Universal Library of Multidisciplinary Book

ISSN: 3064-9986 | Volume 3, Issue 1

Open Access | PP: 83-92

DOI: https://doi.org/10.70315/uloap.ulmdi.2026.0301008

Operational Structure of Early Drug Discovery Organizations

Nataliia Koval

Abstract

The methodology examines the organizational architecture of the early stage of drug discovery as a governed system for the accelerated reduction of scientific uncertainty. The relevance of this methodology is determined by the rising cost of drug development, the high cost of operational failures, and the need to transform Early Drug Discovery from a collection of fragmented experiments into an integrated circuit of reproducible learning. The aim of this work is to formulate theoretically and practically grounded principles for the design and reorganization of R&D structures that ensure synchronization among scientific hypotheses, execution, data, managerial decision-making, and capital allocation. The novelty of the methodology lies in the integration of systems engineering, the theory of constraints, lean manufacturing, and QMS into a unified operational model of a Discovery organization, incorporating a five-layer governance framework, Handover Gate interface gateways, the ELN/LIMS linkage, and principles for managing distributed teams and CROs as nodes of a single research system. The conclusions are that the effectiveness of early Drug Discovery is determined by the quality of scientific ideas, the resilience of cross-functional interfaces, the traceability of data, the cadence of DMTA cycles, and the protection of critical flow constraints. The methodology will be useful to R&D leaders, operational executives, architects of research processes, and managers of biotechnology organizations.

Keywords: *Early Drug Discovery, Operational Model, DMTA Cycle, Handover Gate, ELN, LIMS, QMS.*

INTRODUCTION

The traditional perception of the drug discovery process as a sequence of brilliant insights generated by individual scientists no longer corresponds to the realities of an industry in which the cost of developing a single approved drug usually exceeds USD 3.5 billion (Fernald et al., 2024). Under conditions in which the scientific uncertainty of biology is irreducible, the only lever for improving efficiency is the reduction of operational instability. A paradigm shift is therefore observable: a transition from viewing Drug Discovery as a series of discrete experiments toward understanding it as a functional system for the conversion of capital (Kim et al., 2022). In this model, investments are directed toward executing research cycles whose outputs are structured data that generate intellectual property and increase the company's market value.

The advantage of leading biotechnology companies today lies not only in the speed of synthesis or the power of algorithms, but also in disciplined execution systems (Huanbutta et al., 2024). Experience demonstrates that organizations rarely fail because of bad science (Sampathkumar & Kerwin, 2023). Far more often, collapse is caused by the fragility of the operational system, which is incapable of ensuring the reliable transfer of context and responsibility across functional node boundaries (Galson et al., 2020). When the elements of the conversion chain from hypothesis to decision are disjointed, capital dissipates into operational noise, and learning cycles become unjustifiably protracted, thereby diminishing the organization's financial runway (Galson et al., 2020).

The purpose of this methodology is to provide theoretically grounded principles for the construction and reorganization of R&D structures, ensuring that scientific work advances at maximum speed while all system elements function in synchrony. The central emphasis is placed on creating an operational framework capable of sustaining a high density of learning per unit of time and capital. The methodology integrates systems engineering, the theory of constraints, lean production principles, and quality management systems within the specific context of biotechnology research. Implementing the proposed approaches enables an organization to transform from a collection of laboratories into a single engine of learning, in which every action is directed toward the predictable reduction of project uncertainty.

CHAPTER 1. STRUCTURAL FOUNDATION AND THE DESIGN OF CROSS-FUNCTIONAL INTERFACES

Operational Foundation

An effective organization at the Early Drug Discovery stage

Citation: Nataliia Koval, "Operational Structure of Early Drug Discovery Organizations", Universal Library of Multidisciplinary, 2026; 3(1): 83-92. DOI: https://doi.org/10.70315/uloap.ulmdi.2026.0301008.

must possess a clearly defined operational foundation that serves as a protective mechanism against excessive process variability. Structural fragility is the result of the absence of a rigid linkage between levels of control (Zhao et al., 2024).

Within this methodology, research execution is visualized as a five-layer management system, in which each layer performs its own role in stabilizing the scientific search. Five-layered Early discovery framework is shown in Table 1.

Table 1. Layered Early Discovery Framework

Management Level	Functional Focus	Effectiveness Criterion	Primary Layer Risk
Layer 1, Scientific Hypothesis	Therapeutic logic and program design	Soundness of the target profile	Open-ended research without milestones
Layer 2, Execution Engine	Chemistry, biology, ADME, CRO operations	DMTA cycle throughput	Local optimization and queue accumulation
Layer 3, Data Integrity	Converting outputs into a corporate asset	Traceability and reproducibility	Data fragmentation and loss of context
Layer 4, Decision Governance	Discipline in advancing or terminating projects	Speed of Kill/Go decision-making	Retaining projects due to emotional attachment
Layer 5, Capital Allocation	Financial resourcing based on validated data	ROI per learning iteration	Resource burn driven by narratives, not facts

The first layer, the scientific hypothesis, is the domain of the highest uncertainty. The task of the operational structure at this level is to constrain the research's ambitions through clear time boundaries and evidentiary criteria. Without this, the organization risks endlessly researching while in fact failing to move closer to value creation. The second layer, the execution engine, integrates internal and external synthesis and testing capacities. At this level, it is critically important to avoid local optimization, in which one function, for example, chemistry, is accelerated at the expense of others, creating bottlenecks at the entry to biological testing. The third layer ensures the conversion of results into a corporate-level evidentiary base by embedding standards of traceability and auditability long before the beginning of formal preclinical studies. The fourth and fifth layers transpose operational activity into the strategic plane, ensuring that funding follows only validated progress.

