FLM PROVIDER EVALUATION GUIDE

10 Questions to Ask Before Signing the Contract

Including: evaluation scorecard, red flags checklist,
and how Hype Telecom answers each question with auditable data

Hype Telecom LLC

hypetelecom.net | hypetelecom@hypetelecom.net | +1 (321) 947-2184

2281 Lee Rd, Suite 201, Winter Park, FL 32789

CONTENTS

SECTION 1 & 2

Why This Guide Exists

Choosing an FLM provider is one of the most consequential decisions a data center operator makes — and one of the most poorly evaluated.

Most procurement processes focus on pricing, SLA terms, and geographic coverage. These matter. But they're surface metrics. The questions that actually predict whether an FLM provider will deliver or disappoint are different — and most operators don't ask them until after the first SLA breach.

This guide compiles the 10 questions that, in our experience operating across the Americas for hyperscale clients, consistently separate reliable FLM partners from providers who look good on paper but struggle in execution.

Each question includes context on why it matters, what a strong answer looks like, what a weak answer reveals, and the specific red flags that should trigger deeper due diligence.

How to Use This Guide

→ During procurement: use the 10 questions as a structured evaluation framework alongside your standard RFP.

→ During contract review: verify that the provider's answers from procurement are reflected in contractual commitments, not just verbal assurances.

→ During annual review: re-evaluate your current provider against these criteria to identify performance drift before it becomes an SLA risk.

→ The printable scorecard (page 13) allows you to compare up to two providers side by side on a standardized 50-point scale.

CATEGORY 1 Performance Data

Performance data is the foundation. Without verified metrics, every other claim is unsubstantiated. These two questions establish whether the provider can back their promises with evidence.

Q1 What is your average incident resolution time?

Why it matters: Resolution time is the single most important operational metric in FLM. It directly determines SLA compliance, client satisfaction, and the financial impact of every incident. Average resolution time as a single number can be misleading — it must be broken down by priority tier, by site, and over a meaningful time period.

What a strong answer looks like: The provider can produce documented resolution time data by priority tier (P1, P2, P3, P4) across all active contracts for the past 12 months. The data is auditable — not a marketing figure, but operational records that can be independently verified. Strong providers will proactively share this data because it demonstrates their capability.

What a weak answer reveals: Vague responses like 'we resolve most incidents quickly' or 'our SLA is 4 hours' without supporting data indicate either poor performance tracking or performance that doesn't withstand scrutiny. A provider that quotes their SLA target instead of actual performance is telling you they don't track — or don't want to share — the real numbers.

> 💡 **Insider tip:** *Ask for the distribution, not just the average. A provider with a 3.5-hour average might have 80% of incidents resolved in 2 hours and 20% taking 8+ hours. The average looks good. The 20% is where SLA breaches live.*

- 🔴 The provider cannot produce resolution time data broken down by priority tier.
- 🔴 Data is provided as a marketing summary, not as operational records.
- 🔴 The provider quotes their SLA target rather than their actual measured performance.

Q2 What is your SLA breach rate over the last 12 months?

Why it matters: SLA breach rate reveals the provider's consistency under pressure. A provider with excellent average resolution time but frequent breaches on high-priority incidents is unreliable precisely when reliability matters most. Breach rate, broken down by priority tier, shows whether the provider's system works in the moments that count.

What a strong answer looks like: Zero — or a specific, low number with documented root cause analysis for each breach. The best providers treat every breach as a process improvement event, not an acceptable failure rate. They can show you the breach, explain what caused it, and demonstrate the corrective action taken to prevent recurrence.

What a weak answer reveals: Responses like 'we're within acceptable parameters' or 'our breach rate is very low' without specific numbers indicate either frequent breaches being normalized or inadequate tracking. If a provider can't tell you their exact breach count for the past 12 months, they aren't managing their SLA performance — they're hoping it works out.

> 💡 **Insider tip:** *Ask what happens contractually when a breach occurs. Do penalties auto-accrue? Is there a cure period? How are breaches communicated to you? The financial and reporting framework around breaches reveals how seriously the provider takes them.*

- 🔴 The provider cannot provide an exact breach count.
- 🔴 Breaches are described as 'within acceptable parameters' without specifying what's acceptable.
- 🔴 No documented root cause analysis exists for past breaches.

CATEGORY 2

The provider's delivery model — how their teams are structured and where they're located — determines whether they can realistically meet aggressive SLAs or whether travel time consumes the resolution window before anyone arrives.

