WORLDVUE®

The Vendor Evaluation Playbook:

How to Choose a Technology Partner When Every Demo Looks the Same

SECTION 01 — WHY THE STANDARD EVALUATION PROCESS FAILS

The Demo Is Designed to Win, Not to Inform

Before we build the framework, it is worth understanding exactly why the standard hospitality tech evaluation process consistently produces poor outcomes, not because operators are unsophisticated, but because the evaluation format is structurally biased toward the vendor.

The Five Structural Biases of the Standard AI Demo

1) **Controlled environment performance.**
Every demo runs on the vendor's infrastructure, with their data, at their chosen load level. The platform that responds in under a second in the demo may take four seconds in your property's network environment. Issue #18 documented this gap. It is not a bug. It is a feature of demo design.

2) **Curated case studies.**
Vendors present their best implementations. You are not seeing the 40% of deployments that underperformed, the integrations that took nine months instead of three, or the properties that churned after year one. You are seeing the highlight reel.

3) **Feature depth over integration breadth.**
Demos showcase what the platform does in isolation. They rarely show how it behaves when connected to your specific PMS, your CRM, your housekeeping platform, and your loyalty system simultaneously, which is the only environment that matters.

4) **ROI projections built on assumptions you cannot verify.**
The 340% ROI projection in the leave-behind was built on assumptions about your current baseline, your staff adoption rate, your network performance, and your competitive set. Change any of those assumptions and the projection changes fundamentally. You are rarely shown the model, only the output.

5) **The wrong people in the room.**
Demos are attended by GMs, owners, consultants and revenue managers, the people with budget authority. The people who will live with the decision daily, front desk supervisors, housekeeping leads, the IT manager, the AI Ambassadors from Issue #16, are rarely present. Vendors know who controls the contract.

None of this is nefarious. It is rational vendor behavior. The problem is that buyers have adapted to the demo as the evaluation, rather than treating it as what it actually is: the opening of a sales process that requires a rigorous evaluation before any decision is made.

> “We watched three vendors demo the same platform category. All three scored over 90 on our internal rating sheet. Then we asked each of them to connect to our test PMS environment. One of them never got it working in the two weeks we gave them. That told us everything we needed to know.”
>
> - *Director of Technology, 380-room full-service hotel, Southeast*

SECTION 02 — BEFORE THE DEMO: WHAT TO ESTABLISH FIRST

The Evaluation Begins Before the Vendor Enters the Room

The properties that make the best tech vendor decisions start their evaluation before they invite any vendor to present. They do three things that most operators skip entirely.

Step 1: Define Your Outcomes Before You Define Your Requirements

Most RFPs and vendor briefings are built around feature requirements: "We need a platform that does dynamic pricing, preference management, and automated pre-arrival communications." The problem with feature-first evaluation is that nearly every vendor in a given category can check every feature box. Features are table stakes. Outcomes are differentiators.

Before you contact a single vendor, your evaluation team should answer these four questions in writing:

The Four Outcome Questions to Answer Before Any Vendor Contact

1) What specific guest experience problem are we trying to solve — and what does success look like in measurable terms? (Not "improve personalization." Specifically:"Increase our AI-Enabled Proactive Service Rate from 12% to 35% within 12 months.")

2) What is our current Culture ROI Multiplier baseline, and what multiplier do we need this platform to help us reach? (If you have not run the Issue #17 diagnostic, do that first.)

3) What are the three staff behaviors we need to change for this platform to succeed — and are we prepared to invest in the change management to change them? (A platform that requires staff to check a new interface 40 times per shift will fail if staff currently check it zero times.)

4) What does our network infrastructure assessment show about our readiness to support this platform at peak load? (If you have not run the Issue #18 diagnostic, do that before you speak to any vendor.)

Step 2: Score Your Own Readiness Before You Score the Vendors

The 3.7× properties shared a discipline that most operators find uncomfortable: they evaluated themselves before they evaluated the vendors. They asked: are we ready to deploy what we are about to buy?