For example, a company launches an early-stage program targeting a novel anticancer target. At the first layer, the team immediately fixes the working hypothesis, a 12-week time horizon, and two criteria for transitioning to the next stage: confirmation of the relationship between the target and the disease, and the existence of initial active compounds meeting a specified selectivity threshold. At the second layer, chemists, biologists, and an external contract site operate in a unified two-week cycle so that newly synthesized molecules enter biological validation immediately, without queue accumulation. At the third layer, each result is entered into a common system that includes the protocol version, experimental conditions, and author-specified details, so that after one month, it is possible to determine precisely why a compound demonstrated an effect and whether that effect can be reproduced.

The linkage between layers is subsequently manifested in managerial decision-making. If, by week twelve, the hypothesis has been confirmed only partially, but the data remain unstable, then the fourth layer requires halting the program until the gaps in the evidentiary base are eliminated, since advancing without reliable data increases the risk of false selection. In such a case, the fifth layer reallocates money and capacity to another program in which each iteration is already generating a stable increase in knowledge. This approach enables a leader to see the entire chain. The first layer defines the search boundaries. The second sustains the cadence of execution. The third converts results into a verifiable basis for inference. The fourth protects the portfolio from protracted decisions. The fifth retains resources in directions where learning genuinely increases the program's value.

Within such an architecture, the Discovery process is regarded as a tightly coupled network of nodes for the transfer of responsibility. The structural robustness of the system is determined by the reliability of the interfaces between functional units: computational modeling, synthesis, analytical biology, and pharmacology. The transition to such a model requires abandoning department-level management in favor of managing the value stream.

Interface Architecture

Cross-functional interfaces are the most vulnerable points of any research organization (Xie et al., 2025). Failures at functional boundaries, such as the transfer of compounds without metadata or the interpretation of biological signals in isolation from experimental conditions, generate cascading variability that necessitates rework (Samuel & König-Ries, 2021). The present methodology requires that interfaces be treated as formal operational constructs with rigidly specified input and output parameters.

The central mechanism of interface control is the Handover Gate. It is an operational barrier that confirms that the output of one function meets the decision-readiness standards for the next function. A Handover Gate is considered viable only when seven control fields are present: transfer object, package format, acceptance criteria, distribution of

responsibility, time frame, deviation protocol, and version-traceability requirements. Figure 1 shows the architecture of the interface gateway.

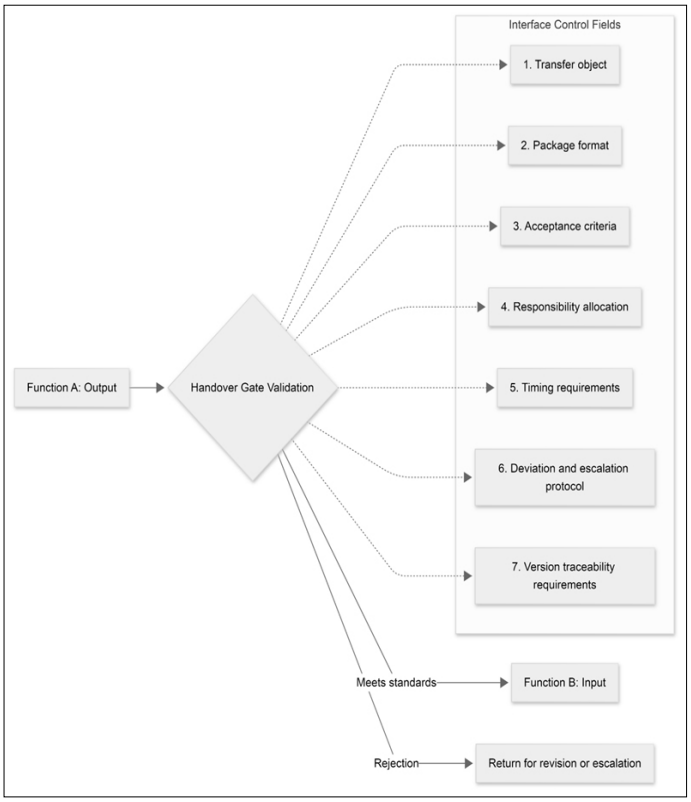

Fig. 1. Architecture of the interface gateway in the Discovery system

The application of a responsibility-allocation matrix at the boundaries between functions enables elimination of situations in which accountability for the result is diffused across teams. Consider the following example. A chemistry unit transfers ten new compounds to biology for a cell-based target inhibition assay. The transfer is permitted only after completion of a batch card indicating the code of each compound, confirmed purity of not less than 95 percent, solubility in the selected medium, stock-solution concentration, solvent used, preparation date, permissible storage period, the name of the responsible employee, and the version number of the file containing analytical results.