Q3 Are your technicians dedicated to my environment or shared across clients?

Why it matters: The provider assigns named, dedicated technicians to your environment with documented training on your specific hardware and configurations. Backup technicians are cross-trained and familiar with your environment — not cold replacements. The provider can tell you who is assigned to your account by name, what certifications they hold, and how long they've been supporting your infrastructure.

What a strong answer looks like: A dedicated technician who knows your specific hardware, configurations, network topology, and facility layout resolves incidents fundamentally faster than a shared generalist arriving cold. The difference isn't skill level — it's context. A technician who has resolved 50 incidents in your environment recognizes patterns, knows shortcuts, and carries institutional knowledge that a rotating pool cannot develop.

What a weak answer reveals: Responses like 'we assign the next available qualified technician' or 'our team is multi-client by design' indicate a dispatch model that prioritizes provider efficiency over client-specific expertise. This model works for routine tasks. It fails for complex incidents where environment-specific knowledge determines resolution speed.

> 💡 **Insider tip:** *Ask about technician turnover rate on your account. High turnover — even with dedicated assignment — means the institutional knowledge advantage is constantly being reset. The best providers have low turnover because they invest in their teams.*

- 🔴 The provider cannot name the technicians assigned to your account.
- 🔴 'Next available' dispatch model with no dedicated assignment.
- 🔴 No documented training on your specific hardware and environment.

Q4 Where are your teams physically located relative to my sites?

Why it matters: In a sub-4-hour SLA environment, geography is destiny. If the nearest qualified technician is 90 minutes away — or a flight away — more than a third of the SLA window is consumed by travel before anyone touches the equipment. Pre-positioned local teams eliminate this structural time penalty.

What a strong answer looks like: The provider can map their team locations to your sites with specific distances and estimated mobilization times. For each of your facilities, there is a named technician within 30 minutes of travel time. The provider has permanent local presence in your markets — not a promise to 'deploy resources as needed.'

What a weak answer reveals: Responses like 'we have national coverage' or 'we can mobilize to any location' indicate a centralized dispatch model where teams travel to your sites rather than living near them. National coverage with centralized dispatch is a geographic coverage claim, not a response time capability.

> 💡 **Insider tip:** *Ask for mobilization time data — not estimates, but actual historical data showing time-to-arrival for the past 12 months at each of your sites. The gap between estimated and actual mobilization time often reveals the real delivery model.*

- 🔴 'National coverage' without specific team locations mapped to your sites.
- 🔴 Mobilization time estimates based on best-case scenarios, not historical data.
- 🔴 The provider must fly technicians to your secondary sites.

CATEGORY 3 Operational Continuity

24/7 coverage is easy to promise and difficult to deliver without gaps. These questions expose the operational discipline — or lack thereof — behind the '24/7/365' claim that every provider makes.

Q5 What happens during shift transitions — is there a coverage gap?

Why it matters: The provider has a documented shift handoff protocol that includes complete status transfer: all open tickets, pending analyses, scheduled maintenance windows, emerging patterns, and any priority issues the incoming team must monitor. The handoff is documented — not verbal — and takes a specific, defined amount of time. There is overlap between shifts, not a gap.

What a strong answer looks like: Shift transitions are the most vulnerable moments in 24/7 operations. The outgoing team is wrapping up. The incoming team is settling in. If the handoff isn't structured and documented, tickets fall through cracks, emerging patterns go unnoticed, and the client's infrastructure experiences a window of reduced oversight that doesn't show up in any SLA report.

What a weak answer reveals: Responses like 'our shifts overlap by 15 minutes' without describing what happens during those 15 minutes indicate an informal handoff that depends on individual discipline rather than systematic process.

- 🔴 No documented shift handoff protocol.
- 🔴 Verbal handoffs without written status transfer.
- 🔴 No shift overlap — one team leaves before the next arrives.

Q6 Do you perform remote pre-analysis before dispatching a technician?

Why it matters: The provider describes a structured pre-analysis process: when a ticket arrives, their coordination team reviews scope and equipment history, performs preliminary assessment, and generates a briefing for the field technician that includes preliminary diagnosis, recommended tools and parts, and facility-specific access instructions. The technician receives this briefing before leaving for the site.