Readiness Dimension	Ready to Deploy	Needs Preparation First
Network infrastructure	Completed Issue #18 diagnostic; 7+ Green answers	Fewer than 7 Green; failure modes unaddressed
Staff AI sentiment	STES above 7.0; AI Ambassadors identified	STES unknown or below 6.0; no ambassadors in place
Integration landscape	PMS, CRM, and RMS documented and API-accessible	Point-to-point integrations; no integration map
Change management capacity	Leadership committed to Superpower Briefings and 90-day plan	No change management budget or plan identified
Data quality	Guest profile completeness above 60%; clean historical data	Fragmented guest data; no unified profile layer
Measurement framework	KPIs defined; baseline metrics established	No baseline; will measure “when we see results”

If you have three or more “Needs Preparation First” answers, you are not yet ready to evaluate AI vendors. You are ready to address your preparation gaps. The vendors will wait. Your guests cannot afford for you to deploy a platform your organization is not ready to use.

Step 3: Build Your Evaluation Team Deliberately

The people in the room when you evaluate AI vendors should reflect the people who will live with the decision. For most properties, the right evaluation team includes:

- The General Manager or Department Head — accountable for outcomes and culture
- The Director of Technology or IT Manager — accountable for infrastructure and integration
- One front-line representative from the primary use case department, the person whose daily work changes most
- The Revenue Manager or Director of Sales, if the platform touches pricing or distribution
- One designated skeptic, someone whose job in the evaluation is to find what is wrong, not what is right

The designated skeptic is the most important person on the list and the one most often absent. Give them explicit permission to ask the uncomfortable questions. Vendors who cannot handle a skeptic in the evaluation room cannot handle the inevitable challenges of implementation.

SECTION 03 — THE MASTER EVALUATION FRAMEWORK

The 40-Question Vendor Assessment: What to Ask, What to Listen For, and What to Watch Out For

The following framework is organized into six evaluation categories. Each category contains questions designed to reveal something the demo will not show you. The red flags are specific they are the answers that, in our observation of 200-plus property deployments, consistently predicted underperformance, failed implementations, or vendor relationships that did not survive year two.

We recommend running this assessment across every vendor under consideration, using the same questions in the same order, so that your scores are genuinely comparable.

CATEGORY 1: INFRASTRUCTURE COMPATIBILITY AND NETWORK READINESS

The Question	What a Strong Answer Looks Like	The Red Flag
Can you assess our specific network environment and tell us before we sign whether our current infrastructure can support your platform at the performance levels in your demo?	A specific offer to conduct a pre-sale infrastructure assessment. Willingness to share minimum latency, bandwidth, and segmentation requirements in writing.	"Most properties don't have issues." or "Our platform is cloud-based so infrastructure doesn't matter." Any answer that avoids the specifics of your environment.
What are the specific bandwidth and latency requirements for your platform at our projected device load and occupancy level?	A written specification: minimum WAN bandwidth, acceptable latency thresholds, QoS requirements, IoT segmentation needs.	"It's pretty light on bandwidth" or a generic spec sheet not tailored to your property profile.
Which of our existing systems does your platform need to integrate with, and can you demonstrate a live connection to a reference property's PMS in the same family as ours?	A live integration demonstration on a reference environment, not a recorded demo. Specific documentation of PMS compatibility.	"We integrate with all major PMS platforms" without being able to demonstrate a live connection to yours specifically.
What happens to your platform's core functionality when our WAN connection drops for 90 seconds during peak occupancy?	A specific explanation of offline behavior, data queuing, graceful degradation, and recovery sequence.	"It's cloud-based so connectivity is assumed." No answer about degraded-mode operation.