Before launching the assay, the biological group verifies these data against the acceptance criteria. If solubility data are unavailable or purity falls below the threshold, the samples are returned immediately for correction, and the experiment is not started. Such an order should be instituted for every boundary between functions. First, the organization approves the mandatory composition of the transfer package. Then it fixes the acceptance thresholds and the response time of the receiving party, for example, within one working day. After that, it appoints the employee authorized either to accept the object into work or to reject it with an indication of the reason. As a result, only verified compounds enter cell-based assays, the number of false biological signals declines, and repeated cycles of synthesis and testing arise substantially less often.

Informational Traceability

In contemporary R&D organizations, digital platforms such as electronic laboratory notebooks (ELN) and laboratory information management systems (LIMS) cease to be secondary documentation tools and become primary drivers of structural resilience. They constitute the organization's nervous system, coordinating object identities, definitions, and decisions (Higgins et al., 2022).

In this hybrid manner, the ELN and LIMS together fulfill the principle of active coordination: the ELN holds the intellectual memory of the scientific narrative, hypothesis, strategy, and logic. LIMS systems track samples, plates and reagents through the DMTA cycle and act as a physical anchor. These integrated systems allow the transition of compounds between laboratories, whether cross-border or moving from internal laboratories to partner contract laboratories, while retaining the context and version history of each compound. This is especially important for programs using machine learning, since the quality of predictive models depends directly on the absence of identity drift in the training dataset.

The economic significance of traceability is revealed in the prevention of micro-delays. Individually, a two-day delay in verifying a compound version may appear insignificant, but across eight iterative cycles with six responsibility transitions, cumulative losses reach 96 working days. At the average operational expenditure level of a biotechnology startup, USD 1.5 million per month, such structural losses are equivalent to USD 4.5–6 million of capital effectively burned while awaiting data. Thus, the engineering design of the digital landscape constitutes a primary element of the company's financial strategy. Functional roles of ELN and LIMS are illustrated in Table 2.

Table 2. Functional Roles of ELN and LIMS in Experimental Data Management

Data parameter	Role of the ELN system	Role of the LIMS system	Structural effect
Object identity	Scientific justification of the structure	Registration of digital birth and batch ID	Elimination of duplication and sample confusion
Experimental context	Description of conditions, controls, and replicates	Linking data to specific equipment	Ensuring reproducibility of results
Traceability	History of hypothesis modification	Movement of material between network nodes	Readiness for audit and IP protection
Decision readiness	Annotation of conclusions	Automatic verification of gate criteria	Reduced decision latency

This logic should be implemented through a unified digital registration rule for every object from the moment it appears in the workflow. For example, immediately upon receipt of a new compound, a unique batch identifier is assigned to it in the laboratory information system; then, in the electronic notebook, the origin of the structure, the experimental design, the setup conditions, the composition of controls, and the rationale for the next step are recorded. At every transfer of material, results, or files between units, the employee must confirm the alignment of the identifier, record version, experimental parameters, and sample status. After which, the system automatically blocks further movement in the event of any discrepancy. Such an order prevents batch mixing, loss of context, and the need to reverify already established data in advance, reduces hidden pauses between cycles, and preserves the integrity of the dataset on which subsequent decisions and computational models are built.

CHAPTER 2. PROCESS-FLOW ENGINEERING AND QMS AS A STABILIZATION INSTRUMENT

Discovery as an Iterative Learning System

The fundamental unit of progress in drug discovery is the completed Design, Make, Test, Analyze, Decide cycle, abbreviated DMTA. An organization's effectiveness is determined by the frequency and reliability of these cycles. It is important to understand that an incomplete cycle, for example, when compounds are synthesized but not tested due to the absence of biological capacity, merely consumes capital while creating the illusion of activity.

Scientific convergence, that is, movement toward the optimal molecule, requires dense feedback loops. If the results of pharmacokinetic assays are returned to design chemists with unpredictable delays, the precision of the next design round declines, increasing the total number of iterations required.

In practice, this logic should be implemented by strictly managing the duration of each complete research loop. For example, for one series of molecules, the team predefines the maximum time for the entire passage from design decision to final conclusion, say three weeks, with five days allocated to synthesis, four days to initial biological testing, two days to results review, and one day to formulation of the next decision. If, at the start of the cycle, the biological unit does not confirm the availability of a free testing window, the launch of the series is postponed immediately, as releasing compounds without subsequent testing merely accumulates costs and severs the learning chain.

To increase the speed of convergence upon a workable molecule, a rule should be introduced mandating the return of results to the project group within a fixed time period and in a standard format, including activity, solubility, stability, deviations, and a brief conclusion regarding the causes of success or failure. Then, during the weekly review, only one next decision is taken for each series: continue the direction, adjust the structure, or terminate the work. Such an order reduces signal lag between links, decreases the number of empty repetitions, and makes each new loop a source of useful knowledge suitable for the next step.

Thus, an Early Drug Discovery organization should be designed as a feedback control system in which the time required for a signal to traverse all nodes is a rigidly controlled parameter.

Management of Structural Constraints

According to the theory of constraints, the overall speed of project passage through any R&D organization is limited by a single node, the constraint (Tang et al., 2024). Attempts to accelerate other departments without relieving the bottleneck lead to uncontrolled queue growth and increased Lead Time.

