SOLVING THE VRU PROBLEM: TURNING VAPOR RECOVERY FROM LIABILITY TO ASSET

Michael Chavez & Brandon Dyck
Platinum Control

ABSTRACT

For decades, VRUs have been deployed as compliance equipment and treated as commodity hardware – sized on rough estimates, lightly engineered, and minimally monitored. The results have been predictable: inconsistent runtime, chronic loading instability, ever-increasing maintenance costs, and a long-standing belief that VRUs simply “don’t work.” This paper examines why that legacy persists and outlines a modern engineering and measurement framework that significantly changes VRU performance, economics, and reliability.

The approach centers on precise sizing, real-time operational data, and flexible deployment strategies designed to match dynamic vapor loads. A comprehensive operational dataset – including pressures, temperatures, load signatures, runtime behavior, oil level, and measured vapor flow – enables predictive maintenance, drift detection, and stable runtime across a wider range of operating conditions. Rather than relying on a single indicator, this data ecosystem provides the analytical foundation for a proactive VRU program. The same infrastructure also supports accurate emissions accounting, turning VRUs into valuable compliance assets as regulators move toward measurement-based methane reporting and as operators work to reduce the significant financial exposure tied to modern methane enforcement.

Field deployments in the Permian Basin show significant improvements in uptime, maintenance cost, and equipment longevity, with operators experiencing fewer cycling events, reduced downtime, and lower LOE. These results indicate that a measurement-driven VRU program can prevent both oversized and vapor-constrained installations, significantly reducing risk while improving economic return.

This paper will present the engineering principles, measurement insights, and early field learnings behind this next-generation approach – and why solving the long-standing VRU problem ultimately turns vapor recovery from a perceived liability into a measurable, high-value asset.

INTRODUCTION: WHY VRUS “DON’T WORK”

Vapor recovery units (“VRU’s”) have long occupied an ambiguous position within oil and gas production facilities. Unlike primary compression, artificial lift, or processing equipment, VRUs have historically been justified on regulatory or environmental grounds rather than as production-critical assets. This framing has strongly influenced how they are specified, installed, operated, and evaluated over time.

In many cases, VRUs are sized using limited production data, static assumptions, or conservative design margins intended to avoid short-term compliance risk rather than optimize long-term performance. Engineering rigor is often constrained by timelines, cost pressure, or regulatory urgency. Once deployed, instrumentation is frequently minimal, leaving operators without direct insight into how the unit is behaving under real operating conditions. Ongoing performance is inferred indirectly through vessel pressures, flare activity, or anecdotal runtime observations rather than verified through direct measurement.

Over time, this approach has produced predictable results. Operators encounter inconsistent runtime, frequent cycling, elevated maintenance requirements, and premature mechanical wear. Units may shut down during peak vapor events or struggle to operate continuously during lower-load conditions. Maintenance interventions become reactive, downtime increases, and confidence in the equipment erodes. These experiences reinforce a long-standing industry narrative that VRUs are unreliable, fragile, or fundamentally incompatible with real-world operating conditions.

However, the commonly repeated adage that “VRUs don’t work” is only true in environments where performance is not measured and accountability is diffuse. When VRU operation cannot be observed directly, responsibility for poor performance is easily displaced. Cycling, shutdowns, or reduced capture are attributed broadly to “the VRU” rather than to specific, correctable causes such as application mismatch, operating envelope violations, facility constraints, or upstream and downstream interactions.

When performance is measured – including runtime, cycling behavior, pressure response, and load stability – ambiguity disappears. Performance becomes observable. Trends become defensible. Accountability becomes unavoidable. The discussion shifts from whether VRUs work to who is responsible for correcting deviations when they do not.

This paper contends that the historical perception of VRU failure is not a failure of vapor recovery itself, but a failure of engineering discipline driven by limited measurement, limited visibility, and limited ownership.

HISTORICAL FRAMING AND THE COMMODITIZATION OF VAPOR RECOVERY

Vapor recovery was not originally introduced as a marginal or experimental function. Early VRU deployments were selective and typically justified on clear economic and engineering grounds. Operators evaluated gas composition, expected vapor volumes, sales constraints, and recovery economics before installing equipment. In these early applications, VRUs were treated as engineered compression systems designed for specific operating conditions and economic objectives (Nelson, n.d.).