CATEGORY 2: IMPLEMENTATION REALITY

The Question	What a Strong Answer Looks Like	The Red Flag
Can you show us the implementation timeline for your last five comparable property deployments not your standard timeline, the actual timelines?	Willingness to share actual (not projected) timelines. Honest acknowledgment of what caused delays and what was done to address them.	Only offering the standard implementation timeline document. Unwillingness to discuss specific past deployments by name or in anonymized detail.
Who is our named implementation manager, and what is their current portfolio of active implementations?	A named individual with verifiable experience. An honest answer about how many implementations they are currently managing.	"You'll be assigned an implementation manager after signing." Inability to name the person who will own your deployment.
What percentage of your implementations in the last 12 months went live within the contracted timeline?	A specific percentage with an honest explanation of what caused the exceptions.	A deflection: "most of our implementations go smoothly" without a specific number.
What does your go-live definition actually include and what remains the property's responsibility after go-live?	A precise definition of go-live. A clear delineation of vendor vs. property responsibilities post-launch.	A vague go-live definition that could be met by technical completion while leaving the property to figure out adoption independently.
Can we speak with a property that went through implementation in the last six months including one that had a difficult implementation?	Immediate willingness to provide references, including one that experienced challenges.	Only offering references from flagship success stories. Reluctance to provide a reference from a challenging deployment.

CATEGORY 3: STAFF ADOPTION AND CHANGE MANAGEMENT SUPPORT

The Question	What a Strong Answer Looks Like	The Red Flag
What is your average Staff Technology Enthusiasm Score at properties 90 days after go-live, versus at properties 12 months after go-live?	Familiarity with STES as a metric (or an equivalent measure). Data on adoption trajectory over time.	Unfamiliarity with the metric. An answer focused entirely on utilization rates or feature activation rather than staff sentiment.
What change management resources do you provide, and are they included in the contract or billed separately?	Specific change management deliverables: training materials, Superpower Briefing templates, AI Ambassador program support, 30/60/90-day adoption plans.	"We provide training documentation and a video library." Change management treated as documentation, not process.
What is your staff training approach for a property where 40% of front desk staff have never used AI-assisted tooling before?	A differentiated approach that acknowledges the starting point. Reference to role-specific training sequences and peer mentorship models.	A one-size-fits-all training program. No differentiation for staff with different technology comfort levels.

When a property's STES drops or adoption stalls at month three, what is your intervention process?	A specific, named intervention protocol. Evidence of having successfully re-launched adoption at a named property.	"We would work with the property to understand the issue." No defined process, no precedent.

CATEGORY 4: DATA, PRIVACY, AND INTEGRATION ARCHITECTURE

The Question	What a Strong Answer Looks Like	The Red Flag
Where does our guest data live, who owns it, and what happens to it if we terminate the contract?	Clear contractual language: property owns all guest data, full data export within 30 days of termination, no vendor use of property data for model training without explicit consent.	Vague language about "data security" without addressing ownership, portability, or post-contract access.
How does your platform handle a guest who requests deletion of their data under applicable privacy regulations?	A specific, documented data deletion workflow. Confirmation that deletion propagates across all systems and partner integrations.	"We comply with all applicable regulations" without describing the operational process.
What is your integration architecture point-to-point APIs or a hub-and-spoke data layer and how does that affect data latency when all integrated systems are active?	A clear explanation of integration architecture. Specific data latency figures for multi-system environments.	An inability to explain integration architecture at a technical level. A claim that integration latency is negligible without supporting data.
What is your SOC 2 Type II certification status, and can you share your most recent penetration test results?	Current SOC 2 Type II certification. Willingness to share penetration test summary with NDA. Evidence of proactive security posture.	SOC 2 Type I only, or certification in process. Reluctance to discuss security posture details.

CATEGORY 5: COMMERCIAL TERMS AND PARTNERSHIP MODEL

The Question	What a Strong Answer Looks Like	The Red Flag
What are all the costs we will pay in year one, year two, and year three — including implementation, integration, per-room fees, data storage, support tiers, and any fees triggered by usage thresholds?	A complete, itemized multi-year cost model with no material omissions. Willingness to include all potential fees in a schedule.	A first-year price with vague language about subsequent years. Fees that emerge after
What are the exit terms if the platform does not perform against the ROI projections in your proposal?	Performance-based exit provisions. Willingness to tie contract terms to agreed outcome metrics.	Standard annual or multi-year commitments with no performance provisions. Exit penalties that make switching prohibitively expensive.

