and original evidence does not leave designated data regions. Business collaboration, analysis, and verification all observe the relevant data boundaries.

These data boundaries also apply to excerpts, prompts, vector databases, logs, and vendors' remote access. Personal information and important data within China are processed within China under Chinese law. Only de-identified, aggregated disclosure materials may be provided abroad, with the prescribed procedures completed. Licensed institutions are responsible under applicable legal obligations for client due diligence, list screening, and suspicious transaction reporting; the OS provides monitoring and materials. Suspicious transaction reports do not enter the ordinary approval queue, are kept strictly confidential, and are not disclosed to the subjects concerned.

When an error occurs, records must link it to affected tasks, rule and material versions, and related operations and decisions, supporting determination of the scope of remediation and retesting. Corrections are made by appending records, preserving the original supporting evidence.

Records for audits and inspections are provided under the applicable procedures and data permissions, with access and the scope of provision recorded. The development and validation status of interface, replay, and tracing capabilities is accurately labeled.

Read-only queries for audit or regulatory purposes provide decisions, their bases, and rule versions under applicable procedures. They confer no right to modify records and do not expand existing access to materials.

(3) Technology Platforms and Support for Continuous Operations

The agent platform, knowledge base, and IT infrastructure provide technical support for the three systems; they do not constitute an additional business or accountability system:

Agent platform: organizes AI task execution and links position, task, and individual run records.

Knowledge base: provides rules and business materials approved for use, retaining cited content and versions.

IT infrastructure: supports interfaces, permissions, event records, and task pipeline operations.

Model and cloud service vendors must undergo due diligence, with replacement contingency plans in place. Critical steps use cross-checks across different large language models to reduce dependence on any single vendor.

Task records retain inputs, rules, outputs, verifications, and human decisions, supporting reconstruction of the process, investigation of problems, and retesting of results.

A 7×24-hour task pipeline, described as a "Lights-Out Factory," is the design target for continuous operations. Authorized tasks continue to execute according to rules, while matters requiring human decisions are submitted promptly. Actual operating capability must be tested for load, continuity of task execution, and exception handling.

(4) Compliance Constraints and Responsibility Boundaries

The following requirements apply throughout the design, development, and operation of the OS. System functions, levels of automation, and internal authorizations must all comply with applicable regulatory requirements and the allocation of licensed functions.

The business responsibilities of NUMA and licensed institutions do not transfer to others through the use of AI or third-party services.

In conducting business or adjusting products or rules, NUMA must comply with applicable laws and the rules of relevant licensed trading venues, and complete applicable approvals, deliberations, and necessary procedures. Activities that have not met the relevant requirements will not be conducted.

This part presents a draft OS framework; it does not indicate that compliance approval has been obtained. Completion of system configuration, deployment, or simulation testing does not replace the qualifications, authorizations, and procedures required to conduct business.

Steps requiring a license are executed by licensed trading venues, licensed brokers, and licensed custodian banks within their licensed scope and under their rules. The OS organizes materials, connects interfaces, transmits authorized instructions, and reconciles results. Internal state records must not replace valid business results from licensed institutions.

NUMA intends to serve only Professional Investors through licensed institutions. Investment activities will not be open to the public or retail investors. NUMA will not raise funds within China or accept funds from within China for investment.

Client funds are held only at licensed custodian banks. The OS transmits funds instructions under authorization and reconciles results. Funds operations must not exceed valid authorization.

Investors make investment decisions independently and bear investment risks. Compass measurement, condition-based scanning, and dedicated representative services provide information and explanations of rules. AI does not provide investment advice and must not make choices for investors through default actions.

The OS carries forward the October 3 white paper's inspection approach of "check people, inspect systems, and review standards":

Check people: verify named assignment of responsibility, personnel qualifications, authorizations, and responsibility for specific matters.

Inspect systems: examine task execution, controls, exception handling, and related records.

Review standards: examine applicable rules, versions, completion requirements, and verification bases.

Inspection records are provided under applicable procedures and data permissions. Capabilities still under development or not yet validated are accurately labeled.

NUMA does not promise investment outcomes, but must remain responsible for its own registration, data, systems, and services. Registration should be accurate; use of materials and transmission of instructions should comply with authorization; and exceptions should be handled properly. Investors' assumption of investment risk does not relieve NUMA of responsibility for its systems and services.

When errors or exceptions are identified, their impact is determined, relevant records are corrected, and necessary remediation and retesting are completed. The original supporting evidence, handling decisions, and results are retained to support inspections and subsequent improvements.

(5) Continuous Improvement and Version Management

In developing the OS, we first define business objects, processing workflows, state changes, evidence records, responsibility, and reconciliation mechanisms, and then use these as the basis for designing and improving operating interfaces. Each processing step should have a basis, the person responsible for handling differences should be clear, and decisions should be traceable. Each control point must specify the conditions for proceeding and when work should be returned, suspended, or submitted to a human, providing a shared standard for execution and inspection.

System interfaces, cockpit interfaces, and dashboards should correspond to these workflows, records, and responsibility arrangements. Interfaces transmit instructions and results; user interfaces display tasks, outputs, responsibilities, and rules, and support review of the processing history. Displayed states must be supported by the relevant business records, evidence, and reconciliation results. Subsequent extensions and revisions should also retain links to the original records.

The draft is intended to support modular integration of external data, rules, AI employees, operating interfaces, and display views. Each integration specifies its module, business purpose, inputs and outputs, permissions, costs, responsible person, and rollback plan, and links to the relevant functional block, main task, and microtask. Data applicability, rule compatibility, and permission boundaries are validated before integration.

New components are tested first, then run in shadow mode with results compared. Once they meet the requirements, they are approved under authorization procedures and brought into use. Changes involving workflows, permissions, or interfaces undergo impact assessment and retain version records. If problems occur during the switch, the component is disabled or rolled back under a predefined plan.

Modular integration prioritizes reuse of existing interfaces and control mechanisms, with the specific scope of changes determined by validation results.

When an industry is added, a new team is established in the Industry-Expert AI Agent Cluster. An assembly checklist, autonomy-level table, and table of differences in operating definitions are completed, and the team is configured and validated within the existing microtask structure. Expansion must comply with permission, responsibility, and rollback requirements. Actual effectiveness remains subject to validation results.

The current microtask breakdown remains preliminary and will be adjusted based on test results. Subsequent versions are updated based on validation results, retaining rules, configurations, and evidence from before and after revisions. The microtask breakdown and technical configuration are revised under verified and approved engineering plans; adjustments affecting business rules or responsibility arrangements still follow the relevant procedures.

Simulation and actual operations are labeled separately. A successful simulation supports conclusions only for the conditions tested. Version updates must still record problems, remediation, and retesting results.

This issue introduces the Responsibility Matching System, Transaction Matching System, and Information Matching System, explaining their components, functions, and responsibility arrangements. In December, we plan to introduce the specific methods that operate on these systems - the "engines" - to explain how responsibility is matched, business is processed within the market, and information is communicated outside it. Simulation tests both whether the responsibilities and interfaces of the three systems are clear and whether their operating methods are effective. The related designs will continue to improve accordingly. The operating plan appears in Part III.

Part III | Market Preview

Simulation Testbed Operating Plan

Reader's Guide

This issue's "Market Design" introduces market participation arrangements that begin with the “First Investor”, while "Market Progress" sets out the components, functions, and responsibility arrangements of the NUMA Operating System (OS). Building on these, the Simulation Testbed will organize simulated work involving contracts, systems, and personnel to observe how the market arrangements and execution framework operate together. This also continues the October 3 white paper's vision of "continuous operation in a simulated environment." This column further sets out the scope, organization, and phased plan for those operations.

The Simulation Testbed is intended to generate synthetic contracts, operating data, and participant profiles, and organize full-lifecycle processing from contract entry to exit under the Operating System's rules, with human and AI employees participating in their assigned roles. Humans assume authorized responsibilities for judgment, sign-off, and supervision; AI employees execute tasks within rules and authorizations. All activities described here take place within the Simulation Testbed. No real transactions are conducted, no real funds are handled, and no investment payments are accepted.

This simulation round examines four dimensions: feasibility, bottlenecks, robustness, and cost-effectiveness. Its focus is whether tasks in the Transaction Matching System can be completed reliably, whether Human–AI Cockpits can provide continuous supervision and maintain continuity when personnel change, and whether off-market intermediation can both bring the two sides into the market and maintain communication and participation boundaries. Findings will be used to revise rules, adjust the microtask breakdown, refine cockpit team configurations, and provide a basis for subsequent validation of Compass methods.

Internal simulation is planned to begin on November 3, 2026. Around December 3, we plan to further introduce the Operating System's operating methods and publish an interim report on the Operating System and Simulation Testbed. The target date for consolidating results from this stage and publishing the fourth white paper is around January 3, 2027. Each stage will report actual coverage, operating findings, revision progress, and unfinished matters. Specific dates for external participation depend on actual progress.

I. Positioning, Objectives, and Outputs of the Simulation Testbed

The Simulation Testbed is an operating and testing environment isolated from real business. It uses synthetic contracts, operating data, and participant profiles to organize full-lifecycle processing under the NUMA Operating System's rules and observe how business processing, system execution, and human–AI collaboration perform under different conditions.

All employees are expected to establish their own Human–AI Cockpits, supervise their assigned microtasks and transactions, and assume the corresponding responsibilities for authorization, judgment, sign-off, and supervision. AI employees execute tasks within rules and authorizations. The entire test takes place within the Simulation Testbed, and employees fulfill these responsibilities through their Human–AI Cockpits. Rules, task configurations, and operating records explain how the two correspond.

