FROM CAPSID ACCESS TO IND: WHAT ACTUALLY HAPPENS NEXT

A Practical CMC Guide for Early-Stage Gene Therapy Teams Working with Engineered Capsids

Michael J. DiBiasio-White, PhD ,Founder & Principal, ViroSpark BioConsulting

You have capsid access. Now what?

Getting access to an engineered AAV capsid is a significant milestone. But capsid access is not a development plan. It is not yet a product, not yet a process, not yet a control strategy, and not yet an IND path.

What comes next is a linked system of decisions: capsid, payload, dose, potency assay, plasmid strategy, CDMO fit, tox material, GMP material, analytics, stability, comparability, and Module 3. All that become connected very quickly. One wrong assumption early does not just create a technical problem, it burns runway. For a pre-commercial team operating on Seed or Series A capital with investor milestones tied to manufacturing progress, a single avoidable mistake, a misdirected process development campaign, a CDMO that does not fit, a plasmid delay, a documentation gap that stalls the IND, can cost three to six months and a significant portion of the budget. That is not a setback. For some programs, that is the difference between reaching the next value-inflecting milestone and running out of time.

Most early-stage gene therapy founders know their biology. What they often have not done before is take a program from promising science through manufacturing and into an IND. The path ahead is not four clean steps. It is fifteen to twenty decision points, and the order and quality of those decisions determines how fast the program moves, how efficiently it uses capital, and whether the work done at each stage supports what comes next, or has to be repeated.

Some developers will work with an integrated manufacturing partner that has already onboarded the capsid technology, where the primary new variable is the payload. Others will use their own CDMO, where more of the process risk falls directly on the developer. In both cases, the developer still owns the product strategy, the regulatory narrative, and the choices that connect one stage to the next.

This paper walks through what that path looks like, through the lens of programs where getting these decisions right, or catching them early, made a measurable difference in speed, cost, and outcome.

AVOID SCALING THE WRONG THING

The instinct after capsid access is to make material. The better instinct is to make decision-grade material: Material made and characterized in a way that answers the questions standing between you and the next development move.

A capsid can look promising in discovery and still create problems in development. It may package your payload inefficiently. It may produce more empty capsids than expected. It may behave differently at process scale. It may require a dose that is not compatible with realistic manufacturing economics.

Before committing to full process development, which at a CDMO typically represents months of timeline and a significant financial commitment, the team needs a manufacturability screen. Not an optimized process. A focused assessment of whether this capsid-payload combination is developable- essentially a manufacturability screen.

We ran this screen across a multi-candidate program with four novel capsids on the same backbone. The panel assessed capsid titer, vector-genome titer, stepwise recovery, aggregation propensity, and impurity clearance across all four candidates under comparable conditions. One candidate produced the highest vg titers, the metric most teams would use to select a lead. But it also showed poor downstream recovery and disproportionate aggregation that would have made scale-up extremely difficult.

Without the screen, that candidate advances into full process development. The team commits three to four months of CDMO time and a substantial portion of their development budget to a program with a fundamental manufacturability limitation that does not surface until late in PD, or worse, during tech transfer. At that point, the team is not just behind on timeline. They have spent capital they cannot recover on work that does not carry forward, and they are starting over with a different candidate while their runway shortens.

With the screen, which cost a fraction of a full PD campaign in both time and money ,the team down-selected in weeks. They entered PD with a candidate that had a realistic manufacturing path and gave their manufacturing partner a starting dataset that compressed early development rather than creating it.

Your first goal is not to optimize the process. Your first goal is to avoid scaling the wrong thing.

BUILD THE MATERIAL STRATEGY AS ONE CONNECTED ARC

The founder's next question is usually: what material do we need, and when?

The typical breakdown is RUO material for early research and animal work, development material for process and analytical optimization, tox material for IND-enabling studies, and GMP material for clinical dosing. Most teams treat these as separate campaigns, separate purchase orders, separate timelines, sometimes separate vendors.

This is where avoidable rework begins.