What a strong answer looks like: Pre-dispatch analysis is the single highest-leverage process improvement in FLM. When the coordination team reviews the scope, checks equipment history, assesses likely root causes, and briefs the field technician before dispatch — the technician arrives prepared. They know what they're likely to find. They carry the right parts. Resolution starts immediately rather than after 30–60 minutes of on-site discovery.

What a weak answer reveals: 'We dispatch the nearest available technician immediately' sounds fast. It's actually slow. It means the technician arrives cold — without diagnosis, without the right parts, without context. The apparent speed of immediate dispatch is an illusion when it results in longer on-site resolution time.

> 💡 **Insider tip:** *Ask whether the coordination team's pre-analysis is standardized and documented — not ad-hoc. At Hype Telecom, our coordination team performs structured pre-analysis before dispatch, with the field technician briefed prior to departure.*

- 🔴 No pre-analysis process — 'we dispatch immediately.'
- 🔴 Pre-analysis exists but isn't standardized or time-bound.
- 🔴 The field technician doesn't receive a briefing before arriving on site.

CATEGORY 4

Documentation quality predicts whether the relationship improves over time or stagnates. Scalability determines whether the provider can grow with you or become a bottleneck when you need them most.

Q7 What documentation do you provide per resolved incident?

Why it matters: Each resolved incident produces a structured report that includes: timestamp of ticket receipt, pre-analysis findings, technician assignment and arrival time, root cause identification, resolution steps taken, parts replaced (with serial numbers), verification testing results, total resolution time, and any recommendations for preventive action. Photo documentation at key milestones is standard, not optional.

What a strong answer looks like: Incident documentation serves three purposes: it provides the client with a complete record of what happened and why, it enables pattern recognition across incidents, and it creates accountability for the quality and thoroughness of the resolution.

What a weak answer reveals: 'We provide a ticket closure summary' or 'all incidents are logged in our system' without describing the documentation depth indicates minimal record-keeping. If you can't reconstruct exactly what happened during an incident from the documentation alone, the documentation is insufficient.

- 🔴 Incident reports are one-line ticket closures, not structured analysis.
- 🔴 No photo documentation standard.
- 🔴 Documentation doesn't include root cause analysis or preventive recommendations.

Q8 Can you scale to additional sites without degrading performance?

Why it matters: The provider describes a specific growth model: how they add capacity in new markets (hiring locally vs. relocating), what the timeline is for standing up a new market, whether they have existing local presence in your expansion markets, and how quality standards are maintained during scale-up. The provider can reference specific examples of successfully scaling with existing clients.

What a strong answer looks like: Most FLM relationships begin at a manageable scale. The real test comes when you need to expand. Providers who scale by stretching existing resources thinner deliver degraded performance at the expanded scope. Providers who scale by adding local capacity in new markets maintain consistency.

What a weak answer reveals: 'Yes, we can scale anywhere' without describing the mechanism is a sales answer, not an operational one. If the provider's growth model is to stretch their existing team across more sites, every new site degrades performance at every other site.

> 💡 **Insider tip:** *The best predictor of scalability is existing local presence. A provider with pre-positioned teams in your expansion markets can add your sites to an existing local operation. A provider who must build from scratch in a new market faces a 3–6 month ramp with elevated risk.*

- 🔴 'We can scale anywhere' without a specific growth mechanism.
- 🔴 Scaling requires relocating existing technicians, thinning coverage at current sites.
- 🔴 No examples of successfully scaling with existing clients.

CATEGORY 5 Training & Escalation

The final category determines whether the provider's team can handle the complexity of your environment or whether they'll struggle with anything beyond routine tasks.

Q9 What training do your technicians receive on my specific hardware?

Why it matters: The provider describes a specific onboarding training program for new clients: a defined period where technicians learn your hardware configurations, your facility layout, your monitoring tools, and your specific procedures. This training is documented and verified. Ongoing training updates when hardware configurations change. Technicians hold OSHA 10/30, CompTIA A+/Network+/Server+, plus relevant OEM certifications (Corning, Ciena, etc.).

What a strong answer looks like: Hyperscale environments use specific hardware configurations that differ significantly from one client to another. A technician certified on generic server hardware may not know the specific firmware quirks, configuration parameters, or failure patterns of your exact equipment models. Environment-specific training is the difference between a technician who resolves and one who escalates.

What a weak answer reveals: 'Our technicians are fully certified' without specifying training on your specific environment indicates general competence without specific expertise. Certifications prove foundational knowledge. Environment-specific training proves operational readiness for your infrastructure.