This approach began to shift with the introduction of EPA Subpart OOOO (U.S. EPA, 2012) and subsequent regulatory expansions, including Subpart OOOOa (U.S. EPA, 2016) and the more recent Subparts OOOOb and OOOOc (U.S. EPA, 2024). Compliance-driven timelines accelerated VRU deployment across a much broader

range of facilities and operating conditions. As operators moved quickly to satisfy regulatory requirements, VRUs were increasingly deployed as standardized packages rather than site-specific compression systems engineered for long-term operation.

Subsequent changes in federal and state regulatory frameworks reinforced this pattern. VRUs were installed, deprioritized, or removed based on evolving compliance needs rather than operational fit. Over time, this volatility trained the industry to view VRUs as equipment that could be turned on or off as regulations evolved, rather than as continuous-duty assets requiring stable operation and lifecycle management.

This framing also distorted economic evaluation. As vapor recovery became associated primarily with regulatory obligation, decision-making increasingly emphasized lowest upfront cost rather than lifecycle performance. Equipment selection favored capital efficiency over operational robustness. In practice, this often resulted in systems less tolerant of dynamic vapor loads and more prone to intermittent operation or cycling.

Reduced runtime directly limits vapor capture while fixed costs remain, undermining the economic premise of vapor recovery. Maintenance costs increase, reliability declines, and the perception of VRUs as a necessary liability is reinforced. At its core, the belief that VRUs “do not work” stems from how vapor recovery has historically been framed and deployed, not from any fundamental limitation of the technology itself.

MEASUREMENT VERSUS ESTIMATION: WHY ACCURACY IS THE WRONG DEBATE

Historically, most VRU-related emissions and performance assessments have relied on estimation rather than direct measurement. Engineering calculations, emission factors, runtime assumptions, and inferred performance have been accepted as sufficient for regulatory reporting and internal analysis. In many cases, these estimates were the only practical option available given instrumentation limitations and cost constraints.

However, estimation-based approaches introduce inherent uncertainty that becomes problematic as scrutiny increases. Estimates depend on assumptions about operating conditions, gas composition, runtime consistency, and system behavior that often diverge from reality. This gap between estimated and actual emissions is not theoretical – large-scale assessments of the U.S. oil and gas supply chain have found that actual methane emissions are approximately 60% higher than EPA inventory estimates, driven in large part by the failure of facility-level estimation methods to capture emissions released during abnormal operating conditions (Alvarez et al., 2018). When performance deviates, it is difficult to determine whether the deviation reflects actual system behavior or limitations in the estimation method itself.

This has led to a common but flawed framing of the problem as one of “accuracy.” Discussions often focus on whether a given measurement approach is sufficiently precise or whether instrumentation meets a particular accuracy threshold. In practice, this misses the more important distinction between precision and defensibility.

Defensibility, not absolute accuracy, is the critical requirement. A measurement does not need to be custody-transfer grade to be valuable. It must be supportable, repeatable, calibrated, and transparently reported within known uncertainty bounds. An estimated value with unknown error margins is far less defensible than a measured value with explicitly defined uncertainty. Even estimates with known error margins, while more rigorous, still fall short of measurement-based approaches in critical ways. If actual production deviates from the assumptions underlying those estimates – as it frequently does – permitted volumes, compliance calculations, and reported emissions all become unreliable regardless of how well-bounded the original estimate appeared. The gap between estimated and actual runtime, in particular, compounds through every downstream calculation that depends on it.

Measurement closes the credibility gap between what is assumed and what actually occurs. It allows operators to demonstrate consistency over time, identify deviations, and explain performance with confidence. As emissions reporting frameworks, audits, and corporate governance processes evolve, this defensibility becomes increasingly important.

Importantly, measurement also reveals operational truths that estimation obscures. Runtime assumptions collapse when actual cycling behavior is observed. Systems sized and calculated on the basis of continuous operation produce fundamentally different thermal and mechanical profiles when measurement reveals the unit actually operates in batch or intermittent cycles. Temperature stability, in particular, is highly sensitive to batch operations – conditions cannot remain consistent when a unit cycles on and off, yet every downstream calculation assumes they do. Load calculations are challenged when measured pressures and flows deviate from design expectations. These insights are not academic; they directly inform engineering decisions, maintenance strategies, regulatory and permitting decisions, and operational accountability.

