



WORKING PAPER · SEPTEMBER 2026

The Request for Value

The one page procurement never asked for, where it goes in the process, and how to issue it Monday

Michael Brey

Value Lattice

CONTENTS

Summary	3
What this means for a buyer, in one page	3
1 The request that is missing	4
2 What an RFP asks, and what it cannot	6
3 The Request for Value	7
4 Why the definition has to be the buyer's	9
5 Where it goes in the process	10
6 Who writes it, and what it forces	11
7 The seller's objection, and why adaptation pays	12
8 Attribution, the hard part	13
9 Issuing the first one	14
10 From a step to a standard	14
Appendix A. The Request for Value	16
About the Author	17

Summary

Procurement's vocabulary has a request for information, a request for quotation, and a request for proposal. Each of them asks the supplier something. None of them makes the buyer's own side write down what the purchase has to produce once it is running, how that will be measured, and what evidence will count. That document does not exist in the standard sequence, so the definition of success arrives from the vendor inside the proposal, and the organization negotiates price and terms around it.

This paper describes the missing request. The Request for Value, or RFV, is a single page the business owner writes before the market is asked anything. It has five components: the outcomes the purchase must produce, the metrics that will measure them, the conditions that must hold for those outcomes to be possible, the evidence suppliers must supply, and the criteria on which the comparison will be made. It augments the RFP rather than replacing it. It adds no scoring matrix and no weighted-criteria overhead. It can be adopted on one decision in one category, with no program behind it, no platform under it, and nobody hired to run it.

The paper's argument is that the RFV is the cheapest step in procurement and the one whose absence costs the most, because everything downstream, the contract, the governance cadence, the renewal, and the after-the-fact verification, has to reference a definition of success that somebody on the buyer's side wrote. Where no such definition exists, verification is impossible in principle, and the value of the purchase becomes whatever the seller's dashboard says it is.

Sections 1 through 4 make the case and describe the instrument. Sections 5 through 8 place it in the process, address the seller's objection, and treat attribution, which is the hard part. Section 9 sets out how to issue the first one. Appendix A is the page itself, reproduced as a form.

What this means for a buyer, in one page

If you buy enterprise technology or services, this paper argues five things you can act on.

Your process has no step where your side defines value. The RFI asks what the market offers. The RFQ asks what it costs. The RFP asks how the supplier would deliver. All three point outward. The step where the business owner writes what the purchase

must produce, in the organization's own terms, does not exist in most procurement sequences, and that is why the definition arrives pre-written from the vendor.

The missing step is one page. Five components, written before any supplier is asked anything: outcomes, metrics, conditions, evidence, comparison criteria. The discipline is in the clarity, not the length.

It goes between the requirement and the RFP, and it travels. The RFP attaches it. The contract references it. The governance cadence reviews against it. The renewal tests it. The same page, unchanged unless somebody writes down why, is what makes each of those later documents comparable to the first.

It changes what vendors compete on. With an RFV in the packet, vendors respond to the buyer's dimensions rather than constructing their own. They compete on alignment and evidence rather than on narrative. Sellers who adapt early will do better than those who defend the old arrangement.

You can issue the first one Monday. One category where outcomes can be measured within a quarter. One decision that is upcoming, not retrospective. One provisional baseline, marked as provisional. One page. The rest of the standard grows from there, but the page comes first, and it is enough to start.

The rest of the paper is the argument and the instrument.

1. The request that is missing

Every procurement professional learns three acronyms early. RFI, RFQ, RFP. Information, quotation, proposal. The sequence is taught as a progression: understand what the market can offer, then find out what it costs, then ask how each supplier would deliver it. A good buyer chooses the method by the maturity of the requirement rather than by habit, and the profession's own training says so.

Look at the three from a different angle and something else appears. Each request is addressed to the supplier. Each asks the supplier for something the supplier holds: its capabilities, its price, its approach. That is proper. Procurement's job is to reach the market, and the market answers with proposals.

