



WHITE PAPER

When Platform Processes Break: CMC Realities of Manufacturing Novel Engineered AAV Capsids

A Practical Perspective from the Developer Side of the Table

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

Introduction

AI-designed and rationally engineered AAV capsids are delivering tissue tropism, immune evasion, and potency profiles that natural serotypes cannot match. For early-stage gene therapy developers, access to these capsids represents a genuine competitive advantage.

But novel engineered capsids don't behave like AAV9 or AAV5 on a manufacturing floor. They challenge platform assumptions across process development, purification, analytical characterization, and GMP readiness. For resource-constrained teams, these challenges compound into costly rework, missed milestones, added change orders and delayed FIH timelines. Often because the problems aren't visible until the work was started.

This paper offers an operator-level perspective on the technical CMC realities of manufacturing novel engineered capsids, drawn from direct experience including leading a novel capsid program from platform concept through FDA-accepted IND with no CMC comments, and current work supporting early-stage and cross-border AAV programs through PD, manufacturing, and IND strategy.

Platform Assumptions Don't Always Apply

Most CDMOs have built their AAV capability around natural serotypes (AAV9, AAV8, AAV5, AAV2) refined over time into platform approaches with standardized production systems, purification trains, hold strategies, and analytical panels. One of the biggest challenges in AAV manufacturing today is around novel capsids and speed of development.

For most developer's the use of an engineered capsid, is driven by better specificity and the promise of a lower dose. The assumption around AAV manufacturing being more platform is relatively true for commonly used serotypes and so a lot of CDMOs have built platforms around these. But engineered capsids routinely deviate in ways that derail timelines if not caught early: production behavior that doesn't match AAV9 experience, batch-to-batch variability requiring adaptable strategies and surface interactions creating yield losses; often observed during small scale RUO generation.

Where I've seen this break:

On one novel capsid program, I led discovery through IND, the capsid did not tolerate extended 4°C holds due to interactions with buffers and process-related contaminants post-TFF. The CDMO's validated hold conditions were built around natural serotypes that tolerate those holds without issue. Their GMP suite also lacked -80°C capability, so freezing through the hold wasn't an option. We engineered an alternative: a room-temperature hold for a shorter duration with a modified affinity load, supported by hold-time stability data, demonstrating equivalent process clearance. That was a sponsor-side process decision; one that required understanding both the capsid's behavior and the CDMO's operational constraints simultaneously.

On the same program, the capsid was significantly more prone to aggregation than platform serotypes at equivalent vg/mL. The critical insight: vg/mL doesn't tell the full story. Total capsid particle concentration (the inclusion of empty and partials) can be substantially higher than the genome-containing titer suggests. This drives aggregation at lower apparent vg/mL thresholds than expected. Increasing the affinity pool volume to reduce total capsid concentration, allowed extended hold times without aggregation.

Published work confirms these issues are not unique. Engineered variants can produce unexpected analytical signals including temperature-dependent aggregate species that are variant-specific, with mitigation strategies that cannot be universally applied [4]. Pan-serotype affinity resins like POROS CaptureSelect AAVX have expanded purification options, but characterization has shown that engineered capsids may exhibit reduced or absent affinity depending on surface modifications [3].

We've seen across multiple teams that having someone on the sponsor side who has seen the failure modes, understands the uniqueness of novel capsids and how to apply this towards expedited PD; helps cut down surprises as teams move further down that development pipeline towards their first GMP batch.

Analytical Blind Spots with Novel Capsids

Analytical strategy is where platform assumptions create some of the most dangerous blind spots, because analytical gaps discovered late are among the most expensive and time-consuming to fix.

Potency. The most deferred and early CMC decision. Cell line, readout, and mechanism relevance all need consideration early. FDA guidance expects a quantitative, mechanism-relevant potency assay with increasing rigor toward licensure [6]. Teams that defer entirely scramble when IND timelines tighten and at that point, the work cannot be done faster.

Full/empty/partial capsid ratio. Standard approaches (AUC, SEC-MALS, cryo-EM) may require optimization for engineered variants. AEX-based separation of full, empty, and partial populations is capsid-variant-dependent [7]. Polishing strategy should be evaluated during PD, not assumed it fits.

Aggregation and stability. Engineered capsids may aggregate under conditions routine for natural serotypes. Structural modifications can produce temperature-dependent aggregate species that are variant-specific and unpredictable from wild-type behavior [4]. Stability across all stages; in-process, drug substance, drug product, in-use clinical, cannot be assumed from published serotype data. Leveraging early production history, even in the research setting can be beneficial in understanding the complexities early on.

What most teams miss entirely: in-process analytics during manufacturability screening. Most teams focus on titer, aggregation, and percent full after final bulk drug substance. What gets missed is how vectors behave during production; precipitates at hold points, unexpected affinity recoveries, full/empty/partial ratios at post-affinity (not

just post-polishing), capsid titers alongside vg titers for the total particle picture, host cell DNA/HCP/plasmid DNA clearance at intermediate steps, and process observations across constructs.

