

Procurement Checklist

# THE SOFTWARE OBLIGATION BEHIND EVERY EDGE DEVICE

---

A pre-purchase checklist for IT, OT,  
and procurement teams

---

May 26

## **Before You Sign the PO: Software Lifecycle Questions for Industrial Edge Hardware**

For plant IT managers, OT leads, and procurement teams evaluating industrial edge devices.

The hardware decision is visible. The software management obligation that comes with it usually isn't. Then, months after deployment, the fleet has drifted and no one can say with confidence which software version is running on which device.

These questions surface that obligation at the point of purchase, when you can still do something about it.

## Deployment & Day-One Readiness

What state is the device in when it arrives, and how much work is required before it is manageable?

- Can the device run containerized workloads natively**  
Or is the runtime something you have to install, qualify, and maintain yourself?
- Does the device ship with a management agent pre-installed?**  
Or is that a post-delivery configuration step that requires a specialist?
- What does first-time setup look like for a non-specialist on-site?**  
How long does it take? Does it require inbound network access?
- Does remote management require inbound firewall rules to be opened?**  
Or does the device communicate outbound-only, within existing OT security architecture?
- Is the software stack open to third-party applications?**  
Or are you locked to the hardware vendor's ecosystem from the point of installation?

## AI Workload Readiness

If the device will run AI inference (image classification, predictive maintenance, anomaly detection, computer vision QA), the questions about workload support are different from general software questions. Most modern edge hardware is being specified with these workloads in mind.

- Does the device support GPU or AI-accelerator hardware?**  
NVIDIA, Hailo, Intel, AMD, or proprietary. Confirm the driver stack is supported by the OEM, not left for you to integrate.
- Can container workloads access the GPU or accelerator directly?**  
GPU passthrough into containers is non-trivial. Confirm it works in the field, not just in a vendor demo.
- How are AI models updated independently of the application?**  
Updating a model should not require re-flashing the device or redeploying the full container.
- Does the device have sufficient memory and storage for production AI workloads?**  
Inference at the edge is memory-hungry. Specifications that look fine on paper may not hold up under load.
- Is the AI workload lifecycle managed centrally across the fleet?**  
Or is each model update a per-device configuration task?

## Fleet Management at Scale

Managing one device is not the same as managing 30, and managing 30 devices is not the same as managing a fleet of 30,000.

- Can you push a software update to all devices simultaneously from one place?**  
Or does each device require a separate intervention?
- Can you roll back an update fleet-wide if something goes wrong?**  
How long does rollback take, and what is the failure mode?
- How do you ensure every device in the installed base runs the same software version?**  
What happens when devices are deployed at different times?
- Does the fleet management approach work across mixed hardware vendors?**  
Or are you locked into one vendor's tooling once you commit?
- If you add twenty more devices next year, does your management approach scale?**  
Or does each new device become a new configuration project?
- Who is responsible for software management on these devices once deployed?**  
IT, OT, the OEM, the integrator? And is that defined in the contract?

## Security & Compliance

Compliance frameworks are pushing operators to demonstrate they can patch and account for connected infrastructure within defined timeframes. The relevant standard varies by sector: NIS2 for industrial and critical infrastructure, FIPS 140-3 for defense and federal, HIPAA for healthcare, PCI for retail. The underlying question is the same. Can you?

**Can you push a security patch to every device in your fleet within 48 hours?**

Not in principle, in practice, with the infrastructure you will actually have.

**Can you produce an accurate software version inventory across your fleet on request?**

This is the audit question. Know the answer before the auditor asks.

**Was the device developed under a recognized secure development framework?**


IEC 62443 for industrial OT, FIPS 140-3 for defense and federal, equivalent standards in healthcare and retail.

**How does the vendor deliver security updates, and how quickly?**

What is their disclosed vulnerability response process?

**Does the device support role-based access control for software management?**

Can different teams or sites have different levels of access?



## Vendor Relationship & Total Cost

OEM pricing on post-deployment changes is where the real cost of a locked architecture becomes visible. Ask these questions before you need the answer.

- Is the software stack on your hardware dependent on a specific cloud provider or are you agnostic?**

See: [Cloud IoT vs. Dedicated Edge Management Comparison Chart](#)

- What does it cost to add a new data point or application after installation?**

Understand the costs of change requests for measurement points, not unusual in OEM-locked architectures.

- Are you paying for data access or for the hardware?**


Some OEM models charge ongoing subscription fees for basic access to data the device is already collecting.

- What happens to your management capability if you change hardware vendors in five years?**

Is your management layer portable, or tied to this vendor's ecosystem?

Hardware that can answer all of these questions confidently at the point of purchase is a different procurement decision than hardware that can't.

The software management obligation doesn't disappear. The question is whether the infrastructure for handling it arrives with the device.





**PORTAINER.io**