This simulation round focuses on four objectives:

1. Feasibility: test whether the financing-side workflow connects from off-market intermediation and “First Investor” qualification through contract exit or termination, and whether the investment-side workflow connects from account opening and qualification verification through simulated cash-flow receipts. Test whether business that fails to meet the conditions can be returned, suspended, or terminated under the rules.

2. Bottleneck identification: identify where waiting, send-backs, and reprocessing build up, distinguishing causes such as insufficient materials, unclear rules, system execution, position handoffs, and the handling of judgments.

3. Robustness: observe whether critical controls over simulated funds processing, sign-off authorization, cash-flow distributions, and disclosure remain effective when intake rises, personnel are absent, or AI judgments are wrong, and whether exceptions can be identified and handled promptly.

4. Cost-effectiveness: record elapsed time, resource consumption, and human effort across the full business lifecycle, include failures, rework, and exception handling in cost statistics, and assess the effects of workflow and configuration adjustments under comparable conditions.

Every simulated case must record its current state, responsible person, and basis for handling. Completed or terminated cases retain their results and reasons; unfinished cases identify the blocked step and subsequent arrangements. Sign-off must be based on actual judgment. Returns, disagreements, and exception handling are also included in operating records.

This round is intended to produce operating files that can be reviewed, records of problems and rule revisions, cost data, and lists of tested and untested matters. These will support improvements to the Operating System, adjustments to the division of work between humans and AI, and identification of issues requiring further validation in Compass data collection and scoring. Records retain inputs, rule versions, execution, verification, and human decisions to support replay and comparison. Test results will also be used to adjust the microtask breakdown, refine cockpit team configuration methods, and confirm the boundaries of off-market intermediation.

II. Operating Architecture and Coverage

The Simulation Testbed uses this issue's market arrangements and OS plan as the basis for testing through synthetic cases, role interactions, and task execution. Figure 3-1 shows business workflows on both sides, Human–AI Cockpit supervision, simulated execution support, and a closed feedback loop for testing. The following sections explain which business steps and organizational responsibilities are intended to be covered.

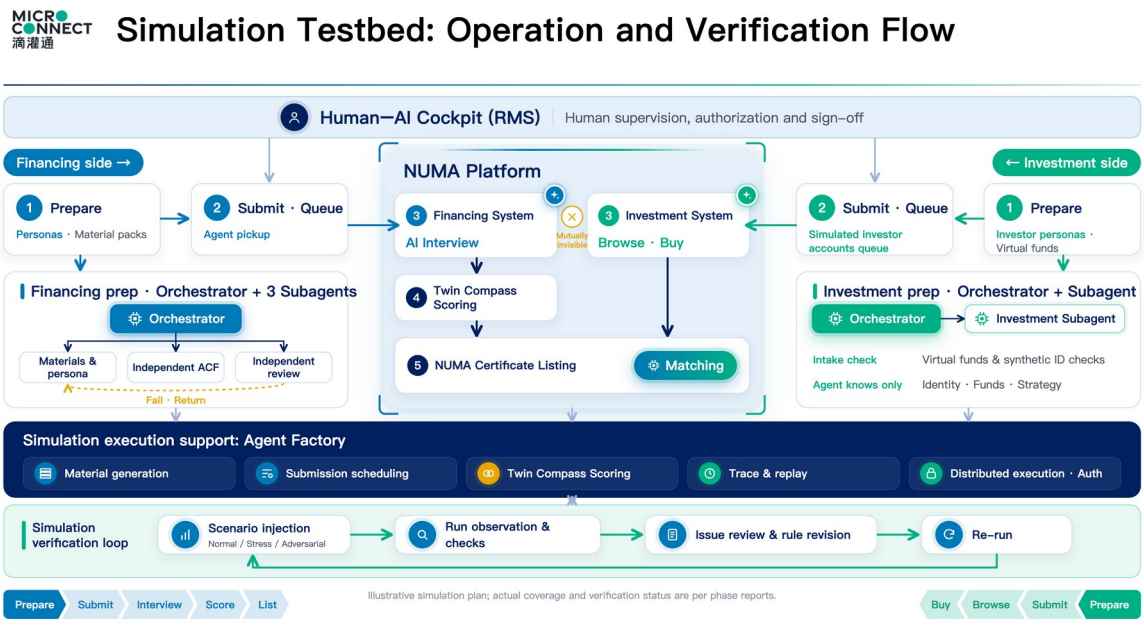


Figure 3-1 Simulation Testbed Operations and Testing Workflow

Simulation will directly test whether the responsibilities of the three systems are clear, whether their interfaces connect, and whether tasks, information, and responsibility can be handed over smoothly at each step. Around December 3, 2026, we plan to further introduce the specific methods operating on these systems - the "engines." The systems define components, functions, and responsibilities; the engines explain how responsibility is assigned, business is processed within the market, and information is communicated outside it. The relevant operating methods will be progressively included in this simulation round as preparations allow. Their design and testing progress as of December will be presented in further detail then.

(1) Relationship to the OS and Target Coverage

This round is intended to cover Human–AI Cockpits, qualifications, authorizations, and responsibility handoffs in the Responsibility Matching System; the Five Zones and Eleven Functional Blocks and their 1,530 microtasks in the Transaction Matching System; and financing-side intermediaries, investment-side intermediaries, and the Matching Hub in the Information Matching System. The Five Zones and Eleven Functional Blocks and microtasks organize work within the Transaction Matching System; they are not organizational units at the same level as the three systems. Figure 3-1 identifies the principal execution, supervision, and testing relationships. Capabilities will be integrated progressively, and stage reports will distinguish what has been integrated, tested, and remains to be improved.

The Orchestrator Agent and Subagents in the figure collaborate on tasks to prepare materials, perform checks, and carry out business operations. Work in the Transaction Matching System is broken down successively into functional, main task, and microtask layers. AI assignments begin at the microtask layer. Humans assume responsibility for supervision, judgment, and sign-off through Human–AI Cockpits. Task configurations and responsibility records establish the specific relationships among execution agents, microtasks, and cockpits.

(2) Role Operations and Comprehensive Observation

Figure 3-1 shows role operations on both sides and Human–AI Cockpit supervision. Under this round's plan, AI plays the roles of “First Investors”, underlying business operators, Professional Investors, licensed institutions, and custodian banks using synthetic profiles. Employees play the financing-side and investment-side off-market intermediaries. Internal employees supervise microtasks and transactions through their own Human–AI Cockpits. Roles and interactions not detailed in the figure will be described in stage reports according to actual integration and testing progress.

Within the same case, an employee playing an off-market intermediary must not use an internal position to read information beyond that role's permissions, nor concurrently perform internal judgment or release duties incompatible with the simulated role. Role assignments, visibility, and arrangements to avoid conflicts of duties are recorded together. If a configuration does not meet the requirements, it is adjusted before the related work begins.

Roles on both sides use only information they are authorized to read. "Private Information Isolated" in the figure means that private role data and information about operations outside the other role's permissions are isolated from one another; market information disclosed under the rules remains available to the relevant roles. Independent review during financing preparation examines case materials. Operational observation links business events, task processes, rule versions, and handling results for tracing and retrospective review. Observation data is managed separately from information used in role operations, preventing roles from learning future outcomes or information outside their permissions in advance.

(3) Business Scope and Operating Period

The illustrated workflow begins with preparation on both sides and highlights submission, financing inquiries, Twin Compass scoring, listing, and investment-side browsing and purchasing. Allocation is simulated under the valid rules of the relevant version, without assuming that all business uses the same allocation method. This

round's target scope also includes “First Investor” qualification, contract registration, settlement and holdings, cash-flow distributions, ongoing monitoring, Contract Pools, and exit arrangements.

Testing is based on the business arrangements expressly adopted in this issue and covers Pro Rata Cash-Flow Distribution and Priority-Based Cash-Flow Distribution, as well as individual contracts and Contract Pools. Applicable conditions, including cash-flow scope, distribution proportions, priorities, and recovery caps, are checked against the relevant rules and registration records. Execution details that remain undecided are listed as matters to be clarified; they are not independently settled within the simulation.

This round is planned to run from November 3, 2026 to January 3, 2027, with the specific pace adjusted according to readiness and stage test results. Steps not integrated, simplified, or not yet tested are identified in stage reports. Overall coverage targets do not substitute for the actual scope of operations.

III. Simulation Methods and Scenario Design

This round is intended to create different business conditions through case generation, role interactions, time compression, and scenario testing, then analyze problems through event replay and version comparison.

(1) Architecture Mapping and Process Simulation

Using this issue's OS plan, the responsibilities and interfaces of the three systems are implemented in simulated operations. Functions within the Transaction Matching System are broken down successively into functional, main task, and microtask layers, establishing relationships among rules, positions, processes, and data. When problems emerge during operations, these relationships help identify the module, task, or responsibility arrangement involved so that revisions can be made.

(2) Role Modeling and Autonomous Interaction

Simulated roles such as “First Investors” and Professional Investors are assigned profiles, objectives, information permissions, and behavioral constraints to reflect differences in expertise, preferences, and judgment approaches. Roles participate in inquiries, declarations, and choices based on visible information, with specific behaviors and outcomes emerging through interaction.

Both normal behavior and behavior that departs from business rules are included in the simulation to observe how internal positions identify, judge, and handle them.

(3) Synthetic Cases and Data Generation

Contracts, operating data, and cash-flow paths are generated according to industry characteristics and scenario assumptions. Case background facts are retained as a verification reference. Submitted materials may contain deliberate omissions, contradictions, or deviations to observe whether the system can identify and handle them.

Material generation maintains the necessary separation from scoring rules to prevent the generation process from directly catering to expected scores. Genuine generation errors are returned for correction, while deliberately incomplete or noncompliant materials are retained as test samples. The two are labeled separately.

