



WHITEPAPER 13

Embedded Panel PC Architecture

When to Choose x86, Arm or a Modular Compute Platform

Balancing performance, power, software, security and lifecycle in industrial HMI computing



A practical decision framework for processor architecture, graphics, operating systems, AI acceleration, I/O, cybersecurity, thermal design and software-porting risk.



Created by Crystal Display Systems Limited

Technical whitepaper | August 2026

Executive summary

Industrial panel PCs must run a defined software and I/O workload for years in a constrained thermal and maintenance environment. Processor architecture affects compatibility, power, graphics, security and the availability of future replacements.

x86 remains strongest for legacy Windows and general-purpose applications; Arm offers efficient integrated HMIs; COM architectures create a controlled compute-refresh boundary. The complete software and lifecycle ecosystem decides the result.

Core recommendation

Choose x86, Arm or COM from software portability, sustained thermal performance and lifecycle ownership. A lower-power CPU does not reduce programme risk if its BSP and security support are weak.

How to use this whitepaper

- Use the requirement table to convert panel PC computing expectations into measurable pass/fail criteria.
- Use the comparison matrix to shortlist architectures without allowing one headline specification to dominate.
- Apply the weighted scorecard using the actual operating environment, duty cycle and commercial lifecycle.
- Validate the preferred approach in the final mechanical, optical, electrical and software implementation.

Contents

1	Define the workload and software lifecycle before selecting compute architecture
2	Compute architectures for embedded panel PCs
3	Comparative decision matrix
4	Engineering trade-offs that change the answer
5	Application-led recommendations
6	Qualification and sourcing checklist
7	Beyond fixed CPU architecture choices
8	Working with Crystal Display Systems

Key decision principles

- Select architecture from the complete software and I/O workload, not benchmark peak performance alone.
- Budget sustained fanless performance at the highest ambient, including display and power-converter heat.
- Treat OS/BSP maintenance, driver ownership and security updates as lifecycle costs.
- Use modular compute only when mechanical, firmware and supply interfaces are genuinely standardised and controlled.
- Validate boot, storage endurance, network recovery and worst-case application load as part of the HMI system.

1. Define the workload and software lifecycle before selecting compute architecture

A useful panel PC requirement states concurrent applications, I/O, graphics/AI load, OS and drivers, boot/update, power, ambient temperature, security, service life and upgrade expectations. These inputs determine architecture more reliably than CPU benchmark tables.

Requirement domain	Questions to resolve
Workload	HMI framework, browser, database, PLC/fieldbus, vision, AI, video decode, graphics resolution, real-time tasks and peak concurrency.
Software estate	Operating system, drivers, legacy binaries, container/runtime, update method, licensing, certification and development toolchain.
I/O and expansion	Ethernet, serial, CAN, GPIO, USB, camera, storage, wireless, PCIe/M.2 and deterministic fieldbus requirements.
Power and thermal	Supply voltage, average/peak power, fanless limit, enclosure temperature, compute throttling and display heat interaction.
Security and management	Secure boot, TPM/root of trust, device identity, patching, remote management, fleet monitoring and decommissioning.
Lifecycle	Processor roadmap, module availability, OS support horizon, memory/storage sourcing, software porting and field replacement strategy.

A practical weighting model

Assign each criterion a weight from 1 (low importance) to 5 (critical). Score each candidate from 1 (poor fit) to 5 (strong fit), then multiply score by weight. The arithmetic does not replace engineering judgement, but it exposes where the compute architecture depends on assumptions, incomplete test evidence or optimistic supplier data.

Do not average away a fatal constraint

A candidate that fails a non-negotiable requirement for panel PC computing should be removed before weighted scoring. Typical fatal constraints include unsafe operation, an unachievable environmental limit, incompatible interfaces, unacceptable field reliability or no credible supply route.

Recommended process

- Set pass/fail gates for the mandatory panel PC computing requirements before comparing desirable features.
- Agree criterion weights with electronics, mechanical, software, product, service and procurement stakeholders.
- Score only from controlled data, supplier evidence or representative testing; mark unknowns explicitly.
- Run sensitivity checks by changing the highest weights and identifying whether the preferred architecture changes.
- Document assumptions, test conditions and revision status so the decision can be revisited after a component or firmware change.

2. Compute architectures for embedded panel PCs

2.1 x86 panel PC

x86 platforms provide broad Windows and Linux compatibility, mature graphics, virtualisation and extensive industrial I/O ecosystems. Performance spans low-power processors to high-end Core-class devices. Power and thermal density can be higher than Arm for equivalent embedded tasks.

- Strengths: legacy software compatibility, mature drivers/tools, strong desktop-class graphics and wide ecosystem support.
- Limitations: higher power in many configurations, shorter consumer-derived CPU cadence and licensing/patching complexity.
- Best fit: Windows applications, legacy x86 binaries, rich HMIs, database workloads and general industrial computing.

2.2 Arm-based HMI platform

Arm SoCs integrate CPU, GPU, video, camera and often NPU functions at low power. They are well suited to Android, embedded Linux and purpose-built HMIs, but depend heavily on the board-support package, kernel/device tree and vendor update commitment.