What none of the three does is turn around and ask the buyer's own organization for anything. There is no request in the standard sequence that requires the business owner, the person who will live with the outcome, to write down what the purchase must produce, how anyone will know, and what would count as proof. That document is assumed. It is assumed to exist somewhere in the business case, or in the requirement, or in the sponsor's head. In practice it is usually absent, and its absence is invisible because the proposals arrive on time and the comparison proceeds.

The consequence is structural. When the buyer has not written the definition of success, the definition is supplied by whoever writes the most complete document in the packet, and that is the vendor. The proposal states what the solution will achieve. The vendor's business case, often built by the vendor's own value engineering team, states the return. The organization reviews those numbers, adjusts an assumption or two, and negotiates hard on price and terms. What counts as value, how it is measured, and when it may be declared achieved was never an item on the list, because there was no request that put it there.

A definition is not neutral. Whoever writes it decides what gets counted and what gets left out. If value is declared against the seller's definition, the outcome was chosen when the template was.

This paper's claim is narrow. The gap is not one of competence or intent on either side. It is a missing step, and a missing step can be added.

	<i>The request asks</i>		<i>Addressed to</i>
RFI	What the market can offer	→	the supplier
RFQ	What it will cost	→	the supplier
RFP	How the supplier would deliver it	→	the supplier
RFV	What it must produce, and how anyone will know	←	the buyer's own side

Figure 1. The fourth request. Three requests in the standard sequence are addressed to the supplier. The fourth is addressed to the buyer's own side, and it is the one the sequence has never had.

2. What an RFP asks, and what it cannot

The Request for Proposal is a good instrument for what it was built to do. It structures a competitive comparison of how different suppliers would meet a stated requirement. It carries the institutional weight of fairness: every supplier answers the same questions, in the same format, by the same date. It supports credentials, references, security review, and pricing. Procurement functions have refined it over decades, and nothing in this paper argues for taking it away.

What the RFP cannot do follows from how it is addressed. It asks the vendor how it will deliver. It does not require the buyer to define what success will look like before the vendor answers. The result is a process where vendors construct the dimensions of the answer, and the buyer evaluates responses against criteria that surface during the comparison rather than before it. By the time the wrong criteria are recognized, the comparison is complete, and often the contract is signed.

Three things go missing in that order of operations.

The first is the outcome. An RFP specifies the solution, sometimes in great detail, and says little about the business result the solution was bought for. Reduce days sales outstanding by a stated number. Cut the time from order to ship. Raise first-contact resolution. Those are outcomes. “Implement a modern order management platform with the following capabilities” is a solution, and a proposal can meet every line of it while the outcome never arrives.

The second is the measure. Even where an outcome is named, the RFP rarely states how it will be measured, from what baseline, with what data, at what frequency, and what threshold would trigger a review. Without the measure there is nothing to miss, and nothing to review.

The third is the evidence standard. An RFP asks for references and case studies. It rarely says what a claim must be supported by to count: outcomes at comparable customers under comparable conditions, the distribution of results rather than the best case, the places where the solution underperformed. Absent that standard, every vendor supplies the best three stories it has, and the buyer cannot compare them.

Each of these three could be added to an RFP as a section. In practice they are not, for a reason that has nothing to do with the RFP’s design. They are the buyer’s to write, and the RFP is a document the buyer issues to the market, not to itself. A separate request, addressed inward first and then attached outward, is the cleaner instrument. That is the Request for Value.

3. The Request for Value

The Request for Value is a buyer-authored statement, written before the market is engaged, that defines what a purchase must produce and how the organization will know whether it did. It augments the RFP rather than replacing it. The RFP keeps its role for capability, credentials, and pricing. The RFV adds the outcome and evidence layer the RFP was never built to carry.

It has five components.

Outcomes. The specific business results the purchase must produce, stated in the organization’s own terms and owned by a named business function. Not features, not

capabilities, not the vendor's language. If the purchase succeeds, what is true afterward that is not true now.