When each stage is planned independently, the lessons from one campaign do not carry into the next. Process knowledge generated during RUO work gets left behind when development starts at a new facility. Development decisions are made without considering whether they will translate into GMP. Tox material is generated under conditions that cannot be bridged to the clinical product, creating a comparability burden that may require additional manufacturing runs, additional testing, or in the worst case, a repeat of nonclinical studies.

Every repeated campaign costs money the program has already spent once. Every gap between stages costs time to close retrospectively. The most capital-efficient path is one where each stage is designed to feed the next, where RUO informs development, development decisions carry into GMP, and the documentation captured at every stage supports the IND without reconstruction.

The early decisions that matter: Can RUO and development material be made on the same platform as tox and GMP? How much process change can the program tolerate between tox and clinical? What analytical methods need to exist before tox starts? What documentation must be captured now so it does not need to be rebuilt later?

A CDMO can execute a scope of work. They cannot define the sponsor's product strategy. The developer needs to own the logic that connects one stage to the next, because no one else will, and the cost of discovering that gap at the IND stage is orders of magnitude higher than addressing it upfront.

USE PLATFORM ASSUMPTIONS TO MOVE FASTER, BUT CONFIRM THEM

Most early teams want to leverage a platform process. That is the right instinct, platform assumptions are one of the most powerful tools for compressing timelines and reducing cost. But the platform gets you started. It does not get you finished.

For an engineered capsid program, much of the manufacturing logic may be borrowable: transient transfection, upstream conditions, harvest strategy, clarification, affinity capture, polishing, formulation, and analytical starting points. If the developer is working with a manufacturing partner that has already onboarded the capsid, a significant portion of this groundwork may already exist, which considerable cuts down on development time and cost.

But the payload is always a new variable. Payload size, packaging efficiency, transgene expression level, and genome structure can shift yield, impurity profile, empty/full distribution, downstream recovery, aggregation behavior, potency, and dose. A process that works well for one payload on a given capsid may underperform for a different payload on the same capsid. And the only way to know is to test it.

In one program we worked on, a novel capsid showed measurable instability during a standard 4°C hold step, a condition that is routine for natural serotypes. The fix was not complicated: a room-temperature hold with a modified affinity-capture load and supporting stability data. But discovering the problem without understanding what to look for would have meant lost material, failed runs, and weeks of troubleshooting. In the same program, a post-affinity pool showed elevated aggregation at capsid-particle concentrations that would have been completely acceptable for AAV9. Expanding the pool volume resolved it, but this is not a standard troubleshooting step in most platform processes.

The real value of understanding what is platform-transferable versus what is product-specific is speed. After building the first process in a multi-candidate program, we mapped each unit operation as platform or candidate specific. For subsequent candidates on the same backbone, this architecture compressed PD timelines from approximately four months to two. Upstream conditions, core purification parameters, formulation, and analytical methods were already locked. Only payload-driven variables required new work.

For a company needing to move fast and limited capital, that compression changes the economics of the entire portfolio. It is the difference between getting to that next milestone faster and securing the necessary capital to move forward or watching the timeline creep forward as you step backwards to move forwards again.

START ANALYTICS BEFORE YOU THINK YOU NEED THEM

Founders almost always prioritize process development and defer analytics. The logic is understandable, make material first, test it later. But this sequence creates one of the most common and most expensive delays in early-stage gene therapy programs.

Analytics are what tell you whether process development is working. Without the right measurements at the right time, the team is running PD blind, making decisions based on incomplete data, advancing material that may not meet the quality profile, and building a process they cannot fully characterize when the IND requires it.

Potency is the single biggest trap. FDA's draft guidance on potency assurance for cell and gene therapy products frames potency broadly, encompassing process design, process control, in-process testing, and release testing. For

gene therapy, the potency assay is rarely simple. It may require a cell-based system, a functional readout, or a surrogate that must be scientifically justified. Development of a credible potency assay takes time, often months.

Teams that wait until GMP to address potency face a cascade that is both expensive and difficult to compress: GMP lot release is delayed because the assay is not ready. Stability cannot be initiated on time, which constrains clinical supply. The control strategy in the IND is weak because the most fundamental quality attribute, cannot be adequately measured. That cascade does not just cost time. It costs the capital already invested in the GMP campaign, which sits unreleased while the assay catches up.