(4) Interactive Simulation and Batch Runs

Simulated roles use conversations and page operations to complete inquiries, supply additional materials, read information, and make investment choices. This tests whether business interactions obtain sufficient information and enable subsequent processing.

Under the same business workflow, the volume and concurrency of introduced contracts and orders are increased to observe task backlogs, cross-module handoffs, and result consistency. Problems in interactions and batch runs are recorded separately and analyzed to determine whether they share common causes.

(5) Business Time Compression

Operating and cash-flow cycles are compressed to observe events such as ongoing performance, overdue payments, buyouts, and exits within a limited period, while preserving event order and business dependencies.

Acceleration of business time is recorded separately from actual processing time, avoiding confusion between contract cycles and system processing efficiency. Emergency tests, such as continuity of work after personnel leave their posts, record handoff time, failures, and exceptions in real operating time. Performance under a compressed business clock is not used as a substitute for actual emergency response capability.

(6) Adverse Scenarios and Stress Testing

Scenarios address business load, staffing, AI judgments, and the operating environment. They include surges in intake, absence from critical positions, judgment errors, integration of new industries, business downturns, investor hesitation, and concentrated stress on Contract Pools. Tests of concentrated buyouts, large inflows of off-market intermediaries, content exceeding permitted boundaries, and prompt injection are added in light of the business rules applicable in this issue.

In adversarial scenarios, some roles actively submit noncompliant or misleading materials and make requests beyond their authority, testing whether controls over qualification, authorization, content transmission, and funds processing can be bypassed. Rule cards must not be bypassed. Tests record successful blocks, missed blocks, and false blocks, as well as subsequent handling and retesting results.

Observation focuses on where processes become blocked, whether controls over authorization, simulated funds processing, cash-flow distributions, and disclosure remain effective, and whether exceptions can be identified, resolved, and processing resumed.

(7) Event Replay and Rule-Version Comparison

Original records are used to reconstruct how problems arose, distinguishing case conditions, execution deviations, and rule defects. After rules or configurations are revised, comparable cases are rerun to check whether the original problem has improved and whether new effects have emerged.

For deterministic calculations, result consistency is checked under identical inputs, rules, necessary parameters, and states. For model judgments, conclusions, reasoning, and differences are compared; all free text need not be identical word for word. Historical records retain the conditions and versions in effect at the time, while revised results are recorded separately to enable comparison of effects before and after changes.

IV. Operating Organization and Iteration Mechanisms

The Simulation Testbed is intended to operate around tasks, with specific problems handled through successive levels. Problems common to different tasks are used to revise rules and configurations. Operating records, handling decisions, and revision results are retained continuously to support retrospective review and renewed validation.

(1) Daily Operations and Task Tracking

Contracts, orders, and events are introduced daily, with each position executing and signing off on its work. Each Human–AI Cockpit views assigned tasks and transactions according to its permissions. Run failures, waiting, returns, and exceptions are recorded at the same time in a bottleneck list; matters unresolved that day carry forward to the next round. Responsibility is checked for microtasks and transactions in progress that have entered operations. Overall design numbers do not substitute for actual integration numbers.

(2) Tiered Issue Resolution and Escalation

Problems follow the path below, with each level retaining the basis for identification, handling results, and reasons for escalation.

1. Entry checks. Verify that cases meet generation and introduction requirements. Generation errors are returned for correction; deliberately missing materials, contradictory information, or noncompliant declarations are retained as test scenarios to observe whether subsequent positions can identify and handle them.

2. AI self-checks and peer review. AI employees check task conditions and execution results, and independent verification positions review outputs. Review uses case source files and valid rules, recording incorrect judgments, missed judgments, disagreements, and their handling. Agreement with a human conclusion is not the sole criterion of correctness. Complete case source files are available only for authorized independent verification and comprehensive observation. Roles involved in business processing, review, and decision-making may use only information available to them under their permissions at that point in the business process. They must not use subsequent operating results or other undisclosed facts in advance. If an issue cannot be resolved or exceeds authorization, the supporting basis is retained and the matter submitted to a human.

3. Human takeover. An appropriately qualified and authorized human makes the judgment. Matters requiring authorization, professional confirmation, or discretion go directly to human handling. The outcome must specify whether processing continues, is suspended, or is terminated, with the basis for the decision and subsequent responsibility recorded.

4. Rule revision. Gaps in rules, recurring problems, or issues affecting multiple steps enter a review of rules and configurations. The cause is first identified as a matter of materials, methods, execution orchestration, authorization, or rules, and adjustments are then proposed under the corresponding authority. After revisions are approved, simulation is repeated to check whether the original problem has improved and whether new effects have emerged.

(3) Review and Iteration Schedule

During startup, operating progress and unresolved matters are expected to be checked daily, with a weekly consolidated analysis of problems shared across cases and positions. Daily checks facilitate prompt resolution, while consolidated reviews focus on identifying systemic causes and organizing revisions and retesting. Urgent matters are handled promptly. Rule or configuration adjustments can be implemented once the necessary checks and approvals are completed, without waiting for the consolidated review.

The frequency will subsequently be adjusted according to task load, the pace of problem identification, and validation needs, progressively shortening feedback and revision cycles. Reviews use operating records to improve rules, processes, and human–AI collaboration, while retaining records of each position's performance of duties and responsibility.

(4) Change Control and Version Management

Rules, task configurations, cockpit team configurations, and AI execution arrangements are version-controlled, with reasons for revision, validation status, approval bases, and scope of applicability recorded. Checks and approvals for model or prompt adjustments are also determined by their business impact.

Each test links its run identifier, inputs and necessary state, rule and configuration versions, outputs, check results, and human decisions so that problems can be tied to the specific run. Historical tasks retain their original conditions, while revised results are recorded separately. New versions enter the relevant simulation after retesting. Problems found lead to further corrections and, where necessary, disabling or rollback, with comparable and traceable supporting records retained.

V. Operational Observations and Verification Requirements

The Simulation Testbed is intended to support ongoing observation of processes, efficiency, information, evidence, and Human–AI collaboration in relation to the objectives above. The four objectives define what is to be tested; the five observation dimensions define the aspects to be analyzed. Specific statistical definitions will be established as capabilities are integrated, and interim reports will describe the samples, conditions, and results.

Evaluation Dimension	Definition	Importance	Application of Results
Process Completeness and Task Quality	Observe whether business processes connect, whether tasks meet completion requirements, and whether returns for rework, suspensions, terminations, and exceptions can be handled effectively.	Both normal processing paths and exception exits must work effectively before the validity of processes and controls can be assessed.	Identify process gaps and improve task conditions, control points, and exception-handling rules.
Operational Efficiency and Resource Costs	Observe the time required for business processing, waiting, handoffs, and rework, together with the model, tool, human, and other resources used.	Time and resource use jointly determine business processing costs and help assess whether the system can operate sustainably and accommodate business growth.	Identify the main sources of cost; improve task decomposition, execution orchestration, and resource allocation; and assess the effects of adjustments under comparable conditions.
Business Information Flow and Consistency	Observe whether the information needed by the financing side, the investment side, and internal positions remains complete, consistent, and usable as it is collected, registered, disclosed, and used.	Information gaps or differences in definitions can affect subsequent judgments and create waiting, misjudgments, or duplicate processing along the process.	Define necessary information requirements and improve field definitions, collection arrangements, and handoffs between modules.
Evidence Sufficiency and Applicability of Measurement Methods	Observe whether the evidence required by the Compasses can be obtained, verified, and used to support scoring, and whether the factors and methods are applicable to the relevant business conditions.	The validity of measurement depends on its evidentiary basis and conditions of applicability. Simulation can help identify issues that require further validation.	Distinguish missing materials from data access limitations and problems with method applicability; improve evidence collection and scoring arrangements; and develop a list of items for subsequent validation.
Human–AI Collaboration and Authorization Alignment	Observe whether human and AI employees' responsibilities, capabilities, and authorizations align, and whether human judgment, exception escalation, and task handovers work effectively.	Necessary human judgment and accountability form part of Human–AI collaboration. An unusual increase in manual handling may, however, indicate gaps in rules, capabilities, or configuration.	Improve the division of duties, task assignment conditions, authorizations, and escalation paths to support staff training and adjustments to AI configurations.

Table 3-1 Operational Observation and Evaluation Dimensions for the Simulation Testbed

In addition to the overall observations, this round is also intended to examine the following specific matters in the operational stages that have been integrated.

Operational Stage	Matters to Be Tested	Records to Be Retained
Responsibility Matching Arrangements and Operation	Whether the microtasks and transactions in progress included in the simulation have clearly assigned owners; whether qualification, authorization, and segregation of processing and release duties requirements are met; and whether work can continue when personnel change.	Actual coverage, assigned responsibilities, changes in team configuration and their reasons; actual handoff time, failures, and exceptions; and the basis for sign-off and release.
Transaction Matching Processes and Calculations	Whether qualification, measurement, registration, pooling, allocation, and funds calculations connect; whether rights, proportions, priorities, and recovery limits comply with valid rules and records; and whether discrepancies and exceptions can be identified and handled.	Case working papers, applicable rules, and registration and calculation results; reconciliation discrepancies, misjudgments, and missed issues; and the results of handling, suspension, and retesting.
Off-Market Collaboration and Content Transmission	Whether content and audiences comply with the rules; whether outputs beyond authorized boundaries can be blocked; whether simulated roles can complete the permitted market-entry steps; and how the associated costs change.	The basis for release, blocking, and transmission; simulated market-entry events and their statistical definitions; resources used and outstanding matters.