What is your price escalation policy for renewal years, and is there a cap?	A written escalation cap, typically CPI-indexed or fixed percentage.	"Pricing is reviewed annually" without a stated cap. Escalation language buried in the contract.
How many of your hospitality customers have renewed for a second contract term, and what is your net revenue retention rate?	A specific renewal rate and NRR figure. Willingness to discuss what drives renewal decisions.	Inability or unwillingness to share retention metrics. A pivot to new customer acquisition numbers instead.

CATEGORY 6: LONG-TERM PARTNERSHIP AND ROADMAP ALIGNMENT

The Question	What a Strong Answer Looks Like	The Red Flag
What does your product roadmap look like for the next 18 months, and how do customer properties participate in roadmap decisions?	A specific roadmap with named features and target timelines. A described customer advisory council or equivalent input mechanism.	"We are constantly innovating" without specifics. No customer input mechanism described.
How does your platform perform at properties achieving a 3.7× Culture ROI Multiplier compared to industry average properties and what does your data show about what drives that difference?	Familiarity with the Culture ROI Multiplier concept. Data-backed answer about performance drivers across their customer base.	Unfamiliarity with the framework. An answer that attributes performance differences primarily to the software rather than to the organizational factors that drive adoption.
Who is our named account manager after go-live, and what does proactive account management look like at your company?	A named individual. A described cadence of proactive outreach, performance reviews, and escalation paths.	"You'll have access to our support portal and customer success team." No named individual, no described cadence.
What is the most common reason properties in your portfolio underperform against their projected ROI and what do you do about it?	An honest, specific answer that demonstrates self-awareness. A described intervention process with evidence of success.	Deflection: "performance depends on the property" without accepting any vendor accountability for outcomes.

SECTION 04 — THE SCORING SYSTEM

Turning Qualitative Answers into a Defensible Decision

The evaluation framework above generates qualitative intelligence. To make a defensible vendor selection decision one you can present to your ownership group and revisit honestly in 18 months, you need a scoring system that converts that intelligence into comparable data.

The WorldVue Vendor Evaluation Scorecard

Score each vendor on the following six categories using a 1–5 scale. The category weights reflect their relative importance to long-term platform performance, based on the patterns observed across the WorldVue property dataset.

Evaluation Category	Weight	Score (1–5)	Weighted Score
Infrastructure Compatibility & Network Readiness	25%	___ / 5	___ × 0.25
Implementation Reality	20%	___ / 5	___ × 0.20
Staff Adoption & Change Management Support	20%	___ / 5	___ × 0.20
Data, Privacy & Integration Architecture	15%	___ / 5	___ × 0.15
Commercial Terms & Partnership Model	10%	___ / 5	___ × 0.10
Long-Term Partnership & Roadmap Alignment	10%	___ / 5	___ × 0.10
TOTAL WEIGHTED SCORE	**100%**		

Interpreting Your Scores

Weighted Score	Interpretation	Recommended Action
4.2 – 5.0	Strong partner candidate	Proceed to contract negotiation with performance provisions
3.5 – 4.1	Qualified with conditions	Identify specific gaps; negotiate remediation commitments into contract before signing
2.8 – 3.4	Significant concerns	Request a second evaluation round addressing specific gaps; consider alternatives
Below 2.8	Not recommended	Do not proceed. A platform that scores below 2.8 on this framework will underperform against your investment.

The Automatic Disqualifiers — Regardless of Total Score

- Cannot demonstrate a live integration with your PMS family before contract signature
- Unable or unwilling to provide references from a challenging implementation
- No clear data ownership and portability provisions in the contract
- No named implementation manager before signing
- No performance-linked exit provisions in a contract longer than 12 months
- Unable to answer the network infrastructure and latency requirements question with specifics

SECTION 05 — AFTER THE EVALUATION: NEGOTIATING THE RIGHT CONTRACT

The Contract Is Where Partnership Begins, or Ends

The evaluation process tells you who to negotiate with. The contract negotiation tells you whether that vendor is prepared to be a partner. The 3.7× properties consistently negotiated contracts that reflected shared accountability for outcomes. Average and laggard properties signed standard vendor contracts and hoped for the best.