Metrics. How each outcome will be measured: the indicator, the data source, the baseline value and the date it was taken, the target, the frequency of measurement, and the threshold that triggers intervention. A metric without a baseline is a slogan.

Conditions. What must hold for the outcomes to be possible. Adoption assumptions, data quality, process changes the organization itself must make, dependencies on other systems, the people who must be in place. Conditions are the honest half of the page. They record that the outcome depends on the buyer as well as the seller, and they are where attribution gets settled later.

Evidence. What suppliers must provide for a claim to count. Realized outcomes at comparable customers, under stated conditions, with the distribution rather than the best case. Where the solution has underperformed and why. What the supplier will commit to measuring after go-live and how the buyer will access that data.

Comparison criteria. How the responses will be weighed against one another on the four components above. Not a scoring matrix. A short statement of what will decide the comparison, written before the responses arrive so it cannot be reverse-engineered from them.

Two properties matter as much as the components.

The RFV is short. One page is the target, and two is the limit. The discipline is in the clarity. A page that a CFO can read in four minutes and a vendor cannot misread is worth more than a forty-page requirements document nobody re-reads after award.

The RFV carries no new scoring apparatus. It introduces no parallel evaluation system, no weighted criteria, and no vendor rating grid. Those things add friction without adding signal, especially in the first year, and the profession has enough of them. The RFV adds one document and the requirement that later documents reference it.

The name is deliberate. It sits beside RFI, RFQ, and RFP as the fourth request, the one addressed to the buyer's own side, and it fills the slot the sequence always had open.

4. Why the definition has to be the buyer's

The natural objection is that vendors already do this work, often well and usually free. Vendor value engineering teams build the business case, capture a baseline, model the return, and run a realization review after go-live. If the seller is willing to write the definition and measure against it, why should the buyer duplicate the effort?

The answer is custody, and it has two parts.

The first part is the question being answered. A vendor's business case exists to help a purchase decision get made. It is built to answer whether the buyer will buy. It is competent at that. It has never been tested against the buyer's question, which is whether the value arrived. I know of no vendor value function that has published its projections against its outcomes across its deals, and I spent a career inside them. The reviews existed. Nobody added them up. The figures that do exist come from the buyer's side and are not encouraging: Gartner's 2025 CIO and Technology Executive Survey, fielded in 2024 across 3,186 CIOs, found 48 percent of digital initiatives meeting or exceeding their outcome targets.¹ That figure does not say the vendor cases were wrong. It says that whatever got those initiatives approved was never checked by anyone without a stake in the approval, and that an untested instrument should not be the only one in the room.

The second part is what the definition controls. Whoever writes the outcome decides what gets counted, and the same hand usually sets the baseline, which decides where the measurement starts, and picks the metric, which decides what improvement means. A vendor writing those three things in good faith will still write them in the vendor's favor, because the vendor's product is the lens through which the vendor sees the buyer's business. No misconduct is involved. The seat supplies the lens.

Since 2024 the seat has started to carry a meter. Enterprise vendors have begun charging per outcome: per resolution, per case closed, per result the vendor's system counts. In that arrangement the vendor defines the outcome, the vendor counts it, and the invoice follows the count. For as long as enterprises have bought technology, the buyer's inability to define and measure value in use was a gap in what the buyer knew. Under outcome

¹ Gartner, 2025 CIO and Technology Executive Survey, fielded May to June 2024, 3,186 CIOs and technology executives; Gartner press release, October 2024, reporting 48 percent of digital initiatives meeting or exceeding business outcome targets. Cited with full source in Michael Brey, "The Bracket," Value Lattice working paper, September 2026, section 8.

pricing the same gap shows up on the invoice.² The RFV is the buyer's own meter, written before the vendor's is installed.

None of this argues against vendor help. A vendor that can populate the buyer's RFV with evidence, propose a measurement method the buyer accepts, and commit to reporting against the buyer's metrics after go-live is a better vendor, and the RFV is how a buyer finds that out. The point is only that the page belongs to the party whose money is at risk.