In Discovery, the constraint is dynamic: at the hit-identification stage, it is usually localized in synthetic chemistry. During SAR expansion, the constraint often migrates toward biological assay throughput. At the lead-optimization stage, the critical resource becomes DMPK bandwidth or data-integration capacity.

A key factor in flow stability is managing capacity utilization. In an experimental environment with high internal variability, for example, synthetic failure, the need for assay repeats, and the aspiration toward 100 percent personnel utilization is destructive. At utilization above 90%, even minor disruptions cause exponential growth in queue waiting time. The present methodology recommends maintaining the load on the critical constraint at 75–80%, with the remaining 20% serving as a buffer to absorb variability and support process improvement. System load regimes and their impact on lead time are shown in Table 3.

Table 3. System Load Regimes and Their Impact on Lead Time

Load level	Queue dynamics	Impact on Lead Time	Recommended action
< 70%	No queues	Minimal cycle time	Risk of underutilizing capital
75–80%	Manageable micro-queues	Stable, predictable pace	Optimal operating point of the system
85–90%	Growth of work in progress	Noticeable cycle-time stretch	Incoming tasks need to be filtered
> 95%	Exponential queue growth	Unmanageable delays	Flow collapse and loss of control

In practice, this logic should be applied by continuously identifying the current bottleneck and directly subordinating all incoming work to it. For example, if during the compound-series expansion stage the biological group can reliably test only 30 samples per week, then this limit serves as the basis for the weekly plan for chemists, analysts, and coordinators.

In such a situation, the release of new compounds should be restricted in advance to the volume that biology can actually process without queue accumulation, while all requests above this threshold should be shifted to the next cycle. At the same time, it is useful to introduce a reserve of capacity equal to approximately one-fifth of the available resource, so that the system can absorb repeats, experimental failures, and urgent validations without an immediate increase in delays.

To avoid the hidden disintegration of flow, the program leader should review once per week exactly where unfinished work is accumulating, how many objects are waiting in the queue, and the actual passage time through the constraining node. If the queue begins to grow, the first action should be to reduce the system input, rebuild priorities, and temporarily refuse secondary tasks. If the constraint has shifted, for example, from chemistry to distribution and excretion studies, the plan is immediately rebuilt around the new limit. Such an order makes it possible to preserve a predictable tempo, maintain cycle controllability, and protect the program from a situation in which formal busyness grows faster than real scientific progress.

Quality Management System in R&D

In the context of Early Drug Discovery, the quality management system serves as a stabilizing layer of the architecture and helps build trust in the evidence being generated. Unlike a manufacturing quality management system, which is oriented toward strict compliance with specifications, R&D-QMS focuses on reproducibility, data comparability over time, and traceability of changes.

The methodological foundation of QMS at the discovery stage includes three key elements. The first element, Backbone Alignment, ensures the interpretability of changes in the methods of one function across other functions. This means that any changes must be recorded in a way that allows them to be properly accounted for in adjacent processes. For example, the replacement of a reagent in a biological assay must be reflected in the protocol version used in subsequent ADME modeling.

The second element, Constraint Protection, is associated with applying more stringent quality and control standards at the stage that determines the limiting speed of the entire DMTA cycle. Such an approach is necessary to prevent defective input data from entering the process bottleneck and to reduce the system's overall efficiency.

The third element is disciplined feedback. At its core lies the application of CAPA procedures, corrective and preventive actions, for analyzing operational failures and preventing their recurrence. Such a mechanism enables not only identifying the causes of deviations but also reducing the likelihood of their systematic reproduction in the future. Figure 2 illustrates the QMS methodological foundation at the discovery stage.

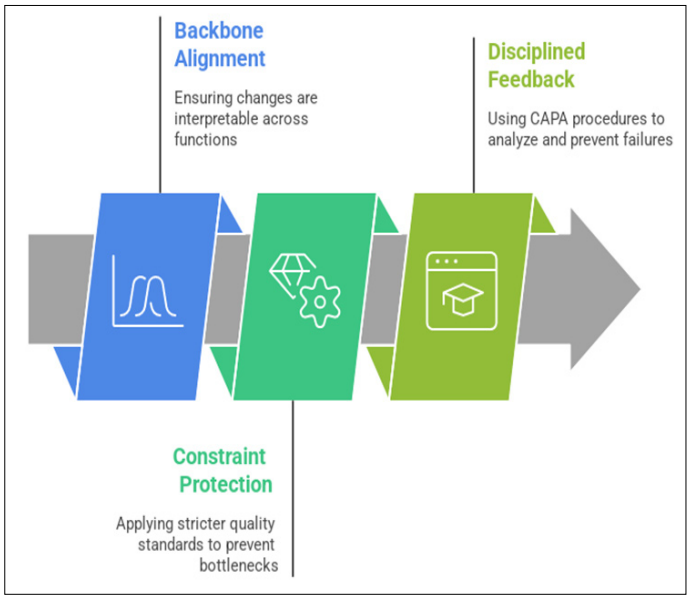

Fig. 2. QMS Methodological Foundation at Discovery Stage

SOPs, that is, standard operating procedures, are regarded in the present methodology as the fixed, best available way of performing a task at the current moment, ensuring the organization's institutional memory amid personnel turnover.