Beyond potency, in-process analytics are one of the most underused tools for moving faster. Measuring capsid titer, vg titer, recovery, aggregation, and impurity clearance at each process step, not just on final product, generates data that makes every subsequent decision faster and cheaper. The manufacturability screen described earlier worked precisely because it applied this approach: measuring at every step, across multiple candidates, under comparable conditions. That data did not just select a lead, it gave the manufacturing partner a roadmap that reduced the scope and cost of early PD.

The early analytical question is not "what assays do we need for release?" It is: what assays do we need to make decisions now, and which of those must mature into the tools that support release, characterization, comparability, and stability later? Getting that right upfront prevents the program from paying twice, once for development-stage data that does not carry forward, and again for IND-ready data that has to be generated from scratch. An accelerated path is to understand what assays can be platform or fit for use and use that for early development insight as you begin to move closer to your GMP batch.

LOCK THE PLASMID STRATEGY BEFORE IT LOCKS YOU

Plasmids are treated like a procurement detail. They are not. For transient-transfection AAV, plasmids are foundational raw materials and a plasmid decision gone wrong will cascade through the entire program timeline.

If plasmid supply slips, the engineering run slips. If the engineering run slips, analytical development slips. If analytical development slips, tox release slips. If tox slips, IND-enabling studies slip. If tox studies slip, the IND date becomes fictional. That is not a hypothetical sequence. It is the most common single-point-of-failure timeline in early-stage gene therapy and it is entirely preventable with early planning.

Beyond supply, grade transitions create hidden cost. If early process development uses research-grade plasmids and the team later switches to GMP-grade, the transition is not just a procurement change. It is a comparability event. Differences in topology, supercoiling, impurity profiles, and lot consistency can affect transfection efficiency, viral productivity, and product quality. For a well-characterized natural serotype, the impact may be absorbable. For a novel capsid-payload combination where the process window is still being defined, it can force additional manufacturing runs and months of bridging work, all of it unbudgeted.

We evaluated five plasmid suppliers for one program across cost, lead time, quality grade, cell-bank strategy, analytical package, geography, and audit requirements. The cheapest option used a rolling cell bank with a minimal analytical certificate, functional for early research but a liability downstream. The recommended path was a master cell bank approach with a fuller analytical package aligned to IND expectations, plus a sourcing plan that supported a clean grade transition from development through GMP.

The upfront cost difference was modest. The cost of not doing it, based on what we have seen in other programs, would have been additional manufacturing runs, a comparability assessment no one budgeted for, and three or more months of delay at exactly the point in the program where delay is most damaging to financing and milestones.

Build the raw material plan early enough that it supports the critical path. Do not wait until the CDMO asks for it.

DEFINE YOUR CDMO STRATEGY BEFORE ASKING FOR PROPOSALS

The most expensive CDMO mistake is not choosing the wrong one. It is going to them before you know what to ask for.

A founder approaches a CDMO and says "can you make our AAV?" The CDMO says yes and sends a proposal. The proposal looks reasonable. The price fits the budget. The timeline works. The team signs.

Six months later, the team discovers that analytical development was not included in the scope. Development reports were not part of the deliverables. The process includes a CDMO-proprietary step that cannot be replicated at another facility without a license. The batch records generated will not support Module 3 without significant rework. Deviation investigations are out of scope. The cheapest proposal has become the most expensive decision in the program, not because of what it included, but because of what it did not.

This is process debt. It does not appear on the invoice. It appears six to twelve months later as rework, as delays, as additional campaigns, and as budget that was supposed to fund GMP but is now funding gap-filling work that should have been covered in the original engagement. For a company operating on limited capital, process debt is not an inconvenience. It is a direct threat to the program's ability to reach the next milestone.