Table 3-2 Verification Requirements by Operational Stage

The corresponding verification and release requirements are zero violations of mandatory authorization conditions and segregation of duties, zero noncompliant releases of requests beyond authorized boundaries or in breach of rules, and zero discrepancies in three-way reconciliation. Testing must accurately record actual violations, discrepancies, blocking failures, remedial action, and retesting; these requirements must not be presented as results already achieved. Deliberately introducing noncompliant inputs tests whether the system can identify and handle them correctly; it does not in itself mean that the system has breached a rule.

Any comparison with manual methods will use comparable cases, rules, information conditions, and statistical definitions. It will compare simulated conversion, costs, and boundary controls, recording both improvements and areas without improvement. Simulated conversion is measured only by predefined events within the testbed and does not mean that actual clients or funds will necessarily enter the market.

VI. Simulation Testbed Methods and Design Features

The Simulation Testbed is intended to bring contract generation, role interaction, business execution, and Human–AI accountability into a single operational framework. These stages will influence one another, and the framework will be revised continuously in light of records and comparative results. The specific approaches are as follows.

(1) Link the Entire Process to Observe How Effects Propagate

Connect contract materials, inquiries and evidence collection, Compass measurement, investment choices, and subsequent contract performance into a continuous process. Information obtained, judgments made, and actions taken at each stage become conditions for subsequent work. Investor roles may also choose not to participate in accordance with their own objectives.

This design allows observation of how information gaps, judgment errors, or handoff problems affect subsequent outcomes and helps identify systemic issues that may be difficult to detect by examining an individual stage in isolation.

(2) Let Evidence Drive Inquiries and Business Needs Shape Information Collection

Inquiries develop progressively according to the industry, operating model, existing materials, and evidence gaps. The information to be collected depends on the facts that need to be established for the current business

activity. New information then supports factor judgments and updates to Compass readings under the applicable rules.

This approach can bring inquiries closer to business realities, reduce inapplicable reporting, and make evidence gaps easier to identify through interaction. Fields and evidence collection paths may be adjusted, but factor applicability and scoring continue to follow approved rules. Changes in methods are recorded separately, with their versions and supporting basis.

(3) Combine Autonomous Interaction with Human Governance

Simulated roles act according to their respective profiles, objectives, and visible information. The Orchestrator Agent and Subagents collaborate to organize specific tasks. Business choices emerge under the given conditions, and disagreements, refusals, waiting, and failures are allowed during operation.

Humans are responsible for rules, authorization, supervision, and necessary judgments, and AI employees execute tasks within defined boundaries. This enables the simulation to reveal behavioral differences among roles and to show whether existing rules, the division of work between humans and AI, and escalation mechanisms are sufficient to address those differences.

(4) Combine Central Scheduling with Distributed Execution

Manage task organization separately from execution devices: central services coordinate tasks and statuses, distributed devices perform assigned work, and execution processes and results are recorded centrally.

This allows multiple roles to participate at the same time and tasks to continue after device changes. The simulation can expand as resources permit. Resource utilization, execution stability, and cost-effectiveness will be monitored continuously during operation.

(5) Combine Modular Improvement with Version Comparisons

Organize case generation, task orchestration, inquiries and evidence collection, measurement, execution, and other stages as separate modules, so problems can be traced to specific modules and their inputs, rules, or operations. Revisions will identify what is being changed and the scope of the effects, followed by observation of results using comparable cases.

This approach can establish whether a problem arises from data, models, tools, processes, or business rules; clarify what should change and how it should be checked; and reveal the effects of local adjustments on other stages.

(6) Trace Evidence, Execution, and Accountability Together

Retain linked records of business facts, tasks, rule versions, AI execution, and human decisions, so that each result can be traced to its basis, the process that produced it, and its accountability arrangements.

These records support both the investigation of execution errors and the examination of whether rules are applicable, authorizations are sufficient, and judgments have an evidentiary basis. Problem resolution and rule revisions will be based on these facts, and results before and after a change can be reviewed and compared.

VII. Data, Participation, and Simulation Boundaries

Simulated business activities use synthetic contracts, operating data, role profiles, and virtual funds. They do not connect to actual trading accounts or handle real funds. Cases, pages, records, and results shown externally are all clearly labeled as simulated and managed separately from actual business materials.

Information within the Simulation Testbed is tiered by role, function, and permissions. Information used by individual roles is managed separately from data used for an overall view of the simulation; subsequent operating results must not be made available in advance as a basis for role decisions. Records of sources,

parameters, rule versions, and authorizations are retained for data generation, access, and use, enabling checks for information use beyond authorized permissions.

Actual information about human employees and external participants generated through registration, authorization, and accountability records is managed separately from synthetic business data. Its purposes, access scope, and handling arrangements are specified under applicable requirements.

External participation is conditional on internal readiness, resolution of key issues, and any necessary retesting. Specific arrangements will be communicated as work progresses. Before joining, participants must sign written participation terms defining the nature of the simulation activities, data use, and responsibility arrangements. Activities within the testbed do not constitute an investment offer. They do not collect actual expressions of investment interest, reserve investment allocations, accept investment funds, charge participation fees, or promise remuneration for participation in the simulation. Simulation participation creates no entitlement to actual business incentives or remuneration.

When Compass readings, simulated cash flows, and operating costs are shown externally, the samples, rules, and operating conditions will be disclosed alongside them. They will not be presented as investment recommendations or promises of returns. Participation in the simulation and publication of interim reports do not replace the approvals and necessary procedures required for subsequent actual business activities.

VIII. Phased Schedule, Progress Disclosure, and Expected Outputs

This round of internal simulation and testing is planned to begin on November 3, 2026, progressing through internal operation, interim reporting, and consolidation of results. The milestones below are planned dates. The pace will be adjusted according to system readiness, resolution of key issues, and retesting. External participation will be arranged separately under the conditions described in Section VII.

(1) Launch of Internal Simulation

This issue of the monthly publication is planned for release on November 3, 2026, alongside the launch of the internal simulation. Human and AI employees will participate in their assigned roles, with business processes, roles, and operating scenarios integrated progressively to produce initial operational records and an issue list.

(2) December Interim Report

NUMA plans to publish NUMA Operating System OS: The Three Engines and a simulation interim report around December 3, 2026, describing actual coverage, operational findings, and revision progress as of that date. Operational methods that are ready will continue to be tested during this simulation round.

(3) January Consolidation of Results and the Fourth White Paper

Around January 3, 2027, NUMA plans to consolidate the Operating System's structure, operational methods, and verification progress using the simulation records, producing its fourth white paper, NUMA Operating System OS: The Three Systems and Three Engines, and publishing simulation results. Tested, untested, and unresolved matters will be identified separately. Publication of a document will not be treated as completion of all validation.

Expected outputs include structural and accountability arrangements; rules and configurations; verification records for operational methods that have been integrated; issue and revision files; records of operating costs and resource use; and a list of validation statuses for each stage. Interim disclosures will identify actual coverage, control violations and discrepancies, blocking outcomes, simulated conversion, and stages that have not been completed successfully, together with the samples, conditions, and statistical basis. Reasons and follow-up arrangements will be given for delays or supplementary testing.

Passing a simulation and going live in production require separate approvals. Passing a particular simulation supports only the corresponding verification conclusion under the tested conditions. Rules or methods must still complete the relevant approvals and necessary procedures before use in production.

IX. Applicability and Limitations of Simulation Results

The Simulation Testbed can be used to observe processes, controls, and Human–AI collaboration under specified conditions and to identify operational problems, informing improvements to the Operating System and arrangements for subsequent validation. The scope of the conclusions depends on the samples, scenarios, rule versions, and integrated capabilities actually covered.

Synthetic cases and role profiles depend on industry assumptions and behavioral settings. They cannot capture the full range of actual participants' judgments, information asymmetries, constraints arising from their interests, or emotional responses. Increasing case numbers can extend test coverage, but whether the sample assumptions reflect reality still requires separate validation. Employees playing intermediary roles are also familiar with internal rules; their performance cannot directly represent the behavior of actual off-market participants.

Case generation, role interaction, and judgment stages may share correlated model biases. Agreement across multiple stages or smooth process execution may still leave the same type of problem undetected. Results should be interpreted in light of the original conditions, failed cases, and scenarios not covered.

Compressing business time helps observe processes such as contract performance, cash flows, and exit, but cannot fully capture long-term changes in business operations or adequately reproduce external constraints in actual trading, custody, and resolution. The effects of simplifications and stages not yet integrated on the conclusions will be explained with the interim results.

Results within the testbed are records of simulated operation and will not be described as actual business performance or forecasts of future returns. The applicability of Compass methods to actual operating data, system performance under real workloads, and the viability of the relevant business arrangements under real conditions all require progressive confirmation through subsequent validation. Simulation conclusions also do not replace admission, approval, and launch requirements for formal business activities.

Part IV | Continuity and Revisions

Continuity and Updates across the Three Previous White Papers

NUMA has published three successive white papers that build on one another, each with a different focus. The August 3 white paper proposed a market framework for investing in and financing contract cash flows, explaining the initial design of the market concept, contract structure, instruments, and governance. The September 3 white paper proposed measurement rules for the NUMA SHARP Twin Compasses and developed methods for registration, measurement, and portfolios. The October 3 white paper explained NUMA's six-instrument system and the organization, authorization, verification, and working methods of HAT Over MASS.

This issue restates market arrangements from the “First Investor”'s perspective and introduces the draft framework for the Operating System (OS) and the Simulation Testbed operating plan. The following items explain which arrangements continue, which have been supplemented, and which have been adjusted. The three white papers retain their respective accounts as published. Updates expressly identified in this issue apply only to the relevant matters and within the stated scope. Content not addressed here should still be read together with the original documents and their usage notes. Design proposals, verification arrangements, and actual operational status are described separately.