COMMON FAILURE MODES IN LEGACY VRU DEPLOYMENTS

Given the historical framing described above, the failure modes commonly observed in legacy VRU installations are not anomalies. They are predictable outcomes of how these systems have been specified and operated. These issues generally fall into three closely related categories: improper application sizing, intermittent operation driven by cycling, and limited operational visibility.

Vapor generation at tank batteries is inherently dynamic. It is influenced by production rates, fluid properties, tank configuration, ambient temperature, operating practices, and facility design. When VRUs are sized based on static or overly conservative assumptions, they are often mismatched to real operating conditions.

Undersized units struggle to keep pace with peak vapor loads, leading to pressure excursions, shutdowns, or flaring events. Oversized units struggle to operate continuously, frequently cycling as vapor availability fluctuates. This cycling dramatically

accelerates mechanical wear, increases maintenance frequency, and destabilizes operation of both the VRU and the facility on which it is deployed.

Vapor composition further complicates the operating environment. Unlike most compression applications where gas quality is relatively stable and well-characterized, VRUs frequently encounter variable and challenging inlet conditions – including high liquid content, fluctuating BTU values, elevated H_2S or CO_2 concentrations, and temperature-driven changes in vapor behavior. These composition variables directly affect compressor health, oil life, valve performance, and overall system longevity. A unit that is correctly sized for a given volume may still struggle or fail if the vapor it is compressing falls outside the mechanical and chemical tolerances of its configuration. Sizing and measurement frameworks that account only for volume without considering composition risk solving only half the problem.

These dynamics are well-documented in the EPA Natural Gas STAR Program's technical guidance on vapor recovery unit applications (U.S. EPA, 2003; U.S. EPA, 2023).

Compounding these issues is the absence of meaningful operational visibility. Many VRUs are deployed with limited instrumentation, leaving operators without insight into load behavior, runtime trends, or early indicators of mechanical degradation. Issues are addressed reactively after failures occur rather than proactively as performance begins to drift.

Together, these factors create a reinforcing feedback loop. Poor performance reinforces skepticism. Skepticism discourages investment in better engineering or monitoring. The cycle repeats.

REFRAMING VRUS AS CONTINUOUS-DUTY COMPRESSION SYSTEMS

From an engineering standpoint, a vapor recovery unit is a compressor operating under highly variable inlet conditions. Like any compression system, its reliability and longevity depend on operating within defined envelopes, managing load transitions, and maintaining mechanical stability over time.

Historically, VRUs have often been treated as intermittent or secondary equipment. Nameplate horsepower or nominal capacity becomes a proxy for suitability, while dynamic operating behavior receives limited attention. In practice, factors such as suction pressure variability, discharge constraints, vapor composition changes, and operating regime play a far greater role in determining performance and reliability than static ratings alone.

Reframing VRUs as continuous-duty compression assets fundamentally changes the engineering problem. Sizing must reflect the full expected range of vapor loads rather than isolated design points. Control strategies must accommodate variability rather than react to it. The objective is not to eliminate dynamic behavior, but to ensure stable operation across it without repeatedly entering off-design conditions.

This reframing also clarifies responsibility. When VRUs are treated as engineered assets, performance expectations become explicit, operating envelopes are defined, and deviations demand explanation. This sets the stage for measurement-driven accountability.

MEASUREMENT ARCHITECTURE AND DATA ECOSYSTEM

Stable operation cannot be achieved without visibility, and visibility cannot exist without measurement. Measurement is essential not only for fault detection, but for confirming whether a VRU is operating within its intended envelope over time and for establishing accountability when it is not.

A comprehensive measurement architecture includes suction and discharge pressures, temperatures at critical points across the package, compressor load signatures, runtime behavior, oil level, and direct measurement of vapor flow where feasible. Individually, these signals provide limited insight. Collectively, they create a behavioral fingerprint of the VRU system.

Correlated analysis of multiple signals allows operators to identify drift, instability, or degradation long before a failure occurs. Subtle changes in load signature combined with temperature and pressure trends may indicate developing mechanical issues, facility constraints, or application mismatch.

Measurement fundamentally changes the nature of performance discussions. When runtime, cycling frequency, and pressure response are continuously recorded, it becomes difficult to attribute poor outcomes to vague or external causes. Measurement forces specificity. It becomes clear whether instability is driven by sizing, facility constraints, operating practices, or equipment condition.

VRU performance is no longer isolated to an eight-by-fourteen skid. It becomes a system-level responsibility involving tank behavior, facility design, control strategy, and operational discipline.