Before any RFP, the sponsor needs to define: product and capsid-payload, target scale, number of lots across stages, analytical responsibilities, documentation deliverables, data ownership, change control expectations, and IND-supporting outputs. The goal is not to find the cheapest CDMO or the most prestigious one, it is to find the partner whose capabilities and operating model fit the actual stage and scope of the program, with a SOW that is explicit about what is included and what is not.

For a pre-seed academic spin-out with a novel capsid and grant-driven timelines, we built a structured evaluation and matched the team to a suitable partner in under two months. Programs without structured evaluation typically take three to four months of back-and-forth, and sometimes end up switching partners after the first campaign, adding six months and the full cost of a wasted manufacturing engagement to a budget that cannot absorb it.

Do not outsource your strategy. Outsource defined work packages.

BUILD THE TECH TRANSFER PACKAGE BEFORE THE CLOCK STARTS

Tech transfer is where more programs stall than almost anywhere else and where the cost of stalling is highest, because by this point the team has already invested months of development work and significant capital. A transfer failure does not just delay the program. It puts the value of everything that came before it at risk.

The core issue is knowledge asymmetry. The developer understands the capsid biology, the payload, and the intent behind each process step. The receiving facility understands GMP operations, equipment constraints, and regulatory expectations. The gap between those perspectives is where transfer failures live: a process step that worked in development but cannot be executed in the GMP suite, an analytical method that was never formally transferred, a material that was used at research grade with no GMP-qualified equivalent identified.

Whether the developer is transferring into a manufacturing partner that has already onboarded the capsid or into their own CDMO, the developer still owns the payload-specific process knowledge, the analytical context, and the development history. That information must transfer cleanly, because a receiving site forced to guess will make conservative choices that cost time, or wrong choices that cost material and schedule.

Most teams treat tech transfer as a sequential activity that happens after process development is complete. That sequence: finish PD, then write the transfer package, then transfer, adds months to the timeline that are avoidable.

For a Series A client using an engineered capsid heading toward IND, we built the tech-transfer architecture in parallel with PD. That meant conducting a facility-fit assessment before the GMP site was locked, documenting every process modification with rationale as it happened, aligning analytical packages between development and GMP sites during method development rather than after, and retaining in-process samples from development runs for future bridging and comparability. The transfer package was not assembled after PD. It was built as a living deliverable throughout PD.

The IND was accepted by FDA. That outcome was not because the product was simple. It was because the parallel approach eliminated the sequential gap that most programs accept as unavoidable, the months between "PD is done" and "we are ready to transfer." Those months cost money (the team is burning operating capital while no manufacturing is happening), they cost momentum (investors see timeline slippage), and they are almost entirely preventable with planning.

WRITE THE IND WHILE YOU ARE DOING THE WORK

Module 3 requires a coherent narrative: what the manufacturing process is, how it was developed, what changes were made and why, what the product looks like analytically, and how it is controlled. When that narrative is written retrospectively, which is what most teams do, gaps appear. Development rationale is reconstructed from memory. Batch records are incomplete. Process changes that seemed minor turn out to be significant from a regulatory perspective.

Retrospective documentation is also the most expensive and lowest-value way to spend CMC budget. The team is paying senior technical resources to reconstruct information that could have been captured in real time at a fraction of the effort. And the quality of retrospective documentation is always lower, because memory degrades, personnel change, and context is lost.

We have seen both sides. In the engineered capsid program described above, we maintained continuous documentation across every development, toxicology, and GMP campaign. When the IND was assembled, Module 3 was not a reconstruction, it was a consolidation of records that already existed, with narrative already built around each decision point. The assembly phase was weeks, not months.

In a separate engagement, a company preparing for PPQ and Phase 2/3 had accumulated significant documentation backlog across their development and characterization history. We systematically identified what was missing, built the reports, and created a documentation architecture to carry forward. That rescue took months of dedicated work; work that would have been a fraction of the effort and cost if the documentation had been maintained throughout development.

Write the development story while the development is happening. It is the single cheapest investment a team can make toward a clean IND submission.

WHERE ONE WRONG MOVE CREATES A CASCADE

The biggest risks in early-stage gene therapy CMC are not dramatic failures. They are quiet compounding effects from decisions that were not made deliberately, and each one costs time and money the program has already committed.