5. Where it goes in the process

The RFV is a step, and a step has a place. Its place is between the requirement and the RFP.

Before. The requirement is settled enough to know what category of purchase is in view and which business function owns the result. The business owner, not procurement alone, writes the RFV. Procurement's role at this stage is to require that the page exists before the market is engaged, and to hold the line when the sponsor would rather skip it.

With the RFP. The RFV is attached to the RFP packet. Vendors are told that responses will be compared against it and asked to address each of its five components directly: here is how we would produce those outcomes, here is the evidence at comparable customers, here is what we would measure and how you would see it. The RFP continues to collect capability, credentials, security, and price as it always did.

In the evaluation. Responses are compared on the RFV's comparison criteria first, and on price and terms second. A response that cannot address the outcomes or supply the evidence is compared on that failure, whatever its feature list says.

In the contract. The contract references the RFV by name and attaches it. The outcomes, metrics, and baseline become part of what was agreed. Where the organization is ready for it, the vendor's post-go-live measurement and reporting commitment is written in. Where it is not ready, the reference alone still makes the page discoverable at renewal, which is more than most contracts allow today.

² Zendesk, outcome-based pricing announcement, 28 August 2024; Salesforce, pay-per-resolution announcement, 25 June 2026. Discussed in "The Bracket," section 9.

In governance. After award, the cadence that reviews the implementation reviews against the RFV's metrics, at the RFV's frequency, from the RFV's baseline. The document does not change without a written reason and a named person.

At renewal. The renewal decision opens the original RFV and asks one question of it: did the outcomes arrive, as measured by the metrics, from the baseline. That question cannot be asked of an RFP, a contract, or a vendor dashboard. It can only be asked of a page somebody on the buyer's side wrote at the start.

The step's placement is what makes it durable. An office can be defunded and a committee can be disbanded. A required document, attached to the RFP and referenced by the contract, survives reorganization because removing it would require somebody to remove a requirement, and nobody does that on purpose.

6. Who writes it, and what it forces

The RFV assumes the buyer has done the upstream work to decide what value means in the context of the decision. That work cannot be delegated to procurement alone, and it cannot be borrowed from the vendor. It requires the business owner who will live with the outcome to articulate what success looks like before the comparison begins.

This is where the RFV is more than an evaluation document. It is a forcing function. When the upstream work has not been done, the RFV exposes the gap rather than masking it. A business owner who cannot fill in the outcomes line does not have a purchase; they have an appetite. A metric line that cannot name a baseline reveals that the organization has not measured the thing it is about to spend money to improve. A conditions line that stays empty means nobody has thought about what the organization itself must change for the purchase to work.

Each of those discoveries is cheaper on page one than in month nine. Expect the page to take longer to write than anyone planned. The delay is the point. The RFP that follows is shorter and clearer, because the hard question was answered before it was drafted rather than discovered inside the responses.

The division of labor is simple. The business owner writes the outcomes, metrics, and conditions, because those belong to the function that will operate the result. Procurement writes the evidence requirements and the comparison criteria, because those belong to the

function that runs the market process. Finance reviews the baseline and the metric. The sponsor signs it. Four names on one page, before any vendor is called.

7. The seller's objection, and why adaptation pays

The first RFV an organization issues will draw an objection from at least one incumbent, and it will arrive at the CFO rather than at procurement. The process is unnecessarily complex. An outcome-driven evaluation will slow the decision. The vendor already has a proven methodology for value realization and would be glad to run it.

The objection deserves a straight answer, because it is partly right. The RFV does slow the front of the process, by the time it takes to write one page well. It does ask the vendor for something the vendor's sales process was not built to supply: evidence in the buyer's dimensions rather than the vendor's. And it does move the definition of success out of the vendor's hands.