I build manufacturability screening panels capturing these in-process data points across multiple candidates from the earliest RUO productions. On one engagement, I helped a novel capsid company screening capsids A through D across multiple GOIs extract substantially more decision-quality data from initial productions than their standard analytical package provided. This gave them better capsid down-selection data and a clearer PD picture for each candidate.

The value of the screening panel was not just that the team had more data, it changed what they could decide. By capturing capsid titer alongside vg titer, post-affinity recovery, aggregation tendency, full/empty/partial composition, and intermediate clearance signals across capsids A through D, the team could separate "high expressing" candidates from candidates that were manufacturable. That distinction matters for early-stage groups because the best-looking construct at final bulk may not be the one that survives scale-up, polishing, hold-time constraints, or GMP facility fit.

In that case, the screening package gave the team a practical down-selection lens before they spent serious money on PD. They could avoid advancing a candidate simply because it produced attractive vg/mL at small scale and instead prioritize the capsid/GOI combination with the cleanest overall CMC profile. It also gave the CDMO a more useful starting picture: where recovery was being lost, where concentration stress might show up, which analytics needed tighter method attention, and which process observations had to be carried into the next development run.

This is the type of analytical strategy I set up early: a focused characterization package that tells the sponsor what to do next. Early in a program, that usually means pairing final product readouts with in-process readouts: vg titer, capsid titer or total particle readout, residual DNA/HCP/plasmid DNA where appropriate, affinity recovery, aggregation by SEC or orthogonal method, full/empty/partial profile, and simple stability or hold observations under the actual process conditions the team may need later. When those data are captured early, they become the baseline for tech transfer, CDMO comparison and process changes.

Tech Transfer: The Most Underestimated Risk

Tech transfer is the single point in early development where most assumptions collide with reality. Early-phase programs are challenged by transferring processes that do not fit into their external partner facilities, and for novel capsid programs, where some unique handling might need to be done, it gets even more complicated.

The real problem is this knowledge asymmetry. The sponsor has deep capsid science but limited process understanding. The CDMO has GMP infrastructure but no capsid history. Neither side fully understands what the other doesn't know.

How I've seen this play out:

Licensing and contracting create hidden constraints. I have seen teams locked into terms preventing them from switching CDMOs when processes didn't transfer; constraints buried in initial agreements. These constraints can be cell line IP, or process knowledge that either costs to transfer out, or cannot be transferred out. These are the typical tripping points when looking at CDMOs to engage with. In my practice we are always on the lookout for what licensing terms are whether it's during plasmid DNA manufacturing or RUO material generation at a CDMO. Understanding the implications of switching processes or moving from CDMO A to CDMO B and the impact that has on your vector, yield, process and timelines is important.

Brute-forcing material generation is not a strategy. Making RUO material quickly is fine and often needed. Choosing which CDMO to use and what you need this RUO material for is key. The gaps between RUO, PD, and GMP are where

timelines slip and costs go up. Understanding how the RUO material is being generated, how this translates to what PD work you will need to do and how that process needs to shift as you start thinking about GMP. I've chatted with companies who have gone the fast and cheap route for RUO material generation and used it in some early studies, only to find out that a lot of the process information from that workstream wasn't usable and small-scale work needed to be re-done in early PD.

Choosing the cheapest option may not always be the fastest. On a Series A-stage AAV program, the most important sponsor-side work was building the transfer plan earlier. The team needed speed, but the goal was not just to get material made. The goal was to make sure the RUO, PD, tox, and eventual GMP paths were connected enough that early decisions did not create avoidable rework later.

The work started with a facility-fit and process-fit review across the CDMO partners: what production system they would use, what scale they could realistically support, what parts of the process were platform versus adjustable, what analytics were included by default, what information would transfer with the material, and what licensing or process-knowledge constraints would matter if the sponsor needed to move. I mapped the intended use of each material lot to the level of process and analytical information required: discovery screening material, animal study material, PD-enabling material, and material that would meaningfully inform tox or GMP planning.

That approach gave the sponsor a faster path without treating speed and rigor as opposites. They could move quickly into material generation while preserving the information needed for PD, vendor comparison, and IND planning. The practical outcome was less duplicated work, fewer unknowns going into the next CDMO conversation, and a clearer line of sight from early material generation to the first GMP batch.

Raw Materials and Plasmid Strategy: Decisions That Compound

Plasmid source, quality grade, and testing strategy are among the earliest and most underappreciated program choices that get made.

Moving from research-grade to GMP-grade plasmids introduces comparability challenges. Topology, supercoiling, impurity profile, and lot-to-lot consistency can shift transfection efficiency, productivity, and product quality. For a novel capsid with a still-developing process window, these transitions can cause some headache.

The bacterial cell bank. Some plasmid suppliers use research cell banks (RCBs) for GMP-like production. RCBs don't always translate to the master cell bank (MCB) required for later-phase GMP supply. Plasmids from RCBs may not perform identically to MCB-derived material, requiring comparability work that adds cost and timeline. Working with a supplier who establishes an MCB early costs more upfront but de-risks the transition later.