ACCOUNTABILITY AND OWNERSHIP IN A MEASURED SYSTEM

Measurement does more than improve diagnostics; it changes behavior. In unmeasured systems, accountability is diffuse. Performance issues can be attributed to design, operations, environment, or regulation without resolution. In measured systems, responsibility becomes explicit.

When runtime, cycling, and load behavior are visible, it becomes difficult to assert that VRUs “do not work” without addressing the underlying causes. Measurement forces conversations about ownership: who is responsible for maintaining operating envelopes, who controls facility constraints, and who is accountable for corrective action when deviations occur.

Importantly, this accountability must apply in all directions. A service provider operating within a measurement-driven framework cannot credibly hold operators accountable for

facility conditions while ignoring evidence that its own equipment is misapplied. If operational data indicates that a unit is oversized for current conditions – cycling excessively, consuming disproportionate maintenance resources, and generating costs that a properly sized unit would not – the responsibility to initiate that conversation falls on the party with visibility into the data. In a leased equipment model, that may mean proactively recommending a smaller unit at a lower lease rate. For customer-owned equipment, it may mean flagging the mismatch so that the operator's internal teams are not expending resources on a problem that engineering, not maintenance, should resolve.

This kind of self-directed accountability is not incidental to the framework. It is foundational. When operators see that measurement is used to identify and correct misapplication regardless of where the responsibility falls – including when it falls on the service provider – confidence in the system increases. Accountability becomes a shared discipline rather than a one-directional expectation, and the resulting conversations shift from adversarial to collaborative.

The inverse is equally important. When measurement clearly identifies that performance deviations are driven by facility conditions, operating practices, or upstream and downstream constraints outside the VRU provider's control, that evidence must be available and defensible. Operators benefit from this clarity as much as service providers do – it replaces assumption-driven troubleshooting with specific, actionable insight and ensures that corrective efforts are directed where they will actually resolve the issue.

Underlying all of this is a basic requirement: operational data must be accessible to all parties. Measurement-driven accountability loses its meaning if the data is controlled exclusively by one side of the relationship. When operators have direct access to the same runtime, pressure, cycling, and alarm data that informs the service provider's decisions, performance claims can be independently verified, and the framework operates on trust that is earned through transparency rather than assumed through a contract. Without that access, accountability reverts to assertion – and assertion, in the absence of shared visibility, is indistinguishable from the unmeasured systems this framework seeks to replace.

With the addition of modern analytical tools capable of continuously evaluating operating data, oversized or misapplied units can be flagged automatically rather than relying on monthly field checks to identify what is often a subtle and gradual mismatch. This extends the accountability framework beyond human observation and into persistent, data-driven oversight – ensuring that application drift is identified early and addressed before it compounds into the chronic cycling, elevated maintenance, and eroded confidence that characterize legacy deployments.

OPERATIONAL BEHAVIOR OF A STABLE VRU SYSTEM

When VRUs operate within engineered envelopes and are supported by adequate measurement, their behavior changes markedly. Runtime becomes stable rather than

intermittent. Cycling events decrease. Maintenance intervals extend. Operators gain confidence in performance, and intervention shifts from reactive to proactive.

Critically, each of these behavioral shifts is directly observable through the measurement architecture described above. Availability – the percentage of time a unit is running versus shut down or alarmed – is the most immediate indicator of system health. In a stable deployment, availability should trend upward toward sustained high levels as the unit and facility reach equilibrium, rather than eroding over time as is typical in unmeasured systems. Alarm frequency provides a complementary signal: a unit operating within its intended envelope should generate few alarms, and a rising alarm count is among the earliest measurable indicators that conditions are drifting away from design intent.

Consumable consumption rates offer a deeper view into mechanical stress that may not yet manifest as downtime. Oil consumption, in particular, functions as a reliable proxy for internal wear, seal condition, and thermal loading. In a properly matched system, oil consumption should remain low and stable over time. A sustained increase signals that mechanical stress is accumulating – often well before it results in a failure event. Similarly, cycling frequency and duration provide direct evidence of whether a unit is sized appropriately for its vapor load. A unit that cycles frequently is operating outside its intended duty, regardless of whether it appears “available” in a simple uptime calculation.

Taken together, these indicators define what “stable” looks like in measurable terms: high and trending availability, minimal alarm activity, low and consistent consumable usage, and infrequent cycling. When these conditions are present, the VRU is operating as designed. When they are absent or deteriorating, the measurement record identifies where and why – providing the basis for the targeted interventions and fleet-level assessments described in the following section.