In practice, this system should be implemented by requiring the mandatory recording of all changes to methods, materials, and interpretive rules before they affect the next research step. For example, if a biological assay changes a reagent, sensitivity threshold, or signal-calculation method, this change is immediately incorporated into the approved version of the procedure, linked to the date of entry into force, and communicated to the groups that use these data in pharmacokinetic analysis, modeling, and the selection of new structures. For this purpose, it is useful to establish a rule that no result is included in collective decision-making unless it includes a reference to the current procedure version, a confirmed executor, and a record of any deviations during the experiment.

For example, a biological testing group altered the buffer composition in a cell-based assay to reduce background signal. From that moment, the assay leader is obliged immediately to update the current procedure, assign the new version an effective date, and indicate from which experimental series the new composition began to be applied. This information is then transmitted to the groups that use the assay results for pharmacokinetic analysis, cross-series comparison, and the selection of subsequent compounds. If this is not done, the new activity values will be compared with the old ones as though the conditions were equivalent, even though the measurement basis has already changed, and the team may make an erroneous decision to advance a weak candidate.

A simple rule should be introduced. Any change in reagent, protocol, signal threshold, or method of result calculation automatically requires three actions on the same day: update the procedure, link the change to specific data series, and

notify all functions that rely on these results in subsequent steps. If modified data are directed to the most heavily loaded segment of the program, for example, distribution and excretion studies, an additional check of metadata completeness and method version is performed before transfer. Such an order sharply reduces the risk of hidden mixing of incompatible data and protects the program from delays caused by late error detection.

Prevention of Distortions

Distortion of information in the research process rarely arises from large and obvious errors. Much more often, it develops gradually through the accumulation of small inconsistencies, each of which may appear rational at the local level when considered in isolation. Optimizing a data-processing script in one laboratory or changing the plate supplier in another may indeed improve local KPIs. However, in the absence of centralized change management, such decisions disrupt the longitudinal comparability of SAR data.

To prevent such distortions, integrity gates are used at the junctions of DMTA stages. Their function is to prevent ambiguous, incomplete, or poorly controlled data from passing to the next stage of the process. One such mechanism consists of preliminary design verification. It enables molecules to be filtered according to criteria of synthetic accessibility before they enter the synthesis queue.

Another mandatory element is QC validation. This involves verification of the purity and identity of each batch before the start of biological testing. Such a requirement reduces the risk that samples with unconfirmed characteristics will be transmitted into downstream work. Another important mechanism is structured data annotation. In this case, metadata must be completed before information is uploaded to the shared analytical environment.

Investments in the prevention of distortions at early stages yield practical benefits through reduced Decision Latency. By this is meant the time management spends determining what the obtained data actually signifies. Reducing this delay enables greater effort to be devoted to strategic decision-making.

For example, a project group has planned the synthesis of forty new molecules for series expansion; however, before entering the queue, each structure undergoes preliminary feasibility assessment. If the calculation indicates an excessively large number of stages, an unstable intermediate, or dependence on a difficult-to-obtain starting material, such a molecule is immediately excluded from the nearest cycle and replaced by a more attainable variant. After synthesis, each batch obligatorily undergoes confirmation of purity and identity. If the purity value falls below the established threshold, the sample is returned for rework and is not admitted to biological validation. This makes it possible to eliminate a future source of distortion before disputed activity results even appear.

Before results are uploaded to the common analytical environment, a mandatory set of attributes is completed for each measurement point: compound code, batch number, experimental conditions, cell-line type, testing date, plate used, operator, and version of the calculation template. If even one field is missing, the record is temporarily blocked and does not participate in series review. Such an order is especially useful when one laboratory has changed its consumable supplier, and another has updated its signal-processing method. With rigorous verification, the team will immediately identify the source of the discrepancy and will not spend weeks attempting to interpret incomparable data as a single scientific signal.

CHAPTER 3. SCALING, DISTRIBUTED TEAMS, AND THE STRATEGIC MANAGEMENT OF CROS

Geographically Distributed R&D Organizations

International expansion and the use of global talent transform the organization into a distributed network of functional nodes. A typical architecture includes a strategic management and modeling Hub, for example, in the United States, a high-throughput synthesis center in Eastern Europe, and specialized biological sites in Asia or Europe.

A distributed structure confers advantages in cost and access to expertise, but it introduces new systemic risks: desynchronization across time zones, differences in how problems are escalated, and regulatory heterogeneity. The strategic response to these challenges is to unify operational logic while preserving the geographical distribution of execution. This is achieved through the implementation of a unified digital backbone and synchronized communication windows. In a distributed system, readiness must mean the same thing everywhere, which requires rigid standardization of interface gateways.

For example, a project group locates scientific leadership and decision-making in one center, compound synthesis in Eastern Europe, and biological testing in Asia. For such a scheme to function stably, a unified definition of object readiness for transfer must be introduced in advance across all sites. A compound may move from synthesis into biological work only if a complete digital package is available. This package includes batch code, structural confirmation, purity, solubility, storage conditions, release date, protocol version, and responsible executor. If at least one parameter is missing, the transfer is automatically suspended until the gap is eliminated. This is especially important when time zones differ, as an error noticed too late can halt the next stage for an entire day.