Market Design

1. Qualification, admission, and evidence requirements. The August 3 white paper required parties entering the market to have contributed value and assumed risk in advance under their Real World Contracts, and required contracts to be performing normally and generating actual cash flows. The September 3 white paper specified its application to existing contracts that were being performed and already had cash-flow records, and established platform admission standards. The October 3 white paper emphasized actual contract performance and third-party verifiability of cash flows. This issue carries forward those requirements: “First Investor” qualification continues to be based on a contribution made in advance, which need not be monetary. Contracts entering the market must meet platform admission standards as well as legal and regulatory requirements. This issue further explains how gaps in evidence and limitations in historical performance observations should be disclosed. Contracts must still meet the requirements for actual contributions, legal validity, ongoing actual performance, verifiable cash flows, and applicable admission criteria. Disclosure of insufficient records cannot substitute for meeting those criteria.
2. Cash-flow arrangements and their scope of application. Previous white papers described participation in distributions according to contractual cash-flow rights and arrangements for the timing of pooled cash flows. In light of the asset characteristics of Contract-Based Financing (CBF), this issue further explains the application, selection, and registration requirements for pro rata and priority-based cash-flow arrangements, and clarifies the relevant batches, recovery conditions, and protection of existing rights. The two arrangements do not replace one another. The mechanism for cancellation upon cash-flow distributions described in Section 8.4.4 of the September 3 white paper remains premised on pro rata distribution. The proportion of capital contributed, the share of cash-flow rights, the proportion allocated on a priority basis, and the priority of cash-flow distributions should be understood separately. Whether a recovery limit is set is agreed separately from the choice of distribution arrangement. Specific arrangements are governed by what is agreed, registered, and disclosed before the transaction. See Section V(1) of Market Design.
3. Continued retention of interests and special arrangements. Previous white papers explained the alignment of issuers' and investors' interests through retained interests and first-loss exposure. This issue further

explains that, as a general principle, the “First Investor” should continue to retain relevant interests to maintain alignment. Retention arrangements, including the proportion of rights retained, related commitments, and subsequent changes, are disclosed to Professional Investors. Under a priority-based arrangement, the retained portion ranks last, consistent with the meaning of a first-loss arrangement. Where special circumstances warrant a different arrangement, its conditions of application, disclosure, and continuing responsibilities will be discussed separately. See Section V(3) of Market Design.

4. Entry in batches and transfers of positions. The August 3 white paper proposed time-priority arrangements at the pool level for multiple contracts from the same issuer, and the September 3 white paper described how sold interests could be transferred. This issue further specifies batches, transaction terms, and registration under priority-based arrangements. Professional Investors enter in batches in order of purchase, with each batch's terms fixed and registered at the time of transaction. Adjustments by the “First Investor” to the untransferred portion apply only to investors entering afterward. Later investors may acquire an existing registered holder's entire position and assume its priority under rules disclosed in advance and specified before the first purchase. Holders may transfer the front-ranking part of their own interests. Any holder may transfer only a position that the holder owns, and subsequent transactions may not unilaterally change the existing rights of registered holders. See Section V(1) of Market Design and Section V of Market Progress.
5. Contract Pools. The August 3 white paper allowed a Protocol Contract to correspond to one or more Real World Contracts and described transparency requirements for bundling. This issue further explains the fixed scope, unified distribution, priorities, and release conditions of asset-side Contract Pools. A Contract Pool is distinct from the investor holdings portfolio described in the September 3 white paper. It does not create an exception to the principle of progressive listing or the number of listings. Contract-by-contract verification, concentration risk disclosure, and explanations of validation status continue to apply. See Section V(2) of Market Design.
6. Measurement of issuer performance and ongoing disclosure. The three previous white papers progressively developed arrangements for issuer profiles, performance records, contract and industry measurement, and ongoing observation. The August 3 white paper proposed issuer profiles, records, and reputational constraints. The September 3 white paper included issuer performance risk in the Core Compass and explained that relevant records could accumulate across contracts and periods. The October 3 white paper further developed measurement of contract structures and industry business operations. The “First Investor” in this issue is the issuer. This issue further explains how relevant performance scores in the Core Compass and Industry Compass relate to retention arrangements, Contract Performance Records, and the concentration of contracts held by the issuer, and specifies ongoing recording and disclosure of the corresponding information. See Section VI of Market Design.
7. On-market transactions and off-market intermediation. The August 3 white paper invited participation from brokers, financial advisers, independent researchers, and other ecosystem participants, and the September 3 white paper emphasized a two-sided market and access through licensed institutions. This issue further explains the off-market division of work. Through the Matching Hub it operates within the Information Matching System, NUMA continuously produces content, coordinates content direction and publication timing, provides support, and organizes activity. Intermediaries on both sides connect with clients, explain the market, and help them enter the on-market environment through prescribed channels; activities requiring a license are handled by the relevant licensed institutions. Listing, subscription, allocation, and transaction execution remain exclusively on-market activities, conducted by licensed trading venues and their members. Off-market activity is limited to intermediation and does not include transactions. Platform-Funded Early-Stage Market Incentives and Remuneration Allocation Support Services are described separately. NUMA is willing to devote part of its platform revenue to supporting and encouraging market participation; specific incentive methods and rules remain to be announced. Separately, the platform

provides accounting and transfer support for remuneration allocation based on agreements already reached among market participants. Where funds transfers are involved, the relevant licensed institutions execute them under valid authorization. See Section IV of Market Progress for the specific arrangements. These arrangements will be introduced after the relevant approvals are obtained and necessary procedures are completed. See Section VII of Market Design and Section IV of Market Progress.

Market Progress

8. Organizational structure of the Operating System (OS). The October 3 white paper proposed the HAT Over MASS method of Human–AI collaboration and maintained accountability for AI positions through assignment to named humans, departmental reporting lines, and task-flow lines. This issue's Operating System puts that collaborative method into practice through the Responsibility Matching System, Transaction Matching System, and Information Matching System, translating it into specific task, process, and accountability arrangements. The Responsibility Matching System uses Human–AI Cockpits as organizational units to assign task supervision, authorization, and decisions to named employees on duty. The Transaction Matching System carries business and support work across Five Zones and Eleven Functional Blocks, decomposing work internally into Functional, Main Task, and Microtask Layers, with AI employees executing tasks under approved rules and within their authorizations. The Information Matching System organizes content support and intermediation on both sides. The three Systems describe the system's constituent parts and their respective responsibilities; the three layers of work decomposition within the Transaction Matching System describe how work is organized internally. These two sets of hierarchical names are not interchangeable. When an employee leaves duty, the AI agents, microtasks, and transactions in that employee's cockpit are reassigned before the original cockpit is dissolved, so that no AI agent is ever without a responsible human. The principle of human accountability and AI employee execution remains unchanged. See Sections I and II of Market Progress.
9. Classification and collaboration of AI employees. The October 3 white paper organized AI employees into four functions: administration, business, control, and deployment support. This issue describes AI employee classification and collaboration according to how tasks are assigned. Within the Transaction Matching System, AI employees are divided into the Microtask AI Agent Cluster, Industry-Expert AI Agent Cluster, and Client-Dedicated AI Agent Cluster. These respectively handle execution, sign-off preparation, and general judgment; industry-factor assessment; and dedicated services that follow clients. Human–AI Cockpits are separately equipped with Supervisory AI for Microtasks and Supervisory AI for Transaction Tracking. Each AI agent within the Transaction Matching System reports to a cockpit AI agent, and each cockpit AI agent reports to a named human. The two classifications use different criteria and cannot be mapped one to one. See Sections II and III of Market Progress.
10. How Industry Compass work is assigned. The October 3 white paper explained task decomposition, team formation, and the seven-station workflow for the Industry Compass. This issue establishes industry-factor assessment microtasks for individual industries within the OS. Teams of AI agents from different large models undertake industry research, professional judgment, and error identification, supervised through Human–AI Cockpits by employees in the relevant fields. This assignment arrangement does not replace the earlier workflow, independent verification, or accountability arrangements. See Section III of Market Progress.
11. Registration and Compass validation status. The September 3 white paper proposed the Contract Register and Order Register, and the October 3 white paper proposed six instruments. This issue assigns these functions to functional blocks in the Transaction Matching System: L2 handles contract registration; R3 registers orders, positions, priorities, and Contract Pools; and C1 handles Compass readings and updates. Compass readings are also labeled by degree of validation in four statuses: run only in simulation, replayed against historical data, tested on actual contracts, and revised using real-world outcomes. Predictions at

market entry are recorded and compared with actual results when the contract is fully settled. Readings are updated according to the nature of the data and new evidence; daily checks do not mean daily changes to values. Passing a simulation does not mean validation on actual contracts is complete. Historical records and their versions are retained. See Section III of Market Progress.

12. Funds handling and reconciliation. Agenda Item 2 in Chapter XI of the September 3 white paper proposed research into custody of investor funds, transfer of transaction proceeds, and cross-checking funds records against the Registers. As an interim development of that agenda, this issue proposes the following controls. Client funds are held only at licensed custodian banks. The OS transmits funds instructions under authorization and reconciles the results. The Registers, custody accounts, and distribution ledgers undergo daily three-way reconciliation, with zero discrepancies required. New business is suspended if discrepancies are not zero. Decisions involving funds and changes to payment or receiving accounts require sign-off by two humans who are not in a direct reporting relationship. Payments already approved by dual sign-off may be executed within valid authorization. The proposal of these control requirements should not be taken to imply that all relevant capabilities have been implemented or have passed operational validation. Their status will be described according to actual progress. See Sections II and V of Market Progress.
13. Human authorization and sign-off. The October 3 white paper proposed that humans retain control of rules, authorization, and ultimate accountability, subject to checks through managing people, checking systems, and reviewing standards. This issue further proposes three forms of sign-off - Sign-Off by Task Category, Case-by-Case Sign-Off, and Exception-Only Sign-Off - and classifies the matters employees may handle into four tiers according to professional field and risk. It also requires segregation of processing and release duties and sign-off by two humans for decisions involving funds. Approved tasks that meet the conditions and remain within authorization may be executed under the rules without repeated approval of every action. General authorization does not waive applicable requirements for case-by-case, dual, or statutory sign-off. See Section II of Market Progress.