- 🔴 No documented client-specific training program.
- 🔴 General certifications cited without environment-specific training.
- 🔴 No ongoing training updates when hardware configurations change.

Q10 What is your escalation process when a Level 1 technician can't resolve?

Why it matters: The provider has documented, automatic escalation triggers based on time-in-SLA-window, not individual judgment. Time-based thresholds are defined for each priority tier. Escalation fires regardless of the technician's assessment — because in the middle of an incident, the person closest to the problem is the least likely to recognize when it's time to escalate.

What a strong answer looks like: Escalation is where most FLM operations fail. When a Level 1 technician encounters an issue beyond their capability, the speed and structure of escalation determines whether the incident gets resolved within the SLA window or breaches. Manual, ad-hoc escalation depends on individual judgment in high-pressure situations — where optimism bias ('I can figure this out') leads to delayed escalation and SLA breaches.

What a weak answer reveals: 'Our technicians know when to escalate' is an answer that guarantees delayed escalation. Human judgment under pressure defaults to 'let me try one more thing.' Systematic, time-based triggers override this bias and ensure escalation happens before the SLA window closes — not after.

> 💡 **Insider tip:** *Ask to see the escalation matrix — the actual document. It should show, for each priority tier, the specific time thresholds that trigger each escalation level, who gets notified, and what actions are taken. If it doesn't exist as a document, it doesn't exist as a process.*

- 🔴 Escalation depends on individual technician judgment, not automatic triggers.
- 🔴 No documented escalation matrix with time-based thresholds.
- 🔴 OEM vendor relationships are not pre-established — they're ad-hoc during incidents.

SECTION 8 — RED FLAGS SUMMARY

12 Warning Signs During Evaluation

If you encounter any of these, it warrants deeper investigation — or disqualification, depending on your risk tolerance.

1. **Vague performance data.** The provider cannot produce auditable resolution time or SLA breach data by priority tier for the past 12 months.
2. **Marketing metrics instead of operational records.** Performance claims are sourced from sales materials, not from documented operational data that can be independently verified.
3. **'Next available' dispatch model.** Technicians are not dedicated to your environment — they're dispatched from a rotating pool based on availability, not expertise.
4. **Centralized dispatch with long mobilization times.** The provider's nearest qualified technician is more than 30 minutes from your critical sites. Travel time consumes the SLA window.
5. **No shift handoff protocol.** Shift transitions are informal, verbal, or undefined. There is no documented status transfer between outgoing and incoming teams.
6. **'We dispatch immediately' (no pre-analysis).** The provider prioritizes apparent speed of dispatch over the actual speed of resolution. Technicians arrive cold.
7. **One-line ticket closures.** Incident documentation is minimal — no root cause analysis, no photo verification, no preventive recommendations.
8. **'We can scale anywhere' without evidence.** No examples of successfully scaling with existing clients. No specific mechanism for adding capacity in new markets.
9. **General certifications only.** CompTIA certifications (A+/Network+/Server+) are the baseline. Without documented training on your specific hardware, the technician is competent but not prepared.
10. **Escalation depends on judgment, not triggers.** No automatic, time-based escalation thresholds. The technician decides when to escalate — usually too late.
11. **The provider resists data transparency.** Any reluctance to share specific performance data, escalation matrices, or documentation samples is a warning sign.
12. **'We'll figure it out.'** The most dangerous phrase in FLM procurement. It means the provider hasn't thought through how they'll serve your environment and is planning to learn on your infrastructure.

SECTION 9 — EVALUATION SCORECARD

Printable Scorecard — 50 Points Maximum

Rate each question from 1 (poor/no data) to 5 (excellent/fully documented). Compare up to two providers side by side.

#	Question	Provider A	Provider B	Notes
1	Average resolution time (auditable data)	___ /5	___ /5	
2	SLA breach rate — last 12 months	___ /5	___ /5	
3	Dedicated vs. shared technicians	___ /5	___ /5	
4	Team proximity to my sites	___ /5	___ /5	
5	Shift transition protocol	___ /5	___ /5	
6	Pre-analysis before dispatch	___ /5	___ /5	
7	Incident documentation quality	___ /5	___ /5	
8	Scalability without degradation	___ /5	___ /5	
9	Client-specific hardware training	___ /5	___ /5	
10	Escalation process (automatic triggers)	___ /5	___ /5	
	TOTAL SCORE	___ /50	___ /50	

Score Range	Assessment	Recommendation
40–50	Strong provider	Proceed to contract negotiation with confidence.
30–39	Adequate with gaps	Negotiate specific improvements in weak areas before signing.
20–29	Significant concerns	Consider alternative providers or require a remediation plan.
Below 20	Not recommended	The provider is not equipped to meet hyperscale FLM requirements.