These characteristics are not exceptional. They reflect the expected behavior of properly engineered compression systems operating within defined limits.

FIELD OBSERVATIONS FROM PERMIAN BASIN DEPLOYMENTS

Field deployments in the Permian Basin applying this framework have demonstrated consistent improvements across both individual site performance and fleet-level operational efficiency. Two cases are presented below – one illustrating the impact of measurement-driven diagnosis at a single facility, and one demonstrating the value of systematic fleet reassessment. Together, they represent complementary applications of the same underlying principles.

Site-Level Intervention

To illustrate the impact of this framework at the individual site level, consider a 150-horsepower VRU deployed at a production facility operated by a major Permian Basin operator. The site had a history of elevated maintenance demands and was exhibiting a clear pattern of operational degradation at the time of evaluation.

Over an observation period of approximately four months prior to intervention, availability trended downward from near-100% to approximately 92%, with monthly alarm counts escalating from single digits to over twenty-five in a single month. Compressor oil consumption – a reliable proxy for mechanical stress and system health – increased steadily over the same period, with total consumption reaching 262 cumulative gallons. The unit was not yet in outright failure, but every leading indicator pointed in the same direction: accelerating degradation that, left unaddressed, would result in extended downtime, component failure, or both.

Rather than treating these symptoms in isolation – replacing consumables, resetting alarms, reacting to individual shutdown events – the service provider conducted a system-level evaluation of the VRU's operating conditions, including off-skid facility interactions, operating envelope alignment, and mechanical configuration. Working closely with the operator, the team coordinated adjustments to facility setpoints and operating practices that directly influenced VRU performance, while also installing a proprietary oil system modification to address the accelerating oil consumption and associated mechanical stress. The intervention was not a single equipment change, but a comprehensive realignment of the VRU within its operating environment.

Over an equivalent observation period following the intervention, the results were significant and sustained. Monthly availability stabilized above 97% and trended upward, exceeding 99% in the final months of observation. Alarm frequency dropped to near zero. Total oil consumption across the entire post-intervention period fell to 9 gallons – a reduction of over 96% compared to the pre-intervention period. Downtime effectively ceased.

This case reinforces a central premise of this paper: when VRU performance is measured, and when operational visibility enables root-cause diagnosis rather than reactive maintenance, the equipment performs as designed. The improvements observed were not the result of a single component change or increased maintenance intensity. They followed from a disciplined, collaborative process – evaluating the full operating environment, coordinating with the operator on facility conditions, and making targeted engineering modifications based on what the data revealed. The perception of inherent unreliability gives way to documented, repeatable operational stability.

Fleet-Level Right-Sizing

While the preceding example illustrates the impact of measurement-driven intervention at a single site, the same principles apply at fleet scale. A second case involving a separate major Permian Basin operator demonstrates how systematic evaluation of an existing VRU fleet can identify and correct widespread application mismatch – and how the resulting improvements benefit operators and service providers alike.

The operator had accumulated 49 VRUs across a single county over an extended period. As is common in mature or evolving fields, the original sizing basis for many of these units no longer reflected actual operating conditions. In some cases, production had declined over time, leaving units oversized relative to current vapor loads. In others,

units had been placed in anticipation of production that had since come online under different conditions than originally expected. The net result was a fleet in which the majority of units were operating outside their intended envelopes – oversized units cycling frequently due to insufficient vapor supply, undersized units struggling to maintain pace with actual loads, and the associated consequences of both: elevated maintenance demand, reduced availability, and inefficient allocation of field service resources.

A comprehensive assessment was conducted across all 49 units, evaluating each against current production data, facility conditions, and projected operating requirements. The results were significant: only 10 of the 49 units were determined to be appropriately sized and located. The remaining 39 required action – relocation to facilities where their capacity better matched actual vapor loads, or replacement with different sized equipment appropriate to current conditions. In many cases, replacement meant bringing in a smaller unit better suited to the site's actual vapor generation, which in turn rendered the original, larger unit unnecessary. Eight units were released from contract entirely as a direct result of this process – not because vapor recovery was no longer needed at those facilities, but because the properly sized replacements eliminated the need for the oversized equipment they replaced. These releases were recommended proactively by the service provider, even though the replacement units carried lower monthly lease rates. All moves and replacements were completed by the end of the month following project initiation.