Market Preview

14. Simulation operating plan and phased deliverables. Chapter XVII of the October 3 white paper proposed continuous operation in a simulated environment. This issue proposes a Simulation Testbed operating plan. This round of internal simulation and testing is planned to begin on November 3, 2026. All employees will supervise microtasks and transactions through their Human–AI Cockpits, and employees will play the roles of off-market intermediaries. NUMA plans to publish NUMA Operating System OS: The Three Engines and a simulation interim report around December 3, 2026, describing actual coverage, operational findings, and revision progress as of that date. With a target publication date around January 3, 2027, NUMA plans to consolidate the Operating System's structure, operational methods, and verification progress using the simulation records, producing its fourth white paper, NUMA Operating System OS: The Three Systems and Three Engines, and publishing simulation results. Tested, untested, and unresolved matters will be identified separately. External participation depends on readiness, resolution of key issues, any necessary retesting, and other conditions, and is not tied to a specific publication date. Passing a simulation and going live in production require separate approvals. The Human–AI Cockpit in this issue is the work unit through which employees supervise AI, a different concept from the simulation cockpit described in the October 3 white paper for participants to experience. See Section VIII of Market Preview.

Version usage principle: Updates identified in this issue apply only to the relevant matters and within the stated scope. Matters not addressed should still be read together with the original documents and their usage notes. The status of design proposals, system development, internal verification, and actual business operation is described separately. Future arrangements do not constitute commitments regarding launch dates or investment outcomes.

Appendix I

Terminology and Historical Documents

This appendix lists specialized terms used in the final main text for contracts, operational organization, AI collaboration, authorization controls, and simulation evaluation. Definitions are provided to assist reading. Specific rights, business rules, numerical definitions, and conditions of application are governed by the corresponding main text and notices.

Core Contract and Market Terminology

Chinese Term	English Full Name and Abbreviation	Definition
“第一投资人”	“First Investor”	A party that contributes resources, capabilities, or funds in advance, obtains subsequent returns under a contract, and then brings the corresponding contractual cash flows to market.
专业投资者	Professional Investor（PI）	An investor, which may be an institution, that meets applicable qualifications and participates through the relevant licensed channels. The relevant licensed institutions verify client identity and qualifications. Platform access, public reading, or simulation participation does not replace the conditions for formal investment participation.
原生合同	Real World Contract（RWC）	The pre-existing contract between the “First Investor” and the underlying business operator, specifying the contribution made in advance and the corresponding returns. Entry into the market does not change its nature.
原生合同对手方	Real World Contract Counterparty（RWCC）	The underlying business operator that enters into a Real World Contract with the “First Investor”, conducts business under the original contract, and fulfills its agreed obligations.
挂牌合同	Protocol Contract（PC）	The listing document signed by the “First Investor” in respect of cash flows from a Real World Contract, specifying the scope of cash flows brought to market, the distribution method, and retention arrangements.
交易合同	Transaction Contract（TC）	The contract signed when a Professional Investor participates, recording the cash-flow rights acquired and the investor's rights and obligations.
合同余额	Contract Value（CV）	The Contract Value described in the September 3 white paper is retained; in this issue, it indicates the size of an investor's purchase relative to the contract.
等比例回收	Pro Rata Cash-Flow Distribution	Each period's cash flows are distributed simultaneously to all holders according to the shares of rights agreed in advance.
优先顺位回收	Priority-Based Cash-Flow Distribution	Cash flows allocated to priority distribution go first to earlier-ranking positions in the registered order. Where a recovery cap is set, distributions proceed to subsequent positions in sequence once the agreed cap is reached. For positions without a recovery cap, distribution and exit conditions are separately agreed and registered before the transaction. Returns are not guaranteed.
顺位	Cash-Flow Distribution Priority	A position in the order of cash-flow distributions recorded in the Register. The rights, recovery limits, and risks of each position depend on valid agreements and actual cash flows.

Chinese Term	English Full Name and Abbreviation	Definition
合同组合	Contract Pool	An arrangement under which a “First Investor” holding many contracts of the same type brings a batch of contracts to market as a whole. The contracts are specified and fixed when included, with no subsequent replacements or additions, and their cash flows are distributed under a unified arrangement.
“第一投资人”保留比例	“First Investor”'s Retained Share of Cash-Flow Rights	The share retained by the “First Investor” after transferring cash-flow rights. It must be disclosed together with the method used to select the retained portion, enabling investors to make their own judgments.
履约记录	Contract Performance Record	Records of performance, cash flows, and exception handling accumulated continuously in relation to actual contracts and their processing, serving as the basis for measurement and disclosure.
双罗盘	NUMA SHARP Twin Compasses	The Core Compass and Industry Compass jointly provide contract measurement. Returns and risks are assessed separately, and gaps are retained where evidence is insufficient.
通用罗盘	Core Compass	A measurement tool for observing contract factors common across industries.
行业罗盘	Industry Compass	A measurement tool for observing business factors in a specific industry. Its factors, evidence, and scope of application must match the relevant industry and scenario.
罗盘验证状态	Compass Validation Status	An indication of the evidence, validation status, and scope of application underlying measurement results. Unvalidated readings must not be treated as validated results.
IP 授权	Intellectual Property Licensing	Licensing of intellectual property, such as Real World Contract arrangements based on the use of trademarks, brands, or other intellectual property and the corresponding cash flows.
登记簿	Register	An ongoing record of contracts, rights, distribution priorities, and relevant valid agreements. Cash flows are handled according to the registered arrangements, and changes must retain their supporting basis and historical trail.

Operating System and Organizational Structure

Chinese Term	English Full Name and Abbreviation	Definition
牛码新市运营体系	NUMA Operating System（OS）	The operating system that translates NUMA's business rules into tasks, processes, controls, records, and Human–AI accountability arrangements. It puts the HAT Over MASS method of Human–AI collaboration into practice through collaboration among the Responsibility Matching System, Transaction Matching System, and Information Matching System. This issue introduces the framework, which remains subject to simulation testing and revision. It is not equated with the name of the HAT Over MASS theory.
人工智能	Artificial Intelligence（AI）	In this issue, AI employees execute tasks under approved rules and within their authorizations. Named humans undertake necessary judgments and decisions and bear the corresponding responsibility.
信息技术	Information Technology（IT）	In this issue, support work includes systems, infrastructure, permissions, interfaces, event records, and task pipelines.

Chinese Term	English Full Name and Abbreviation	Definition
帽子军团	HAT Over MASS	A theory and method of Human–AI collaboration combining HAT and MASS. It emphasizes human control of rules, authorization, and accountability, with AI agents collaborating on execution through microtasks. The NUMA Operating System puts this method into practice; HAT Over MASS is not a synonym for the OS.
帽子：人管控、AI 运作的信任机制	Human-Commanded, AI-Operated, Trust（HAT）	A trust mechanism that is human-commanded and AI-operated: humans establish and revise rules and bear responsibility; AI executes within its authorization.
军团：微任务智能体与简单且短时的任务	Micro Agents · Simple & Short Tasks（MASS）	Business work is divided into microtasks with clear boundaries that can be verified independently, then completed through collaboration among many agents.
场控引擎、场内引擎与场外引擎	Responsibility Matching Engine (RME); Transaction Matching Engine (TME); Information Matching Engine (IME)	Algorithms running on the Three Systems, planned to be introduced around December 3, 2026. Actual integration and simulation verification progress will depend on readiness.
ABCD 四类职能	Administration, Business, Control, Deployment	The four functions described in the historical white papers: administration, business execution, control and compliance, and deployment support. Their classification differs from that of the three AI clusters in this issue, and they cannot be substituted for one another item by item.
场控基座	Responsibility Matching System（RMS）	The System that organizes task supervision, qualifications, authorization, team configuration, decisions, and continuity of accountability. Its specific organizational form is the Human–AI Cockpit, through which responsibility is assigned to named employees on duty. It is an overall accountability arrangement and cannot be defined simply as a single cockpit.
场内基座	Transaction Matching System（TMS）	The System that carries on-market business and support work, decomposing work across Five Zones and Eleven Functional Blocks and through the Functional, Main Task, and Microtask Layers. AI employees execute tasks under approved rules and within their authorizations, with connections to external licensed business interfaces.
场外基座	Information Matching System（IMS）	The System that organizes central content support, intermediation on both sides, information communication, and handoffs for market entry. It helps the asset and funding sides understand the market and participate through prescribed channels. Content and operations requiring a license are handled by the relevant licensed institutions.
五区十一块	Five Zones and Eleven Functional Blocks	The on-market layout comprises top, left, right, central, and bottom zones, containing Central Control, three financing-side blocks, three investment-side blocks, three central blocks, and Infrastructure. Numbers identify locations and do not indicate processing order.
功能层、主任务层与微任务层	Functional, Main Task, and Microtask Layers	The internal decomposition of work within the Transaction Matching System: eleven functional blocks organize responsibilities, which are divided into 34 main tasks and then 1,530 microtasks. AI agents are assigned at the microtask layer.
T 总控	T: Central Control	The functional block responsible for publishing approved rules, scheduling tasks, coordinating across modules, handling exceptions, and conducting sample checks and backtesting.