The Five Contract Provisions That Separate Partners from Vendors

1) **Performance benchmarks with defined consequences.**
Your contract should specify the AI-Enabled Proactive Service Rate, STES improvement targets, and integration latency thresholds your vendor commits to helping you achieve and what happens if those benchmarks are not met within 12 months.

2) **Named resources, not team commitments.**
Your implementation manager, your account manager, and your technical integration lead should be named individuals in the contract, not “a member of our team.” When those individuals change, you should have the right to approve their successors.

3) **Data ownership in plain language**
The contract should state explicitly: the property owns all guest data; the vendor may not use property data for model training, product development, or third-party sharing without written consent; data is exportable in full within 30 days of contract termination at no additional cost.

4) **A staged payment structure tied to go-live milestones.**
Do not pay the full implementation fee before go-live. Structure payments against verified milestones: integration certification, staff training completion, and 30-day live performance data.

5) **An infrastructure assessment provision.**
If the vendor has not conducted a pre-contract infrastructure assessment, include a provision that implementation does not begin until a mutual infrastructure readiness sign-off is completed. This protects both parties from deploying into an environment that cannot support the platform.

SECTION 06 — THE EVALUATION TIMETABLE

A 45-Day Process That Protects Your Investment

Operators who rush AI vendor selection consistently regret it. The following timetable has been refined through observation of what the 3.7× properties did differently. It is not slow. It is deliberate.

Phase	Timeline	Key Activities	Deliverable
Internal Preparation	Days 1–10	Complete Issue #17 Culture ROI baseline. Complete Issue #18 network readiness diagnostic. Define outcome goals. Assemble evaluation team. Identify your designated skeptic.	Written outcome statement. Network readiness report. Evaluation team charter.
Vendor Briefing	Days 11–15	Send evaluation framework to shortlisted vendors (3–4 maximum). Request written responses to infrastructure, integration, and commercial categories before the demo.	Written vendor pre-responses on file before any demo occurs.
Demo + Live Testing	Days 16–25	Standard demo from each vendor. Followed immediately by: live integration test with your PMS environment. Infrastructure requirements review against your network assessment.	Demo scorecard. Live integration result (pass/fail). Infrastructure compatibility assessment.
Reference Calls	Days 26–32	Two reference calls per vendor: one flagship success, one challenging implementation. Ask the same questions to each reference.	Reference call notes. Comparison of vendor claims vs. reference reality.
Scoring and Selection	Days 33–38	Complete weighted scorecard for each vendor. Identify automatic disqualifiers. Present recommendation to ownership group with documented rationale.	Completed scorecards. Ownership presentation. Selected vendor.
Contract Negotiation	Days 39–45	Negotiate the five contract provisions. Do not sign without named resources, data ownership language, and performance provisions.	Executed contract with all five provisions included.

SECTION 07 — A NOTE ON THE VENDORS WHO WELCOME THIS PROCESS

The Framework Is Also a Filter for the Right Vendors

We want to close with an observation that has become one of the most reliable signals in our work with 200-plus properties: the best AI vendors are not threatened by this evaluation framework. They welcome it.

A vendor who responds to your infrastructure compatibility questions with specific, documented answers is a vendor who has already solved the problem with other properties. A vendor who provides a challenging implementation reference without hesitation is a vendor who learned from that experience and can tell you what they changed. A vendor who builds performance provisions into their contract is a vendor who believes in their platform enough to be accountable to it.

Conversely, a vendor who deflects the infrastructure questions, who offers only glowing references, who resists performance provisions, and who presents a contract written entirely to protect their interests while deferring all risk to the property, that vendor is telling you something important. Listen.

The evaluation framework is not just a tool for choosing the right vendor. It is a tool for eliminating the wrong ones quickly, before they cost you 18 months of underperformance and the cultural damage that follows a failed technology deployment.

Your guests experience every technology decision you make. The evaluation process is where you decide what they will experience, and for how long.

> The vendors who pushed back hardest on our evaluation process were the first ones we eliminated. The vendor we chose sent back our 40 questions with detailed written answers before we even scheduled the demo. That told us everything.
>
> \- *GM, 310-room independent hotel, Pacific Northwest*