Treating RUO material as "good enough" without understanding how it was made creates a weak process history, an unclear tox bridge, a difficult GMP transfer, and a Module 3 narrative that does not hold together. The cost: months of retrospective work and potentially a repeat manufacturing campaign.

Choosing a CDMO before defining the scope creates an incomplete SOW, missing analytics, change orders, and budget burn. The cost: the difference between the original proposal and the actual spend, which can be multiples of the original, plus three to six months of timeline expansion.

Waiting too long on potency means no meaningful release assay at GMP. The cost: a completed GMP campaign that cannot be released, with capital locked in unreleased material while the assay catches up.

Plasmid strategy not locked early means a supplier delay cascades through every downstream activity. The cost: three or more months of delay at exactly the point where the team needs to demonstrate manufacturing progress to investors.

Assuming the platform process will work without confirmation means discovering yield, recovery, or stability problems after months of committed development. The cost: the full investment in misdirected PD, plus the time and capital to start over.

No sponsor-side CMC owner means vendors execute disconnected work packages with no integrated view. The process data, the analytical story, and the regulatory narrative do not connect. The cost: a late-stage rescue project that is always more expensive, slower, and more painful than the upfront planning it replaces.

Every one of these cascades is avoidable, not with a bigger team or more capital, but with experienced CMC judgment applied early enough to prevent the compounding from starting.

THE GAP VIROSPARK IS BUILT TO FILL

Early-stage gene therapy development is defined by constrained capital, tight timelines, and milestones that determine whether the program continues. The CMC decisions described in this paper are the ones that most often determine whether a team stays on track or spends its runway on avoidable rework.

ViroSpark BioConsulting provides practical, operator-level CMC and TechOps support for early-stage gene therapy teams working with engineered capsids. We work alongside developers and their manufacturing partners, not as a replacement for either, but as the experienced layer that connects capsid science to manufacturing reality, process development to regulatory strategy, and early decisions to long-term program success.

A typical engagement starts with a 2–4 week CMC triage: reviewing the capsid-payload concept, mapping existing data, identifying critical gaps, building the material and analytical strategy, and delivering a clear, milestone-driven CMC roadmap. From there, the work extends into CDMO evaluation and SOW structuring, sponsor-side execution oversight, tech-transfer planning, development reporting, and IND assembly, as much or as little as the program needs.

The point is not to add processes. It is to remove waste, wasted time, wasted capital, wasted effort on work that does not carry forward. The fastest path to IND is not the one with the fewest steps. It is the one where every step counts.

Built for pre-commercial budgets. No enterprise overhead. No six-figure minimums. Project-based or retained.

Michael J. DiBiasio-White, PhD Founder & Principal, ViroSpark BioConsulting michaelwhite@virosparkbio.com
linkedin.com/in/michael-dibiasio-white virosparkbio.com

References

1. Iwuchukwu, I. "Conquering the Complexities of AAV Manufacturing." GEN 45(5), 2025.

2. "AAV Gene Therapy and Capsid Engineering: Challenges, Strategies, and Directed Evolution." PackGene Knowledge Library, January 2026.

3. Adams, B. et al. "Scalable Production and Purification of Adeno-Associated Virus." Molecular Therapy: Methods & Clinical Development 32(2), 2024.

4. Mattison, T. "Platform-Based Gene Therapy Manufacturing Strategies for Scale-Up." AGC Biologics, October 2025.

5. FDA. Chemistry, Manufacturing, and Control (CMC) Information for Human Gene Therapy Investigational New Drug Applications (INDs): Guidance for Industry. January 2020.

6. FDA. Potency Assurance for Cellular and Gene Therapy Products: Draft Guidance for Industry. February 2024.

7. FDA. Manufacturing Changes and Comparability for Human Cellular and Gene Therapy Products: Draft Guidance for Industry. 2023.

8. Oyama, H. et al. "High-Efficiency Purification of Divergent AAV Serotypes Using AAVX Affinity Chromatography." Journal of Chromatography B 1234, 2024.