- Strengths: low power, high integration, fast boot and efficient multimedia/AI acceleration.
- Limitations: software porting, fragmented BSP support, architecture-specific drivers and possible dependence on one module vendor.
- Best fit: dedicated HMIs, Android/Linux appliances, battery or sealed systems and cost-sensitive volume products.

BSP ownership

Define who maintains the kernel, bootloader, device tree, drivers, security patches and recovery image for the full product life. Hardware availability without a maintained BSP is not lifecycle continuity.

2.3 Computer-on-module (COM) architecture

COM Express, SMARC, Qseven and proprietary modules place processor, memory and core I/O on a replaceable module over a custom carrier. They can separate application I/O and mechanics from CPU cadence, but connector standards do not guarantee software or thermal drop-in compatibility.

- Strengths: upgrade path, carrier reuse, scalable performance and reduced custom high-speed design.
- Limitations: module cost, vendor-specific pin use, BIOS/BSP differences and thermal/mechanical requalification for upgrades.
- Best fit: long-life products with stable custom I/O and periodic compute refresh.

2.4 Fully integrated custom SBC

A custom single-board design can optimise cost, size, I/O and power at volume. It requires high-speed design, regulatory work, software enablement and lifecycle ownership, and is less attractive at low volume.

2.5 Distributed or split HMI architecture

Separating the display head from a boxed PC or edge controller reduces heat and service complexity at the front panel. It adds network/video links, latency, cybersecurity and two enclosure lifecycles but can simplify upgrades and harsh-environment deployment.

3. Comparative decision matrix

Criterion	x86 integrated	Arm integrated	COM + carrier	Split HMI / custom SBC
Windows / legacy software	Excellent	Limited	Strong	Strong
Power efficiency	Moderate	Excellent	Strong	Strong
Graphics / AI integration	Strong	Excellent	Strong	Excellent
Industrial I/O breadth	Excellent	Strong	Excellent	Excellent
Fanless thermal margin	Moderate	Excellent	Strong	Strong
BSP / driver simplicity	Excellent	Limited	Strong	Moderate
Compute upgrade path	Moderate	Moderate	Excellent	Excellent
Customisation	Strong	Excellent	Excellent	Excellent
Low-volume economics	Excellent	Strong	Strong	Limited
Lifecycle control	Strong	Moderate	Excellent	Excellent

Indicative only. Actual performance depends on the complete system architecture, supplier implementation, environmental conditions and validation method.

Example weighted scorecard

Criterion	Weight	x86 integrated	Arm integrated	COM + carrier
Existing Windows application	5	5	2	5
25 W fanless limit	5	3	5	4
Custom CAN / serial I/O	4	4	4	5
Ten-year platform life	5	4	3	5
Future CPU upgrade	4	3	3	5
Development schedule	4	5	3	4
Unit cost at 500/year	3	4	5	3
Weighted total		120	105	135

For this Windows-based industrial machine, a COM plus carrier architecture scores highest because it preserves x86 compatibility while creating a controlled upgrade path and custom I/O boundary. Integrated x86 remains fastest to deploy; Arm is thermally strong but loses heavily on software porting.

Reading the result

- x86 is the low-risk application platform when existing binaries and driver ecosystems dominate.
- Arm should be selected only with a funded BSP/security-maintenance plan and verified application port.
- A COM architecture earns its premium when the carrier remains stable across at least one credible module refresh.

4. Engineering trade-offs that change the answer

4.1 Peak benchmark is not sustained fanless performance

A processor may meet workload briefly and then throttle in a sealed panel PC. Evaluate sustained clocks, GPU/NPU load and enclosure temperature with the display at maximum heat.

4.2 Software porting includes more than recompilation

Drivers, middleware, installers, licensing, endianness assumptions, SIMD code, device access and update tools can create hidden x86-to-Arm cost.

4.3 Integrated accelerators need supported software

An NPU or GPU is useful only when toolchains, operators, model formats and drivers are maintained for the chosen OS image.

4.4 Storage endurance can limit the platform before CPU life

Logging, databases, swap and updates can consume eMMC/SSD write life. Model write amplification, power-loss protection and field replacement.

4.5 COM standards do not remove qualification

New modules can change power transients, BIOS, thermal solution, display outputs, I/O mapping and EMC. Define an upgrade qualification matrix.

4.6 Security maintenance is architecture specific

Secure boot keys, TPM/root of trust, bootloader, OS and firmware update routes must remain supported. A long-available board with unpatchable software is not a long-life platform.

5. Application-led recommendations

Windows machine controller

Use x86 or x86 COM where legacy applications, vendor drivers and IT management dominate. Size the thermal solution for sustained workload.

Dedicated Android/Linux HMI

Arm can reduce power and cost if the UI, BSP and update stack are controlled and no legacy binaries are required.

Vision-enabled inspection station

Compare integrated NPU, discrete GPU and x86/Arm toolchains based on actual model throughput, camera I/O and thermal load.

Harsh front panel with serviceable compute

Use a split architecture with an industrial monitor/display head and remotely mounted boxed PC where heat, access or contamination make an all-in-one undesirable.