What it does not do is add complexity. It removes a layer. Where the buyer used to receive five proposals with five different definitions of value and reconcile them in a comparison nobody could fully defend, the buyer now receives five responses to one definition. The comparison is shorter, and the decision is more defensible to the executives who were not in the room, because they can read the page.

Sellers who adapt to this early will do better than sellers who resist it, for a reason that has nothing to do with virtue. In a market where buyers hold their own definition, vendors compete on alignment and evidence, and a vendor whose outcomes at comparable customers are real has an advantage it cannot get any other way. A vendor with a thin record has every reason to prefer the old arrangement. The RFV sorts the two, which is exactly what a buyer wants an evaluation to do.

The composite case in *The Value Standard* records the pattern. In week six, the incumbent escalates to the CFO. The CFO defers to the CPO. The CPO holds the line. The vendor escalates once more, then complies.³ The practice survived because one person held a requirement, not because the architecture was elegant.

³ Michael Brey, *The Value Standard: How Buyers Evaluate, Sellers Differentiate, and Real Value Gets Realized* (2026), Chapter 20, "The First 90 Days." The case is a composite.

8. Attribution, the hard part

The strongest argument against outcome-based evaluation is attribution. Is the business system responsible for the five percent inventory reduction, or did the operations team get better at cycle counting? Did the procurement platform raise buyer productivity, or did a new head of procurement approve orders faster? If those questions cannot be answered, the objection runs, then measuring outcomes is theater and the honest measure is the one procurement already uses, which is price.

The objection is correct that attribution is hard. It is wrong that the difficulty argues for not measuring. It argues for the conditions component.

The conditions line of the RFV is where the buyer records, before the purchase, what else has to be true for the outcome to arrive: the adoption the business must achieve, the process the business must change, the data the business must clean, the headcount that must be in place. Writing those down at the start does two things. It makes the buyer's own contribution to the outcome visible and owned, so that a shortfall can be traced to the condition that failed rather than argued about afterward. And it gives the vendor a fair basis on which to be held to account, because the vendor is being measured against an outcome whose preconditions were stated, not against a number that assumed a perfect organization.

Attribution then becomes a governance question with a written record rather than a dispute with none. When the outcome misses, the cadence review asks which condition failed, and the answer is usually on the page. When the outcome arrives, the same page shows what the organization did to earn it. Neither party gets to argue from a number that was never written down.

Where the baseline is weak, the RFV says so. A baseline marked provisional, with the assumption named as a verification target, is more honest than a precise number nobody can defend, and the first review will identify exactly which assumptions failed because the baseline was provisional. That identification is the instrumentation work the organization never did, surfaced as evidence rather than deferred as a requirement.

9. Issuing the first one

The architecture in this paper can be adopted in stages, and the most reliable first stage is one of each rather than enterprise-wide.

Pick one category. The most productive first category is one where outcomes can be measured within a quarter and where vendor turnover is slow enough that a pattern can form. Avoid the category where the incumbent's political weight makes an honest evaluation hard in the first ninety days. That fight can be picked after the practice has produced its first piece of evidence.

Pick one decision. Upcoming, not retrospective. Material enough that the gap between projection and outcome will be visible in the next financial cycle. With a sponsor who will defend the discipline when the incumbent calls.

Pick one baseline. Use the data the organization has. If the source system cannot populate the metric precisely, mark the baseline provisional and the assumption as a verification target. Do not delay the page waiting for instrumentation that is not coming.

Write one page. Use Appendix A. Outcomes, metrics, conditions, evidence, comparison criteria. Four names at the bottom. Keep it to a page; the discipline is in the clarity.

Attach it, reference it, review against it, test it at renewal. The page does the rest of its work by being in the places the process already goes.

What not to attempt in the first ninety days. Cross-category rollout. A scoring matrix. Weighted vendor rankings. A parallel evaluation platform. Each adds friction without adding signal during the early window. They may be appropriate after two or three purchases have been governed through a full cycle. They are not appropriate before.

The practice starts with one decision done well, and it spreads because the page from that decision is sitting in the file when the next one comes up.