For practical implementation, two common communication windows should be established during the week, when all

key functions are simultaneously available to align on risks, deviations, and urgent decisions. In addition, each transfer between sites requires a unified escalation procedure: who records the problem; within how many hours it must be acknowledged; and who makes the decision on whether to return, repeat, or continue the work. For example, if a biological site receives a batch without confirmed solubility, it transfers the case that same day into the established review channel, while the chemistry group is obliged to provide either a corrected package or a decision to replace the sample within a predefined period. Such an order reduces time losses due to clarifications, equalizes work discipline across countries, and renders the distributed organization manageable as a single system.

Strategic Vendor Management

In the contemporary Discovery model, external partners, including CROs, are regarded as integral nodes of the company's operational system. Their participation extends far beyond an auxiliary function and becomes part of the general logic of the research process. For this reason, the traditional approach to vendor management, built on transactional contracts, proves insufficient for fully integrating external execution into the iterative learning cycle.

The proposed methodology describes a four-layer CRO management model in which each layer is responsible for a separate aspect of the resilience and coherence of interaction. At the strategic level, attention is focused on managing the dependency portfolio and diversifying territorial risks. Such an approach enables the system to reduce its vulnerability to external disruptions and to create a more resilient configuration of the partner network.

The operational layer is associated with synchronization of CRO activity with the organization's overall DMTA rhythm. For this purpose, queue-monitoring instruments are used to embed external execution schedules into the unified logic of task movement. As a result, partners' activity becomes part of the common operational dynamics rather than an isolated circuit with its own temporal structure.

The technical layer is focused on harmonizing analytical methods and ensuring direct access to primary data, that is, raw data. This is necessary to ensure the results obtained by external executors are interpretable and fully usable within the company. Under such an approach, the transparency of the analytical process increases, and the risk of losing significant details during information transfer declines.

The data layer ensures the use of unified schemes for result transmission, enabling automatic integration into the corporate platform. As a result, data enter the common system in a consistent form and can be incorporated into downstream analysis more rapidly. Such an organization of exchange creates the conditions for a more connected, governable, and coherent research circuit. The four-layer CRO management model is illustrated in Figure 3.

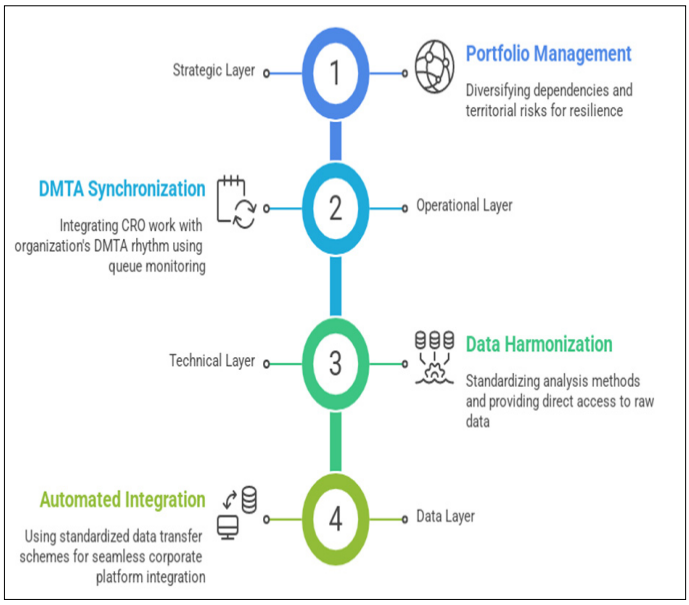

Fig. 3. Four-Layer CRO Management Model

Relationships with key CROs should be designed analogously to a software API: predictable inputs, standardized outputs, and transparent error handling. This minimizes the hidden integration costs that often nullify the economic benefit of outsourcing.

For example, a company transfers a series of twenty compounds to an external contractor for primary biological testing. For such a scheme genuinely to accelerate the program, a unified mode of interaction should be fixed before work begins: exactly which samples are transferred; in what form the results are delivered; which indicators are mandatory for each measurement point; within what time period the contractor must report a delay, material defect, or deviation from protocol; and who on the company side decides on repeat, stoppage, or continuation of work. At the same time, it is useful to require the transfer of primary measurements together with the final tables, so that the internal team may recheck calculations and compare results with data from other directions.

For practical application, work with the contractor should immediately begin as with a full-fledged node of the common research system. If the external executor sees that the testing queue is growing and the result-delivery date is shifting by five days, it is obliged to report this before the overall cycle is disrupted, so that the company can, in time, reduce incoming flow or reallocate part of the series to another site. If results arrive in a format that cannot be automatically uploaded to the common data environment, the package is considered incomplete and returned for correction on the same day. Such an order reduces hidden coordination losses, preserves the tempo of research loops, and enables real acceleration from external execution rather than merely formal budget savings.

Management of Systemic Risks

Systemic risks in an environment with a high share of external

execution are associated with informational asymmetry and the concentration of dependencies. To manage these risks, the organization should use a vendor-management maturity matrix that evaluates partners across four domains: technical stability of methods, metadata accuracy, reliability of deadline adherence, and concentration of work volumes. Table 4 illustrates the risk domains and mitigation mechanisms in distributed operations.