Long-lived OEM platform with custom I/O

A COM plus carrier can protect connector and fieldbus investment while allowing planned compute generations.

6. Qualification and sourcing checklist

	Area	Verification action
☐	Workload	Benchmark representative UI, I/O, network, storage, graphics and AI tasks concurrently at hot ambient.
☐	Software	Verify OS/BSP, drivers, boot, update, recovery, licensing and patch ownership.
☐	Thermal	Measure sustained clocks, throttling, surface temperatures and display/compute interaction.
☐	I/O	Test all ports, fieldbus timing, cable loads, hot plug and electrical protection.
☐	Storage	Run power-loss, endurance, update and filesystem recovery tests.
☐	Security	Verify secure boot, key provisioning, device identity, update signing and rollback/recovery.
☐	EMC / reliability	Test worst I/O configuration, radios, load, vibration and temperature.
☐	Lifecycle	Control module/CPU, memory, storage, OS/BSP and upgrade qualification plan.

Questions to ask a display partner

- Which representative workload proves sustained performance at the maximum fanless temperature?
- Who owns OS/BSP, bootloader, drivers and security patches for the product lifetime?
- What I/O and expansion are native, bridged or dependent on optional modules?

- How are storage endurance, power loss and image recovery handled?
- What changes when a COM or processor module is upgraded, and which tests repeat?
- What lifecycle evidence covers CPU, memory, storage, OS and software toolchain together?

Minimum evidence pack for design approval

- Workload definition and hot-ambient benchmark with throttling and power data.
- Controlled hardware configuration, BIOS/BSP/OS image and driver manifest.
- Thermal report, mechanical integration and worst-case I/O/load qualification.
- Security architecture, key provisioning, signed update and recovery evidence.
- Storage endurance/power-loss test and field-service procedure.
- Lifecycle statement and documented module/architecture migration plan.

7. Beyond fixed CPU architecture choices

Panel PC platforms are increasingly heterogeneous and modular. The engineering question is how much software and lifecycle control accompanies the hardware capability.

Heterogeneous CPU-GPU-NPU SoCs

Integrated accelerators can deliver efficient edge processing, but performance portability and long-term SDK support need proof.

Chiplet and modular compute ecosystems

Future modules may separate compute, I/O and accelerators more flexibly. Signal integrity, thermal interfaces and vendor standards remain critical.

Containerised application deployment

Containers can reduce OS coupling, but hardware drivers, real-time access, graphics and update orchestration still depend on the platform.

Remote graphics and thin-client HMI

Centralising compute can simplify front-panel upgrades but adds network latency, availability and cybersecurity requirements.

RISC-V and alternative architectures

Open instruction sets may improve strategic control, but industrial software, safety, graphics and long-term module ecosystems are still developing.

Adoption gate

A new compute architecture should enter production only when the complete software stack, sustained thermal performance, security maintenance and migration ownership are credible for the service life.

8. Working with Crystal Display Systems

Crystal Display Systems Limited supplies and customises industrial panel PCs, Arm HMIs, embedded computers and complete display-compute assemblies with processor, I/O, touch, enclosure and lifecycle options.

Relevant CDS resources

Resource	Purpose
<u>Industrial Panel PC Solutions</u>	Integrated industrial HMI and panel PC solutions with customisation options.
<u>Industrial Panel PC Range</u>	Current panel PC platforms, processors, I/O and rugged configurations.
<u>Industrial Boxed PCs and Panel PCs</u>	Industrial computing platforms, customisation and long-life support.
<u>Industrial Automation Displays</u>	Display, monitor and computing solutions for automation and machine HMI.
<u>AI-Driven Embedded Innovation</u>	Embedded AI accelerators and edge-processing solutions.
<u>CDS Knowledge Centre</u>	Display guides and technical training documents.
<u>Request a Quote</u>	Submit project requirements for technical and commercial review.
<u>Contact CDS</u>	Contact details and enquiry routes.

Contact Crystal Display Systems Limited

Crystal Display Systems Limited
Address: Unit 6 M2M Park, Fort Bridgewood, Maidstone Road, Rochester, Kent, ME1 3DQ, United Kingdom
Telephone: +44 (0)1634 791 600
Email: info@crystal-display.com
www.crystal-display.com

Conclusion

Panel PC architecture is a software-lifecycle and thermal decision as much as a processor choice. x86 minimises risk for Windows and legacy applications, Arm provides efficient integration for controlled appliances, and modular compute can protect custom I/O investment. The defensible platform is validated under sustained workload and includes an owned update, security and migration plan.

Final engineering rules

- Select from real software, I/O and sustained workload rather than peak benchmarks.
- Close fanless thermal performance with the display and enclosure at worst case.
- Fund BSP, driver, OS and security maintenance for the entire deployment life.
- Use modularity only with controlled interfaces and a tested upgrade path.
- Include storage endurance, recovery and field service in platform qualification.

Disclaimer

This whitepaper provides general engineering guidance and does not replace application-specific design review, supplier specifications, regulatory assessment or formal product qualification. Performance depends on the exact components, implementation, environment and usage profile.