The outcomes following completion were consistent across the fleet. The operator experienced meaningful improvements in VRU availability as units were no longer forced to operate in conditions that induced cycling or frequent shutdown. Monthly LOE decreased at affected facilities as oversized equipment was replaced with appropriately matched units at lower lease rates. Operational disruptions declined, and the operator's field teams were able to redirect attention from VRU issues to other operational priorities.

From the service provider's perspective, the benefits were equally clear. A properly matched fleet required significantly less intervention – fewer emergency calls, fewer parts replacements driven by off-envelope operation, and more predictable service schedules. The same field resources that had previously been consumed by reactive maintenance on mismatched equipment could now be deployed more efficiently across the remaining fleet.

This case underscores a broader point: right-sizing is not a one-time design exercise. Vapor loads change as fields mature, production shifts, and facility configurations evolve. Without periodic reassessment grounded in measured operating data, fleets drift toward mismatch over time – and the resulting inefficiencies compound across every unit, every service call, and every maintenance event. Measurement-driven fleet management treats sizing as a continuous process rather than an initial assumption, and the results demonstrate that discipline benefits both operators and service providers when the engineering is corrected.

EMISSIONS, GOVERNANCE, AND AUDIT DEFENSIBILITY

Methane capture performance is increasingly shaped by internal governance, corporate reporting, and external scrutiny. In this context, stability and observability of VRU operation become as important as vapor recovery itself.

Frequent shutdowns and cycling increase the likelihood of unplanned venting or flaring. Conversely, VRUs operating consistently within defined envelopes reduce these events by maintaining control across steady-state and transient conditions.

Measurement alters the accountability model for emissions. When performance is inferred, emissions estimates are debated. When performance is measured, deviations require explanation. Systems that are continuously monitored are easier to explain, audit, and defend over time.

The operational data described in the preceding case studies illustrates this directly. When a VRU's suction and discharge pressures, runtime, and alarm history are continuously recorded, the cause of an exceedance event can be identified and attributed with specificity. A discharge pressure spike, for example, may indicate a downstream constraint entirely outside the VRU operator's control – a distinction that is invisible without measurement but immediately apparent with it. Conversely, a suction pressure anomaly may point to an on-site facility issue that the operator can address directly. In either case, the measured record transforms an emissions event from an unexplained liability into a traceable, explainable occurrence – and that distinction carries significant weight in audit and enforcement contexts.

This distinction matters for regulatory reporting as well. Recurring "VRU failures" that cannot be adequately explained risk being excluded from consideration as affirmative defense demonstrations or alternative compliance pathways. Measured data provides the granularity needed to distinguish genuine equipment failure from system-level events beyond the VRU's control – and to demonstrate that pattern through a defensible, time-stamped record rather than after-the-fact narrative.

This concern is not hypothetical. In a 2024 enforcement action involving a major operator across fifteen facilities in New Mexico, state regulators identified dozens of VRU-related excess emission events and determined that recurring VRU equipment failures were not eligible for consideration as Abnormal, Dangerous, or unusual Developments (ADDs) under state air quality regulations (NMED, 2024). Without measured data to differentiate root causes across individual events, the operator was unable to demonstrate that the failures were isolated or anomalous – and the full pattern was included in the enforcement action. Notably, some of those events may well have been attributable to the VRU itself. But without measurement, all of them were – because there was no data to argue otherwise.

ECONOMIC IMPACTS OF MEASUREMENT-DRIVEN OPERATION

Improved operational stability reduces maintenance frequency and lease operating expense while supporting more predictable vapor capture. Disciplined sizing and

visibility reduce capital misallocation associated with oversizing or vapor-constrained installations.

The fleet-level assessment described above illustrates these dynamics directly. When a 49-unit fleet was systematically evaluated against actual operating conditions, the resulting right-sizing effort reduced the operator's monthly facility costs at affected sites, decreased maintenance-driven interventions, and freed field personnel to focus on other operational priorities. These gains were not achieved through increased spending or additional equipment – they followed from removing the inefficiencies that mismatched sizing had embedded across the fleet over time.

At the individual site level, the same principles apply. Addressing root-cause operating conditions at a single facility eliminated the recurring cycle of alarm response, consumable replacement, and reactive maintenance that had driven steadily increasing costs. The 96% reduction in oil consumption observed at that site was not merely a consumables savings – it was an indicator that the mechanical stress driving premature wear and component failure had been substantially resolved, with corresponding implications for long-term equipment longevity and total cost of ownership.