Table 4. Risk Domains and Mitigation Mechanisms in Distributed Operations

Risk domain	High-risk indicator	Mitigation mechanism
Technical	Undocumented changes in protocols	Mandatory protocol version confirmation at each handoff
Informational	Absence of metadata and access to raw files	Implementation of automated data-ingestion gateways
Operational	Lack of queue transparency and unpredictable Lead Time	Quarterly review of partner throughput capacity
Structural	>60% of a critical function concentrated in a single jurisdiction	Strategic node redundancy, Dual Sourcing

Particular attention is devoted to the prevention of silent drift, the gradual deterioration of quality or speed that does not manifest as a one-time failure but undermines the long-term effectiveness of the portfolio. Regular audits of interfaces and monitoring of the share of preventable rework enable detection of system degradation before it leads to program stoppage.

For example, a company assigns key biological assays to a single external contractor that processes the bulk of the program's workflow. In order to detect the concealed accumulation of risk in time, four indicators should be tracked monthly for such a partner: whether there have been protocol changes without a confirmed version; whether results arrive with the full set of metadata and source files; how many days execution actually requires compared with the agreed timeline; and what share of the entire critical function is covered by one site or one country. If the contractor begins to delay results more frequently, send incomplete records, or change analytical conditions without explicit documentation, this must be treated as an early signal of systemic deterioration and immediately transferred into the mode of managerial review.

For practical application, threshold actions should be introduced for each such signal. In the event of repeated metadata absence, an obligatory inspection of the data-transfer channel should be conducted. If delays increase, the volume of incoming tasks and the partner's confirmed throughput capacity should be reconsidered. In cases of excessive critical work, part of the flow should be shifted in advance to a reserve executor to preserve program resilience in the event of an external disruption. For example, if one contractor in one jurisdiction performs more than half of all distribution and excretion assays, the company should prepare in advance a second executor with a comparable method and a compatible data-transfer scheme. Such an order enables detection of slow degradation before deadlines collapse and reduces the likelihood that the entire research program will become dependent on a single vulnerable point.

Challenges of Scaling

As an organization grows, the number of interactions between teams increases nonlinearly. Scaling without structural adaptation inevitably leads to an increase in interfaces, a heavier managerial burden, and slower decision-making. The principal challenge lies in maintaining a high density of learning as the volume of activity increases.

Successful scaling requires a transition from informal communication to standardized protocols, specifically Unified Interface Protocols. Such a transformation creates a more stable and predictable environment of interaction in which the exchange of actions, decisions, and results becomes formalized and reproducible. This is especially important under conditions of increasing complexity, when the coherence of processes begins to determine the system's overall governability.

The implementation of the Breathing Organization principle plays a substantial role. Within this approach, internal competencies are concentrated on design and analysis, while external capacities are used for synthesis and testing. Such a distribution of functions enables a more flexible operational configuration and maintains a balance between the depth of internal expertise and the scale of external execution.

An additional element is the formation of Centers of Excellence in key scientific directions. Such centers provide centralized expertise in distributed execution environments. Their significance lies in the fact that they create unified points of scientific coordination through which the quality of decisions, methodological coherence, and the accumulation of specialized knowledge are sustained.

Digital Process Twins constitute an important management instrument. They are used to model the impact of new projects on the load of already existing capacities. For example, a company is running eight programs and plans to open a ninth directed at a new target. Before launch, the portfolio leader collects data on the current load of the synthetic group, biological testing, and data analysis for the coming six weeks, and sees that in two weeks, the biological testing segment will enter overload due to already confirmed series. In such a situation, the ninth program is not launched immediately at full scale. Instead, its first wave is restricted to ten compounds, a testing window is reserved in advance, and a unified transfer package is fixed, including object code,

experimental objective, mandatory quality parameters, result-delivery deadline, and the person responsible for acceptance. At the same time, design and result reviews remain within the company, while additional synthesis is transferred to an external site in accordance with a pre-approved scheme.

Thereafter, all results of the ninth program are reviewed by a standing expert group that applies the same evaluation rules as for the other directions. If, after the first wave, it becomes evident that the actual testing time is rising faster than the permissible threshold, the company temporarily freezes series expansion, shifts part of the load to a reserve executor, and opens the next cycle only then. Such an order provides a concrete mechanism for scaling: new projects are added in a metered fashion, critical nodes are protected against overload, and growth in the number of programs does not turn into a chaotic accumulation of tasks without the capacity to rapidly extract knowledge from them.

CONCLUSION

The methodology provides an integral theoretical and applied foundation for designing organizations at the early stage of drug discovery, as governed by research systems in which the speed of evidence accumulation becomes the central operational variable. Its conceptual value lies in translating Early Drug Discovery from the plane of fragmented laboratory activity into the regime of an architecturally organized process, where hypothesis, execution, data, managerial decision, and capital allocation constitute a single causally connected configuration. From this perspective, scientific productivity is determined by the quality of interfaces between functional nodes, the density of feedback within DMTA cycles, and the degree of controllability of variability throughout the research trajectory. Thus, the methodology establishes a language in which organizational form is regarded as a direct factor in epistemic efficiency, and operational discipline is treated as an instrument for accelerating the reduction of uncertainty.