Chinese Term	English Full Name and Abbreviation	Definition
L1 “第一投资人”服务	L1: “First Investor” Services	Responsible for business inquiries, preparation of materials, and ongoing communication. It also generates and configures Client-Dedicated AI during the contract term, registers sources, and helps collect contracts and evidence of contributions.
L2 合同登记	L2: Contract Registration	Responsible for checking and registering materials on contract elements, versions, signing, and rights arrangements, and preserving the correspondence between Real World Contracts and listing documents.
L3 认定与准入	L3: Qualification and Admission	Verifies “First Investor” qualifications, relevant contract admission conditions, and review opinions, and creates admission records. Internal registration is not equivalent to listing on a licensed trading venue.
R1 投资者服务	R1: Investor Services	Receives Professional Investors through licensed institutions and provides rule explanations, business status information, and notification services. Reception does not mean that subscriptions are open.
R2 投资者准入	R2: Investor Admission	Handles platform admission, permissions, and limits based on identity, qualification, and related verification results returned by licensed intermediaries.
R3 组合登记	R3: Portfolio Registration	A Single Writer registers investor positions, distribution priorities, and Contract Pools based on orders, trade results, and funds reconciliation results.
C1 信息流	C1: Information Flow	Conducts Compass measurement based on contract and operating evidence, updating readings according to the nature of the data and new evidence. Daily checks do not mean daily changes to values. It supports disclosure and scanning, stating the evidentiary basis, validation status, and scope of application.
C2 交易场所接口	C2: Trading Venue Interface	Prepares listing materials, transmits authorized instructions, and reconciles transaction results. Listing, allocation, and transaction execution are performed by the relevant licensed institutions.
C3 资金流	C3: Funds Flow	Transmits authorized funds instructions for settlement, refunds, and cash-flow distributions, and reconciles custody confirmations, contracts, and position records. Client funds are held by licensed custodian banks.
B 基础设施	B: Infrastructure	Provides support for IT systems, risk control and compliance, finance, legal matters, audit, and other functions. Access, records, and permission management follow the relevant requirements.
人机驾驶舱	Human–AI Cockpit	An organizational unit consisting of a named employee on duty and assigned supervisory AI, implementing task supervision and human accountability. It is the specific organizational form of the Responsibility Matching System. One cockpit is planned for each employee on duty. AI tracks tasks, industry teams, and transactions in progress; humans supervise, judge, sign off, and handle exceptions within their qualifications and valid authorizations. Personnel changes follow existing handover mechanisms. Assignment or temporary assumption of work does not automatically expand permissions.
标签库	Tag Repository	A planned repository using standardized tags to describe personnel, microtasks, contracts, investors, industries, content, and other objects, supporting responsibility allocation, business matching, and information continuity across the Systems. Specific tags and their scope of use remain to be refined through design and simulation testing.

Chinese Term	English Full Name and Abbreviation	Definition
市场初期的平台激励	Platform-Funded Early-Stage Market Incentives	During the early stages of market development, NUMA is willing to devote part of its platform revenue to supporting and encouraging market participation. Specific incentive methods and rules remain to be announced.
报酬执行服务	Remuneration Allocation Support Services	Accounting and transfer support for remuneration allocation based on agreements already reached among market participants. Where funds transfers are involved, the relevant licensed institutions execute them under valid authorization. Clients determine which registered intermediaries receive the remuneration, with specific triggers and release arrangements governed by applicable rules. Investment-side remuneration is paid to licensed institutions. No remuneration is earned for recruiting other intermediaries.

Microtasks and AI Collaboration

Chinese Term	English Full Name and Abbreviation	Definition
微任务	Microtask	A unit of work that can be executed and verified independently, with defined inputs, outputs, permissions, end conditions, and verification methods. Each fixed work position performs only one task.
执行类、签字类与判断类微任务	Execution, Sign-Off Preparation, and Judgment Microtasks	Three categories of microtasks handling, respectively, rule-based execution such as checking and calculation; sign-off preparation; and general judgment or industry-factor assessment. Each microtask belongs to only one category.
微任务 AI（驾驶舱督导）	Supervisory AI for Microtasks	AI within a cockpit that supervises a fixed group of microtasks, tracks outputs, queues, returns for rework, and errors, and delivers matters requiring human decisions to the responsible person.
交易跟踪 AI	Supervisory AI for Transaction Tracking	Supervisory AI within a cockpit that continuously tracks each transaction and its Client-Dedicated AI until the transaction is completed or terminated. It may track multiple Client-Dedicated AI agents simultaneously.
微任务 AI 集群	Microtask AI Agent Cluster	The AI cluster within the Transaction Matching System that handles execution, sign-off preparation, and general judgment. It differs from the supervisory role of the same name within a cockpit.
行业专家 AI 集群	Industry-Expert AI Agent Cluster	The AI cluster responsible for industry-factor assessment. This design assigns a ten-member team to each of 500 industries, totaling 5,000 AI agents on standing task assignments.
客户专属 AI 集群	Client-Dedicated AI Agent Cluster	Configured by client, with dedicated AI for each “First Investor” and Professional Investor, providing continuous service from registration and following transactions. These agents are not permanently assigned to a particular microtask.
常驻 AI	AI Agents on Standing Task Assignments	AI agents within the Transaction Matching System assigned to microtasks on an ongoing basis. This design provides for 7,168 such agents, excluding Client-Dedicated AI that grows with client numbers and cockpit supervisory AI. This is not yet the scale of actual operations.
独立判断与检验	Independent Judgment and Verification	For judgment tasks, different large models make separate evidence-based judgments and retain their supporting basis. Independent verification is still required when models agree, with human verification sought when necessary.

Chinese Term	English Full Name and Abbreviation	Definition
数字同权	Delegated Authorization for AI Agents	An AI position's valid business authorization rests on the authorization held by a named human. Data access and operational permissions are limited to the corresponding position and task. Assuming responsibility for a position does not increase a human's approval authority.
动态编组	Dynamic Team Configuration	Microtasks, transactions, and industry teams are assigned according to employee qualifications, authorizations, duty status, and conflict-of-duty recusal requirements. Assignments are redistributed when personnel change, and original records are retained continuously.
任务四级	Four-Tier Task Classification	A four-tier classification ranging from matters governed by clear rules and easy to verify, through matters requiring experience and those requiring professional judgment, to major matters. Task classification does not replace professional qualifications or approval authority.
主控 Agent 与子 Agent	Orchestrator Agent and Subagents	The Orchestrator Agent orchestrates tasks and invokes Subagents, which undertake specific tasks. Invocations, tool results, rule versions, and records are retained during operation.
AI 面试	AI Interview	AI interaction in the financing-side materials and question-and-answer process, used to collect, check, and prepare materials. The process and results remain subject to established rules and human verification.
AI 岗位、任务标识与单次执行标识	AI Organizational Position; Task Identifier; Run Identifier	An AI position is a stable organizational identity; specific tasks and individual runs have separate identifiers. Related records link inputs and outputs, rule versions, and responsible persons, enabling a result to be traced back to the processing that produced it.

Authorization Controls and Records

Chinese Term	English Full Name and Abbreviation	Definition
具名责任	Named Human Accountability	Identifiable humans with the appropriate qualifications and authorizations bear responsibility for judgments, decisions, supervision, and exception handling within the specified scope, with records linking responsibilities to those individuals.
一类一签	Sign-Off by Task Category	Approval of rules, standard terms, and methods as the basis for executing tasks of the same category. Execution must remain within valid authorization, time limits, amounts, and exception conditions.
一单一签	Case-by-Case Sign-Off	Approval based on the materials for an individual case where a separate decision is required.
例外才签	Exception-Only Sign-Off	A specific decision on matters not covered by rules or that depart from the normal path.
签字点	Human Sign-Off Point	A point at which a human with the appropriate qualifications and authorization makes a decision, signs off, and bears responsibility.
经办与放行分离	Segregation of Processing and Release Duties	The processor is responsible for handling the matter and producing outputs, while review and release are undertaken by the relevant independent personnel. The same person switching cockpits does not constitute independent review.

Chinese Term	English Full Name and Abbreviation	Definition
三方对账	Three-Way Reconciliation	Under the funds flow arrangements in this issue, custody confirmations, contracts, and position records are reconciled, with authorizations and transaction results used to identify and resolve discrepancies.
事件账与审计记录	Event Ledger and Audit Records	Linked records of tasks, rule versions, inputs and outputs, executors and signatories, check results, and exception handling, retaining the command relationships in effect at the time and revision trails.
已交付、已完成与已归档	Delivered, Completed, and Archived	Separate verification that outputs have been delivered, tasks have been concluded, and evidence has been retained. When an output undergoes a material change, the continued validity of affected checks and approvals is reassessed. Any required rechecking, decisions, or sign-off follow the applicable rules.
合格合同全程成本	Lifecycle Processing Cost per Contract Meeting Completion Criteria	The principal measure of total cost from market entry to exit. Meeting completion criteria requires business processing to be concluded, evidence to be complete, controls to be effective, accounting records to agree with positions, and related issues to be closed.
第 0 号关口	Gate 0	The gate that checks whether the “First Investor” has made an actual contribution in advance, during financing-side entry and ongoing follow-up. Relevant checks and corrections must have traceable records.
持牌分工	Allocation of Functions among Licensed Institutions	Stages requiring a license, such as trading, brokerage, and custody of funds, are handled by the relevant licensed institutions within their licensed scopes. Internal OS statuses cannot replace their valid business results.
双签	Dual Human Sign-Off	Decisions involving funds require sign-off by two humans who are not in a direct reporting relationship. Disbursement instructions must be transmitted only after valid dual sign-off.
唯一写入者	Single Writer	Under the R3 ledger arrangements in this issue, a Single Writer registers orders, positions, priorities, and Contract Pools, ensuring that records are updated through the prescribed path.