10. From a step to a standard

The RFV is the first act of governance, not the whole of it. What the buyer writes on the page becomes what the buyer commits to at decision, what the buyer governs against during execution, and what the buyer verifies at the end. The loop is define, commit,

govern, verify, and the RFV is where it begins. When any part of the loop runs alone, it works locally and fails structurally, which is the failure the standard exists to prevent.

That is also why the RFV is a standard rather than a framework. A framework prescribes how to evaluate, and frameworks compete: every quarter produces another five-test or seven-pillar cut from another respected source, none of them comparable to the others. A standard fixes the language any evaluation must use to be comparable across evaluators. The accounting profession learned this in its own time. Frameworks for reading financial statements proliferate; comparison across firms is possible only because a common standard anchors the words. The RFV's five components are that anchor for the definition of value. Any evaluation method can be laid over them. What the components do is make every purchase's definition of success legible in the same terms, so that the portfolio can be compared, the vendors ranked, and the investment defended at the level of the enterprise rather than one engagement at a time.

The evidence that the whole loop pays is beginning to accumulate. *The Value Standard* cites Gartner's research on value-based buying, which found that purchases where buyers defined measurable outcomes, defined measurement methods, used agile sourcing, and committed to continuous post-purchase assessment resulted in high-quality, low-regret deals 71 percent of the time, against 3 percent where none of those practices were followed.⁴ The first two of those four practices are the RFV. The page is where the loop starts, and it is the cheapest line in the budget.

Issued consistently, the RFV also becomes something the market can see before it bids: a standing statement of how this buyer defines, evaluates, and governs value, which vendors can read and decide whether they are able to operate inside. *The Value Standard* calls that published position the Preemptive Value Architecture and treats it, with the decision, governance, and audit instruments that follow the page, as the rest of the loop. This paper stops at the page, because the page is where every organization has to start.

⁴ *The Value Standard*, Chapter 15, citing Gartner research on value-based buying. The figures are reported as cited in the book; readers wanting the underlying survey should consult Gartner's technology buying behavior research directly. The five-component definition of the RFV in this paper follows Appendix A of the book, "Request for Value (RFV)."

Appendix A. The Request for Value

The form below is the instrument. Copy it, adapt it, and issue it without permission or attribution. Fill every line or write why it cannot be filled. At this paper's narrow page width it runs onto a second page; at ordinary letter margins it is one.

REQUEST FOR VALUE

Purchase	
Category	
Business owner / Procurement lead / Sponsor	
Date issued	

1. Outcomes. The business results this purchase must produce, in our terms, each with a named owner.

Outcome	Owner
1.	
2.	
3.	

2. Metrics. Indicator, baseline and date, target, frequency, review trigger. Mark provisional baselines.

Outcome	Indicator and source	Baseline (as of)	Target	Frequency	Review trigger
1.					
2.					
3.					

3. Conditions. What must hold, on our side and the supplier's, for the outcomes to be possible.

Condition	Owner
1.	
2.	
3.	

4. Evidence. What a supplier must provide for a claim to count.

Requirement	Required
Realized outcomes at comparable customers, with the distribution of results, not the best case	yes / no
Conditions under which that evidence was produced, and where the solution underperformed and why	yes / no
Post-go-live measurement the supplier commits to, the data we will have access to, and the reporting cadence	yes / no
Other	

5. Comparison criteria. How responses will be weighed on the four components above, written before responses arrive.

Attached to the RFP, referenced by the contract, reviewed at the stated frequency, tested at renewal. Changes only with a written reason and a named person.

Business owner	Procurement	Finance	Sponsor

About the Author

Michael Brey worked in business value engineering for enterprise software vendors, building and defending the business cases this paper argues buyers should own, before writing *The Value Standard*, which sets out the buyer-held instruments of which the RFV is the first. He publishes at valuelattice.com. This is a working paper; corrections to any sourced claim are welcome and will be recorded.