The methodology derives particular significance from its attention to the engineering of boundaries through which objects, data, and responsibility pass. The formalization of the Handover Gate, the integration of ELN and LIMS, and the introduction of criteria for traceability, versioning, and decision readiness create the conditions for transforming the experimental result into a stable corporate asset. Within such a model, data cease to be a secondary trace of scientific activity and become the infrastructure upon which reproducibility, comparability, and the managerial validity of conclusions rest. A substantial strength of the approach also lies in its conjunction with the theory of constraints and the quality management system, by virtue of which the research process acquires the contours of a governed flow sensitive to overloads, queues, hidden delays, and methodological drift. As a result, the organization becomes capable of sustaining a rhythm of learning in which each iteration produces a verifiable increment of knowledge suitable for subsequent project selection.

The methodology is highly applicable in distributed teams, portfolio scaling, and active interaction with external contractors, as it proposes unified operational standardization principles amid the geographical and organizational heterogeneity of execution. Especially productive appears the described approach to CROs as full-fledged nodes of the research system, with predictable inputs, standardized outputs, and transparent handling of deviations. Such a framing enhances the resilience of the entire R&D architecture, reduces the likelihood of silent defect accumulation, and makes external execution part of the common cycle of scientific learning. Taken together, the methodology may be regarded as a mature academic framework for the creation of highly effective discovery organizations, in which managerial logic, digital infrastructure, and research practice are joined into a single coordinated mechanism for converting capital into evidentiary value.

The limitations of this methodology are determined above all by its subject focus and normative density. It is oriented toward organizational design, flow management, data quality, and structural resilience; therefore, its maximum effectiveness is realized in an environment where readiness already exists for formalizing procedures, registering objects digitally, and adhering to unified transfer rules. In contexts marked by high heterogeneity in local practices, differences in the technical maturity of sites, or pronounced dependence on external executors, implementing the described principles requires sufficient institutional capacity and role alignment. In addition, the methodology primarily describes the architecture of execution and mechanisms for accelerating scientific convergence; consequently, it provides a strong framework for organizing the process, yet does not claim to offer an exhaustive description of all substantive features of individual therapeutic areas, experimental platforms, and specialized scientific strategies.

REFERENCES

1. Fernald, K. D. S., Förster, P. C., & Claassen, E. (2024). The pharmaceutical productivity gap – Incremental decline in R&D efficiency despite transient improvements. *Drug Discovery Today*, *29*(11), 104160. https://doi.org/10.1016/j.drudis.2024.104160
2. Galson, S., Austin, C. P., Khandekar, E., Hudson, L. D., DiMasi, J. A., Califf, R., & Wagner, J. A. (2020). The failure to fail smartly. *Nature Reviews Drug Discovery*, *20*(4), 259–260. https://doi.org/10.1038/d41573-020-00167-0
3. Higgins, S. G., NogiwaValdez, Akemi A, & Stevens, M. M. (2022). Considerations for implementing electronic laboratory notebooks in an academic research environment. *Nature Protocols*, *17*(2), 179–189. https://doi.org/10.1038/s41596021006458

4. Huanbutta, K., Burapapadh, K., Kraisit, P., Sriamornsak, P., Ganokratana, T., Suwanpitak, K., & Sangnim, T. (2024). The Artificial Intelligence-Driven Pharmaceutical Industry: A Paradigm Shift in Drug Discovery, Formulation Development, Manufacturing, Quality Control, and Post-Market Surveillance. *European Journal of Pharmaceutical Sciences*, *203*, 106938–106938. https://doi.org/10.1016/j.ejps.2024.106938

5. Kim, Y., Atukeren, E., & Lee, Y. (2022). A New Digital Value Chain Model with PLC in Biopharmaceutical Industry: The Implication for Open Innovation. *Journal of Open Innovation: Technology, Market, and Complexity*, *8*(2), 63. https://doi.org/10.3390/joitmc8020063

6. Sampathkumar, K., & Kerwin, B. A. (2023). Roadmap for Drug Product Development and Manufacturing of Biologics. *Journal of Pharmaceutical Sciences*, *113*(2), 314–331. https://doi.org/10.1016/j.xphs.2023.11.004

7. Samuel, S., & König-Ries, B. (2021). Understanding experiments and research practices for reproducibility: an exploratory study. *PeerJ*, *9*, e11140. https://doi.org/10.7717/peerj.11140

8. Tang, J., Dai, Z., Jiang, W., Wu, X., Zhuravkov, M. A., Xue, Z., & Wang, J. (2024). A Comprehensive Review of Theories, Methods, and Techniques for Bottleneck Identification and Management in Manufacturing Systems. *Applied Sciences*, *14*(17), 7712. https://doi.org/10.3390/app14177712

9. Tariq, J., & Gardezi, S. S. S. (2022). Study the delays and conflicts for construction projects and their mutual relationship: A review. *Ain Shams Engineering Journal*, *14*(1), 101815. https://doi.org/10.1016/j.asej.2022.101815

10. Xie, X., Ye, C., Hu, Q., He, W., & Dai, H. (2025). CrossFunctional Coordination, Status–Authority Asymmetry, and Contingent Exploitation of Digital Technology. *Management and Organization Review*, *21*(5), 864888. Cambridge Core. https://doi.org/10.1017/mor.2025.10100

11. Zhao, Y., Cai, B., Cozzani, V., & Liu, Y. (2024). Failure dependence and cascading failures: A literature review and research opportunities. *Reliability Engineering & System Safety*, *256*, 110766. https://doi.org/10.1016/j.ress.2024.110766