The yield variability problem. I have seen teams switch plasmid grades and experience yield increases (good news) until the downstream process can't handle higher yields, causing lower recoveries, out-of-spec intermediates, or purification re-works. Better plasmid quality has an impact throughout the process.

The sourcing landscape is also evolving. There are now opportunities to reduce cost and time through overseas plasmid manufacturing, but this introduces logistics complexity and IND filing considerations that need to be understood upfront. There are also emerging alternatives to traditional plasmid DNA that can bypass the time, length, analytical requirements, and cost of producing a bacterial MCB, but these come with their own CMC risks around regulatory acceptance, comparability, and process impact. Having someone who knows the options and understands the risk profile of each can save both time and money at the moments you need it most.

A recent example: a sponsor needed the fastest, cheapest path to plasmid supply for Phase 1 AAV vector manufacturing without creating unnecessary IND risk. The request originated as a vendor vetting, but the real work was a plasmid strategy and raw-material risk assessment. I compared five plasmid vendors across cost, timing, quality grade, cell-bank strategy, analytical package, release testing, stability approach, geography, audit burden, and how closely each vendor's GMP-like or GMP-S process aligned with its eventual GMP process.

The recommendation was not simply "pick the lowest quote." We selected a path that used a GMP MCB even where some vendors offered RCB- or PCB-based options, because the MCB path reduced the risk of a later transition, despite adding time compared to the other options. We also prioritized vendors whose GMP-like/GMP-S process was similar or near-identical to their GMP process, so the sponsor was not forced into a major plasmid-process change later. That matters because plasmid topology, impurity profile, supercoiling, and lot-to-lot consistency can all affect transfection, vector productivity, and downstream behavior.

Within this deep dive we also aligned plasmid analytics against the expected IND-facing needs: cell-bank testing, plasmid stability, identity, purity, supercoiled content, residual host-cell impurities, endotoxin, sterility or bioburden expectations, and the documentation needed to support the vector manufacturer and the eventual Module 3. Because the selected option was ex-US, we also had to de-risk auditability, documentation access, shipping/logistics, and third-party QC testing for the MCB. The outcome was a plasmid path that met the initial ask (fast and cost-conscious) while avoiding the trap of a cheap supply chain that would need to be rebuilt before IND.

The Technical Case for Sponsor-Side CMC Support

The challenges in this paper: platform assumptions that break, analytical blind spots, tech transfer failures, and raw material decisions, share a common start: they require experienced judgment on the sponsor side of the CDMO relationship.

The common thread across these examples is that CMC support can aid in catching early decisions that have larger consequences later. This results in outcomes that early-stage teams care about: fewer repeated runs, fewer vendor surprises, fewer late analytical gaps, fewer avoidable change orders, and a cleaner path from early material generation into PD, tox, GMP, and IND. For novel engineered capsid programs, that is often the difference between a strong platform story and a manufacturable therapeutic program.

This is where ViroSpark fits. We operate on both sides of the table, inside the practical decisions that determine whether a promising capsid can move efficiently toward first-in-human development: what to measure, what to ask the CDMO, what to lock now, what to defer safely, and what early shortcut will become expensive later. For developers, the goal is not more CMC, it's the right CMC at the right time.

We are the CMC and TechOps bridge for early gene therapy teams particularly those working with novel engineered capsids. We are not a giant consulting shop. We are practical operators who, de-risk the program and hand back a clean plan that keeps milestones on track. We work inside the SOWs, the batch records, and the facility fit assessments as well as the high-level strategy decks. We bring a former CTO and senior analytical sciences specialist as needed for a plug-in TechOps team when needed.

We are built for pre-commercial budgets accessible to the teams that need experienced CMC guidance the most but have historically been priced out of getting it.

ViroSpark BioConsulting Sponsor-side CMC strategy and execution support for early gene therapy programs moving from discovery to FIH/IND.

Michael J. DiBiasio-White, PhD | MichaelWhite@virosparkbio.com | linkedin.com/in/michael-dibiasio-white | virosparkbio.com

References

- [1] Iwuchukwu, I. "Conquering the Complexities of AAV Manufacturing." GEN, May 2025.
- [2] PackGene Biotech. "AAV Gene Therapy and Capsid Engineering: Challenges, Strategies, and Directed Evolution." January 2026.
- [3] Adams, B. et al. "Structural characterization and epitope mapping of the AAVX affinity resin." Molecular Therapy Methods & Clinical Development, December 2024.
- [4] Mattison, C. "Unlocking Method Insights for Novel Engineered AAV Capsids." Cell & Gene, September 2025.
- [6] FDA. "Potency Tests for Cellular and Gene Therapy Products." Guidance for Industry.
- [7] "Separation of full, empty, and partial adeno-associated virus capsids using AEX." Journal of Chromatography B, July 2024.
- [8] "Ensuring Smooth Tech Transfer of Bioprocess Operations." BioPharma International, April 2020.