SECTION 10 — HOW HYPE TELECOM ANSWERS

Auditable Data — Not Marketing Claims

> We built this guide because we're confident in our answers. Here's how Hype Telecom responds to each of the 10 questions with verifiable operational data.

Q1: Resolution time

Sub-4-hour SLA targets consistently met across active hyperscale FLM engagements. Data available by priority tier, by site, by month — auditable operational records, not marketing figures.

Q2: SLA breach rate

Zero SLA breaches in our active hyperscale FLM engagements over the past 12 months. Documented and auditable by priority tier and site.

Q3: Dedicated technicians

Named, dedicated technicians for engagements meeting minimum volume thresholds; pool-based qualified technicians for short-duration projects. Cross-trained backups familiar with the client environment. Certified to CompTIA A+/Network+/Server+ and relevant OEM standards (Corning, Ciena, etc.). BICSI certification on our 2026 roadmap.

Q4: Team proximity

Pre-positioned teams across 12 U.S. states with 28+ active field engineers and an extended bench of 70+ qualified technicians. Plus deep local presence across Brazil. Mobilization windows depend on geography; pre-positioned coverage across core U.S. states reduces typical mobilization.

Q5: Shift transitions

Documented shift handoff protocol with complete status transfer: open tickets, pending analyses, scheduled maintenance, emerging patterns. Shift overlap ensures zero coverage gaps. 24/7/365 operations.

Q6: Pre-analysis

Yes. Our coordination team performs structured pre-analysis before dispatch. The field technician receives a briefing — including preliminary diagnosis, recommended tools/parts, and facility-specific access instructions — before departing.

Q7: Documentation

Structured incident reports with timestamps, pre-analysis findings, root cause, resolution steps, parts replaced (with serial numbers), verification results, resolution time, and preventive recommendations. Photo documentation at every milestone.

Q8: Scalability

Our distributed model is designed for scale. Adding sites in markets where we have existing local presence takes days, not months. We've successfully scaled multi-site, cross-border deployments across the Americas.

Q9: Training

Client-specific onboarding program for every new engagement. Technicians trained on your exact hardware configurations, facility layout, monitoring tools, and procedures. Ongoing updates when configurations change. OSHA 10/30, CompTIA A+/Network+/Server+, and relevant OEM certifications. BICSI on 2026 roadmap.

Q10: Escalation

Documented, time-based escalation triggers calibrated to the SLA window of each engagement. Pre-established OEM vendor relationships for every equipment type we support. Systematic, not judgment-based.

SECTION 11 — NEXT STEPS

Get a Free FLM Assessment

This guide is designed to help you make a better-informed decision — whether you choose Hype Telecom or another provider. The questions work regardless of who you're evaluating.

If you'd like to see how Hype Telecom performs against these criteria for your specific environment, we're ready for the conversation.

WHAT THE ASSESSMENT INCLUDES

- → 15-minute call with our operations team.
- → Discussion of your current FLM setup, site locations, and SLA requirements.
- → Customized assessment showing how our model maps to your specific environment.
- → No obligation. No sales pressure. Just operational data.

ADDITIONAL FREE RESOURCES

- → Deployment Timeline Template — 5 Excel worksheets, 41 formulas, Gantt, benchmarks and 30-item checklist
  hypetelecom.net/data-center-deployment-timeline-template/
- → Blog — data center deployment guides, FLM best practices, LATAM expansion resources
  hypetelecom.net/blog/
- → FLM Case Study — Zero SLA breaches, OTDR validation, cross-border governance
  linkedin.com/posts/activity-7459713822623363073-X85n
- → 5 Hidden Costs of Slow Data Center Deployment — formulas, benchmarks, $1.8M–$7M scenario
  hypetelecom.net/data-center-deployment-costs-hidden-risks/

Hype Telecom LLC

2281 Lee Rd, Suite 201, Winter Park, FL 32789

+1 (321) 947-2184 | hypetelecom@hypetelecom.net | hypetelecom.net

linkedin.com/company/hype-telecom-llc