Simulation Methods and Observation Dimensions

Chinese Term	English Full Name and Abbreviation	Definition
模拟舱	Simulation Testbed	An operating environment isolated from actual business activities, used to test processes, controls, and Human–AI collaboration under OS rules. In this issue, it does not connect to real accounts, handle real funds, or conduct real transactions.
基于主体的仿真	Agent-Based Simulation	Profiles, objectives, and behavioral rules are set for different external roles to observe business events arising from their interactions with other roles and the OS.
合成数据	Synthetic Data	Contract, operating, cash-flow, and role behavior data generated according to industry and scenario parameters. Parameters, distributions during operation, assumptions, and the basis for revisions are retained continuously.
压缩时钟	Compressed Simulation Clock	Contract durations are compressed proportionately to observe multiple complete cash-flow cycles and low-frequency events within a shorter operating period, while preserving proportions and event order.

Chinese Term	English Full Name and Abbreviation	Definition
情景压力测试	Scenario-Based Stress Testing	Adverse scenarios, such as sudden increases in intake volumes, staff absences, AI judgment errors, and declining operating cash flows, are introduced to observe process resilience under pressure, controls, and exception handling.
回放与对照	Event Replay and Rule-Version Comparison	Processes are reconstructed using retained events, inputs, and rule versions. Performance under different rule versions is compared to distinguish the effects of rule revisions from differences in samples.
黑灯工厂	Lights-Out Factory	A figurative term in this issue for AI task pipelines running continuously. Continuous 7×24-hour operation is a design objective and does not mean that the rules and conditions have been validated.
流程完整性与任务质量	Process Completeness and Task Quality	Observe whether business processes connect, whether tasks meet completion requirements, and whether returns for rework, suspensions, terminations, and exceptions can be handled effectively.
运行效率与资源成本	Operational Efficiency and Resource Costs	Observe the time required for business processing, waiting, handoffs, and rework, together with the model, tool, human, and other resources used.
业务信息衔接	Business Information Flow and Consistency	Observe whether the information needed by the financing side, the investment side, and internal positions remains complete, consistent, and usable as it is collected, registered, disclosed, and used.
证据充分性与度量适用性	Evidence Sufficiency and Applicability of Measurement Methods	Observe whether the evidence required by the Compasses can be obtained, verified, and used to support scoring, and whether the factors and methods are applicable to the relevant business conditions.
人机协作与授权适配	Human–AI Collaboration and Authorization Alignment	Observe whether human and AI employees' responsibilities, capabilities, and authorizations align, and whether human judgment, exception escalation, and task handovers work effectively.
逐级处置与责任升级	Tiered Issue Resolution and Escalation	Issues are handled through successive stages of entry checks, AI self-checks and peer reviews, human takeover, and rule revisions, with escalation according to authority and the nature of the issue. After handling, the decision to continue, suspend, or terminate is made explicit, and the basis, results, and subsequent responsibilities are retained. Approved rule revisions are tested again in simulation.

Note on translated terms: The English column includes historical names, common usage, and descriptive translations used in this publication. Meanings are governed by the Chinese definitions. Entries without confirmed proper names use descriptive translations and are not presented as external standard terminology. Specific statistical definitions for the simulation observation dimensions will be established as capabilities are integrated.

Historical Documents

Publication Date	Document Title and Subject
August 3, 2026	NUMA White Paper - Micro Connect Real World Contract Investment and Financing Master Agreement Market framework, three-layer contract structure, and basic rules. Refer to this document for the foundational arrangements for participation in contract cash flows. Building on that foundation, Market Design in this issue further explains the process of bringing contracts to market through the “First Investor”'s advance contribution and ongoing follow-up.

Publication Date	Document Title and Subject
September 3, 2026	NUMA White Paper - From Contract Cash Flows to Allocable Assets Twin Compasses measurement methods and the two-sided market. Refer to this document for the Core Compass and Industry Compass and the rules for assessing returns and risks separately. The disclosures in this issue concerning the “First Investor”'s Contract Performance Record, retained share, and concentration risk can be read alongside that edition.
October 3, 2026	HAT Over MASS - Microtasks across an Army of Agents. Measurement Infrastructure for the Contract Economy: Theory, Factory, and Samples. Refer to this document for HAT Over MASS organization and verification methods and the boundaries between AI execution and named human accountability. Market Progress and Market Preview in this issue describe the draft OS framework and proposed simulation arrangements, respectively.

Appendix II

Main Tasks and Microtask Examples in the Transaction Matching System

This appendix lists all 34 main tasks across the 11 functional blocks. Numbers in parentheses indicate the number of microtasks under each main task. 5 microtask examples are selected for each main task to illustrate the work involved. These examples are not a complete list and do not cover every branch.

AI configurations by category and human sign-off requirements follow the main text. C1 Compass Readings includes 500 industry-factor assessment microtasks established by industry, which are not listed industry by industry here.

Top Zone · Central Control

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Rule Development (30)	Receive requirements; cross-check against white papers; check for conflicts; approve launch; archive revisions
Central Scheduling (40)	Queue tasks; match qualifications; assess irreversibility; halt operations; archive incidents
Backtesting (30)	Load samples; lock the rule version; attribute discrepancies; calculate blocking rates; store learning records
Consultation and Decision (30)	Generate consultation requests; match positions; check completeness; bind material versions; archive matters

Left Zone · L1: “First Investor” Services

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Reception and Identification (30)	Register sources; link off-market intermediaries; classify business models; compare against exclusion lists; notify intermediaries
Dedicated Assistant (30)	Select industry templates; clone assistants; check configurations; verify images; update files
Post-Investment Follow-Up (30)	Send cash-flow reminders; collect operating data; make preliminary judgments on exceptions; update disclosures; update files

Left Zone · L2: Contract Registration

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Contract Element Collection (30)	Generate interview checklists; assign interview tasks; check interview quality; extract underlying business operator information; record evidence links
Three-Way Validation (30)	Compare contracts with transaction records; compare contracts with business registration records; conclude on consistency; record data sources; inform the “First Investor”
Listing Documents and Registration (30)	Determine the scope of cash flows available for sharing; enter the retained share and retention method; check structural compliance; conduct legal review; retain correction trails

Left Zone · L3: Qualification and Admission

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
“First Investor” Qualification (Gate 0) (30)	Receive applications; compare against exclusion lists; determine routing; qualification conclusions; prepare sample verification reports
Contract Admission (30)	Verify business registration; check credit status; conclude on look-through checks; admission conclusions; refer for calculation of interests available for sale
Interests Available for Sale (30)	Read remaining receivables; read the proportion to be transferred; recalculate independently; interests available for sale that day; retain change trails

Right Zone · R1: Investor Services

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Reception and Education (30)	Verify licensed-institution sources; register sources; assess understanding; provide supplementary explanations; retain records of interest
Dedicated Assistant (30)	Identify investor types; clone assistants; check configurations; check wording for promises; analyze complaints
Notifications (30)	Send transaction execution notices; send non-execution notices; issue external exception notices; review reports; archive reports

Right Zone · R2: Investor Admission

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Licensed Institution Access (30)	Verify licenses; verify membership; assess access; grant institutional access; retain surveillance trails
Qualification Cross-Checks (30)	Receive identity verification results; receive Professional Investor classifications; assess consistency of results; identify related parties; retain declaration trails
Activation and Limits (30)	Read limit rules; calculate limits; review limits; platform limits; retain closure trails

Right Zone · R3: Portfolio Registration

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Orders (30)	Receive settings; validate conditions; process order cancellations; assess abnormal orders; archive surveillance records
Positions and Priorities (30)	Read confirmations of funds received; record positions; review registration; track notice periods; update the pool priority table
Snapshots and Statements (30)	Generate snapshots; calculate concentration; review snapshots; append correction records; retain correction trails

Central Zone · C1: Information Flow

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Compass Readings (530)	Generate evidence lists; retrieve evidence; assess evidence credibility; publish readings; assess industry factors (established by industry, totaling 500, one per industry)

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Daily Checks and Evidence-Based Updates (30)	Read cash-flow transaction records; read receivables schedules; assess discrepancies; determine overdue status; retain determination trails
Scanning and Disclosure (30)	Read investor conditions; validate conditions; check disclosure completeness; verify the scope of Professional Investors; retain status trails

Central Zone · C2: Trading Venue Interface

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Listing Coordination (30)	Generate material checklists; generate contract cards; check material compliance; listing materials; archive document media
Orders and Allocation (30)	Receive through member channels; validate order formats; assess allocation results; reconcile allocation results; archive surveillance records
Trade Reports (30)	Receive transaction results; generate reports for both sides; assess report discrepancies; assess write-back discrepancies; retain invalidation trails

Central Zone · C3: Funds Flow

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Settlement (30)	Generate settlement instructions; verify custody accounts; assess consistency of funds received; disbursement instructions (dual sign-off); retain refund trails
Reconciliation (30)	Read Register balances; read custody account balances; determine whether discrepancies are zero; retain original records; retain verification trails
Cash-Flow Distribution (30)	Read cash-flow receipts; aggregate by contract; recalculate clearing amounts; distribution instructions (dual sign-off); archive records of defaults

Bottom Zone · Infrastructure

Main Task / Microtask Count	Microtask Examples (5 Selected per Main Task)
Risk Control and Compliance (30)	Read limits; validate limits; assess conflicts; suspicious transaction reports; archive submissions
IT Systems (30)	Record events; verify append-only recording; assess exceptions; check data boundaries for external models; retain operations and maintenance trails
Finance, Legal, and Audit (30)	Generate vouchers; post entries automatically; assess accounting discrepancies; independent verification conclusions; archive verification records

Suites 2105–2108
Two Exchange Square
8 Connaught Place, Central, Hong Kong

Micro Connect Group Website: www.microconnect.com
Email: info@microconnect.com