These economic outcomes are not the primary objective of this framework, but natural consequences of applying engineering rigor to systems historically managed as commodities. When VRUs are properly sized, continuously monitored, and operated within defined envelopes, the economics follow the engineering.

These results also underscore a broader economic reality. Even in periods of low commodity prices, properly operated vapor recovery generates meaningful revenue from captured gas that would otherwise be flared or vented. When gas prices are favorable, the economics become even more compelling. When regulatory mandates require vapor recovery regardless of price environment, measurement ensures that the compliance investment also functions as a producing asset rather than a sunk cost. The case studies presented in this paper illustrate that these economics are not theoretical – they follow directly from the engineering discipline and operational visibility described throughout.

CONCLUSIONS AND FUTURE WORK

VRUs have historically underperformed not because vapor recovery is inherently flawed, but because these systems have not been engineered, measured, or operated with the rigor applied to other compression assets. A measurement-driven approach grounded in accountability fundamentally changes VRU behavior, reliability, and value.

The field observations presented in this paper – spanning both individual site intervention and fleet-level reassessment – demonstrate that when VRUs are engineered and sized to actual conditions, monitored continuously, and managed within defined operating envelopes, the outcomes are consistent: higher availability, lower maintenance costs, reduced cycling, and measurable improvements in operational

stability. These are not exceptional results. They are the expected behavior of compression systems that are properly engineered and supported by adequate visibility.

Equally important, this framework changes the nature of the conversations surrounding VRU performance. Measurement replaces ambiguity with specificity. Accountability becomes shared and directional rather than diffuse. Operators and service providers alike benefit from a system in which performance is observable, deviations are explainable, and corrective action is informed by data rather than assumption.

The reinforcing feedback loop described earlier in this paper – where poor performance breeds skepticism, skepticism discourages investment, and the cycle repeats – can only be broken by introducing visibility. Measurement is how the industry exits this loop: it replaces assumption with observation, perception with evidence, and blame with accountability.

Future work will focus on deeper integration of operational measurement with emissions accounting frameworks – connecting continuous VRU performance data to methane quantification methodologies and regulatory reporting structures. As measurement-based reporting frameworks continue to evolve, the operational infrastructure described in this paper positions VRUs not only as reliable compression assets, but as auditable, defensible components of an operator’s broader emissions management strategy.

REFERENCES

Alvarez, R. A. et al. (2018). “Assessment of methane emissions from the U.S. oil and gas supply chain.” Science, 361(6398), 186–188. doi: https://doi.org/10.1126/science.aar7204.

Nelson, R. N. (n.d.). “Analysis of Vapor Recovery Installations.” Paper presented at the Southwestern Petroleum Short Course, Lubbock, Texas, USA.

New Mexico Environment Department, Air Quality Bureau. (2024). Notice of Violation, Compliance Case Involving Fifteen Facilities in New Mexico. May 20, 2024. Available at: https://service.web.env.nm.gov/urls/UqEjFJeK.

U.S. Environmental Protection Agency. (2003). “Installing Vapor Recovery Units on Crude Oil Storage Tanks.” Natural Gas STAR Program Lessons Learned Series, EPA-430-B-03-015. Available at: https://www.globalmethane.org/documents/m2mtool/docs/ll_final_vap.pdf.

U.S. Environmental Protection Agency. (2012). “Standards of Performance for Crude Oil and Natural Gas Production, Transmission and Distribution.” 40 CFR Part 60, Subpart OOOO, 77 Fed. Reg. 49490. August 16, 2012.

U.S. Environmental Protection Agency. (2016). “Oil and Natural Gas Sector: Emission Standards for New, Reconstructed, and Modified Sources.” 40 CFR Part 60, Subpart OOOOa, 81 Fed. Reg. 35824. June 3, 2016.

U.S. Environmental Protection Agency. (2023). “Vapor Recovery Units.” Natural Gas STAR Program. Available at: https://www.epa.gov/natural-gas-star-program/vapor-recovery-units.

U.S. Environmental Protection Agency. (2024). “Standards of Performance for New, Reconstructed, and Modified Sources and Emissions Guidelines for Existing Sources: Oil and Natural Gas Sector Climate Review.” 40 CFR Part 60, Subparts OOOOb and OOOOc, 89 Fed. Reg. 16820. March 8, 2024.